

Multisensor Earth Observation and Unsupervised Anomaly Analysis of Environmental Conditions Preceding the 2026 Langtang Lirung Glacier Collapse and Slope Failure

Debabrata Pruseth
Independent Research Scholar

Abstract

The increasing availability of satellite Earth observation and climate reanalysis creates opportunities for retrospective assessment of glacier hazards, but environmental conditioning must be distinguished from event prediction. We present a quality-aware multisensor framework combining Sentinel-1 ground-range-detected (GRD) synthetic-aperture radar backscatter, Sentinel-2 multispectral imagery, ERA5-Land reanalysis, time-series statistics with autocorrelation-aware inference, and unsupervised anomaly detection to examine conditions surrounding the 26 August 2026 Langtang Lirung glacier-rock slope failure in Nepal. Relative to identical calendar intervals in 1984–2025, mean temperature and accumulated positive degree days (PDD) were at the 100th empirical percentile over 1-, 3-, and 7-day antecedent windows and remained high over 14–60 days; antecedent precipitation was not comparably extreme. Long-term monthly analysis supported an increasing PDD tendency and negative trends in Sentinel-1 VV and VH backscatter, while temperature and optical snow-proxy trends weakened after autocorrelation-aware testing. A consistent Sentinel-1 ascending-orbit event pair (16 and 28 August) retained 95.6% valid glacier coverage; whole-glacier mean changes are not interpreted as a validated failure footprint. Sentinel-2 post-event valid-area coverage was only 2.3%, so quantitative optical change was rejected by design. An unsupervised monthly ensemble combining robust standardized distance, Isolation Forest, and PCA reconstruction error did not identify a threshold-exceeding 2026 month; August was unscored because the month was incomplete. The results identify unusually warm and melt-favourable environmental conditions consistent with conditioning, but do not establish a physical trigger, causal mechanism, predictive precursor, or operational warning capability.

1 Introduction

Rapid glacier and rock-slope failures in high mountains are difficult to diagnose because glacier instability can emerge from interacting thermal, hydrological, mechanical, and geometric processes operating across different spatial and temporal scales. Debris-covered Himalayan glaciers present additional observational challenges because debris cover, heterogeneous surface lowering, supraglacial features, and complex englacial hydrology can produce strongly variable glacier responses. These complexities also constrain the interpretation of individual remote-sensing and environmental variables: optical snow and ice indices may be unreliable over debris-covered surfaces, while radar observations and gridded reanalysis products characterize particular aspects of surface and environmental conditions rather than directly measuring the mechanical state of a potential failure surface.

The 26 August 2026 Langtang Lirung event provides a recent case in which multisensor observations can be evaluated under explicit quality controls. ESA reported a rapid slope failure associated with glacier collapse on Langtang Lirung and documented satellite observations before and after

the event. This study asks a deliberately limited question: what environmental information was observable in the frozen analysis pipeline, which signals were statistically unusual, and what could those data not establish?

We evaluate five research questions. **RQ1** asks whether long-term climate, optical, and radar indicators show statistically detectable change. **RQ2** asks whether environmental conditions immediately preceding the event were unusual relative to matched historical calendar periods. **RQ3** asks whether an unsupervised multivariate model identifies a distinctive pre-event environmental anomaly. **RQ4** asks what quality-controlled Sentinel-1 and Sentinel-2 observations can establish across the event window. **RQ5** asks whether the combined evidence supports environmental conditioning while remaining insufficient to establish a physical trigger, causal mechanism, or predictive precursor.

Two hypotheses organize the analysis. **H1** proposes that statistically detectable long- and short-timescale changes in melt-favourable climate and remotely sensed surface characteristics were present before the event. **H2** proposes that detectable environmental anomalies are not necessarily unique to the event and therefore are insufficient by themselves to constitute a predictive collapse precursor. H1 and H2 are not mutually exclusive: a glacier environment may be measurably conditioned while a generic anomaly detector fails to isolate a unique event signature.

The contribution is methodological rather than a claim of a new collapse-prediction algorithm. Specifically, the framework combines a quality-aware multisensor architecture, matched historical event windows, trend statistics with autocorrelation-aware inference, unsupervised AI, and explicit preservation of negative or insufficient evidence. This separation is important because retrospective environmental explanation and prospective prediction impose different evidentiary requirements.

2 Related work

2.1 Langtang and debris-covered glacier change

Previous work establishes that the Langtang Himal is characterized by spatially heterogeneous glacier change rather than a uniform response. Ragetti et al. documented heterogeneous and accelerating thinning across glaciers of the upper Langtang catchment. Nuimura et al. reconstructed sustained downwasting of the debris-covered ablation zone of Lirung Glacier and reported acceleration after 2000. Miles et al. demonstrated complex supraglacial–englacial hydrological connectivity on debris-covered Lirung Glacier. Together, these studies motivate caution in treating any single spectral or climate proxy as a complete representation of glacier state.

2.2 Catastrophic glacier instability and multisensor observation

Catastrophic glacier detachments and glacier-related mass movements can involve interacting climatic, thermal, hydrological, geometric, and dynamic controls. The 2016 Aru glacier detachments are an important example of abrupt failure following unusual glacier behaviour; other catastrophic glacier events similarly demonstrate the importance of evaluating competing physical mechanisms rather than inferring a trigger from environmental coincidence alone. Multisensor Earth observation is therefore attractive because optical, radar, reanalysis, and inventory data observe different parts of the system. Sentinel-1 supplies all-weather C-band SAR observations, Sentinel-2 supplies multispectral optical observations, ERA5-Land supplies a long-duration gridded land-surface record, and GLIMS supplies standardized glacier geometry.

2.3 Research gap and analytical position

The gap addressed here is not the absence of another classifier. It is the interface among heterogeneous sensor quality, temporal scale, statistical dependence, and AI interpretation in a single-event retrospective analysis. A short-duration thermal extreme can be physically relevant as environmental context without appearing as an extreme monthly multivariate state. Likewise, a post-event satellite acquisition can exist but be quantitatively unusable because cloud or valid-area coverage is inadequate. The framework therefore treats data quality as part of the inference rather than a preprocessing detail. An environmental anomaly is defined as rarity relative to the historical feature distribution; it is not a collapse probability and is not calibrated as a hazard score.

3 Study area and data

The analysis pipeline resolves Lirung Glacier using GLIMS identifier G085544E28246N, with legacy identifier G085547E28252N retained as a fallback. The glacier region of interest is separated from a 15-km contextual buffer, a 25-km downstream visual context, a HydroSHEDS level-12 basin, and a derived 1–5-km stable-control candidate annulus. Suspected source-zone and runout geometries are intentionally unset and are not treated as validated event footprints.

Table 1 summarizes the evidence base and the principal quality constraint attached to each source. Sentinel-2 surface reflectance provides multispectral optical observations; Sentinel-1 GRD provides C-band radar backscatter; ERA5-Land provides daily land-surface reanalysis fields; and GLIMS provides glacier inventory geometry. The climate reference period is 1984–2025, excluding the event year from climatological and anomaly-model training baselines. Sentinel-2 monthly processing begins in 2017, while Sentinel-1 processing begins in 2015. The event date used for event-relative analysis is 26 August 2026, documented separately in provenance rather than silently embedded as an assumed training label.

Table 1: Data sources, analytical roles, and principal quality controls.

Source	Role	Principal products	Principal QA constraint
GLIMS	Glacier ROI	Outline and identifier	Geometry provenance retained
Sentinel-2	Optical state	NDSI; snow/clean-ice proxy	Cloud/SCL masking and valid-area gate
Sentinel-1 GRD	Radar state	VV/VH levels and changes	Pass/orbit geometry, incidence angle, coverage
ERA5-Land	Climate context	Temperature, precipitation, snowfall, runoff, PDD	Temporal completeness and matched-period checks
Terrain/hydrology	Context/screening	Basin, slope/aspect and geometry masks	Not treated as validated source/runout geometry

4 Methods

4.1 Quality-aware processing

The implemented sequence is Raw Data → Quality Control → Feature Engineering → Statistical Analysis → Unsupervised Analysis → Evidence Synthesis. Records retain coverage, observed and

expected counts, completeness, quality status, and reasons. This provenance allows a value to be distinguished from the evidentiary status of that value. For example, a nominal optical index computed from a very small clear-sky fraction is not promoted to an event-change estimate merely because the arithmetic can be performed. The analysis code uses `GOOD`, `CAUTION`, and `INSUFFICIENT` states; incomplete or low-coverage periods are not silently imputed into the principal analyses.

4.2 Sentinel-2 optical proxy

Sentinel-2 digital numbers are scaled to reflectance as $\rho = DN/10,000$. The normalized-difference snow index is

$$\text{NDSI} = \frac{\rho_{\text{green}} - \rho_{\text{SWIR}}}{\rho_{\text{green}} + \rho_{\text{SWIR}}}, \quad (1)$$

using bands B3 and B11. The default snow/clean-ice proxy requires $\text{NDSI} \geq 0.40$ and green reflectance ≥ 0.10 ; scene cloudiness is screened and pixel-level cloud/scene-classification masks are applied. Sensitivity outputs test NDSI thresholds 0.30, 0.40, and 0.50 and alternative coverage thresholds. NDSI is used only as a spectral proxy and not as glacier area or mass balance.

4.3 Climate features and matched antecedent windows

ERA5-Land temperatures are converted from kelvin to degrees Celsius; precipitation, snowfall water equivalent, and runoff are converted to millimetres. Daily PDD is $\max(T_d, 0)$ and accumulated PDD is summed over each interval. Freeze–thaw days satisfy $T_{\min} \leq 0 < T_{\max}$. For an antecedent window of d days, the event-year interval $[t_e - d, t_e)$ is compared with the identical month/day interval in each complete historical year from 1984 through 2025. The analysis evaluates $d \in \{1, 3, 7, 14, 30, 60, 90\}$, admitting only complete historical intervals. This calendar matching avoids comparing late-monsoon conditions with climatologically dissimilar seasons. Empirical percentiles are calculated against the complete historical matched samples. In addition to PDD, the daily climate engineering records mean/minimum/maximum temperature, precipitation, snowfall, snow depth, solar radiation and runoff, and identifies freeze–thaw conditions. Rolling precipitation summaries require complete windows rather than silently accepting partial accumulation.

4.4 Sentinel-1 SAR

The analysis uses Sentinel-1 GRD backscatter rather than interferometric displacement. Candidate pass/orbit geometries are evaluated using usable-month count, median valid-area fraction, temporal coverage, event-window availability, and incidence-angle consistency. The selected event pair uses a consistent ascending relative orbit. Candidate observations are restricted to incidence angles of approximately 30–45 degrees and are screened for valid glacier-area coverage; a candidate track requires sufficient usable monthly history. Track selection is therefore a QA decision, not selection of the orbit that produces the largest event change. Event differences are $\Delta VV = VV_{\text{post}} - VV_{\text{pre}}$ and $\Delta VH = VH_{\text{post}} - VH_{\text{pre}}$. Because whole-glacier means can dilute localized changes, these metrics are not interpreted as a spatial failure footprint.

4.5 Trend and association statistics

For each monthly feature, elapsed time in years is used in an OLS trend with heteroskedasticity- and autocorrelation-consistent covariance (maximum lag 12, or a shorter data-dependent limit). Robust slope is estimated using the Theil–Sen estimator. Trend significance is also evaluated with the

original Mann–Kendall test and the Hamed–Rao autocorrelation-aware modification. False-discovery-rate adjustment is applied to families of lagged-correlation tests using the Benjamini–Hochberg procedure. Snow–climate association is estimated by OLS with HAC covariance (maximum lag 12); coefficients are interpreted as associations, not causal effects. The use of both slope estimators and autocorrelation-aware significance tests is deliberate: an apparent monotonic tendency in a serially dependent environmental series should not be promoted solely because an unmodified trend test is significant.

4.6 Unsupervised environmental anomaly model

All model fitting is restricted to complete observations dated before 1 January 2026. The primary feature strategy uses temperature anomaly, seasonally adjusted PDD, adjusted 7-day precipitation, adjusted antecedent precipitation index, and adjusted VV/VH monthly changes. Features are centered on training medians and scaled by 1.4826 MAD. If MAD is zero, the implementation falls back to IQR/1.349 and then population standard deviation; unusable constant features are rejected. The resulting first detector is a root-mean-square robust standardized distance, not a covariance-aware Mahalanobis distance. Three unsupervised scores are then computed: (i) root-mean-square standardized distance from the training median, (ii) Isolation Forest with 500 trees, and (iii) PCA reconstruction mean-squared error with enough components to explain at least 90% of training variance. Each raw detector score is mapped to its empirical percentile in the training distribution and the three percentiles are averaged. The reference flag is the 95th percentile of the training ensemble score. This quantity measures rarity under the selected representation; it is not $P(\text{collapse})$.

Exploratory environmental regimes are fitted separately with K-means for $K = 2, \dots, 6$, selecting K by maximum silhouette score. These clusters are descriptive state groupings and not hazard classes.

5 Results

5.1 Long-term indicators

Table 2 summarizes the frozen trend output. PDD retains evidence of an increasing trend under both OLS-HAC ($p = 0.0247$) and modified Mann–Kendall ($p = 0.0139$). Temperature has a positive slope but is not significant after HAC or modified Mann–Kendall correction. Snow-fraction and NDSI anomalies show negative point estimates but do not meet the 0.05 threshold under the autocorrelation-aware test. VV and VH anomalies show statistically strong negative trends; these changes are radar-observed surface-state trends and are not uniquely attributable to mechanical instability. Figure 1 shows the seasonally adjusted time series.

5.2 Antecedent climate conditions

Matched-period results (Table 3) show that temperature and PDD were at the 100th empirical percentile for 1-, 3-, and 7-day windows and remained at or above the 92.9th percentile through 60 days. The 90-day thermal percentile fell to 71.4%. In contrast, precipitation percentiles ranged from 16.7% to 71.4% over 1–90 days and therefore were not comparably extreme. Snowfall was unusually low across most windows. Figure 2 places the event-year observations against earlier same-month conditions.

Table 2: Long-term trend diagnostics from the frozen analysis output.

Feature	OLS slope yr^{-1}	OLS-HAC p	Sen slope yr^{-1}	Modified MK p
Temperature anomaly	0.04575	0.2810	0.06352	0.1337
Precipitation anomaly	0.03104	0.1465	0.02787	0.3102
PDD adjusted	0.8394	0.0247	0.3594	0.0139
Snow-fraction anomaly	-0.01404	0.0820	-0.01332	0.0658
NDSI anomaly	-0.01260	0.0995	-0.01297	0.0611
VV anomaly	-0.07182	9.74×10^{-5}	-0.05793	7.06×10^{-6}
VH anomaly	-0.1496	2.54×10^{-8}	-0.1299	2.50×10^{-7}

Table 3: Matched historical percentiles (%) for event-exclusive antecedent windows ending 26 August 2026. Historical sample size is 42 years for each complete interval.

Days	Temperature	PDD	Precipitation	Snowfall
1	100.0	100.0	19.0	11.9
3	100.0	100.0	16.7	2.4
7	100.0	100.0	42.9	0.0
14	97.6	97.6	59.5	14.3
30	92.9	92.9	69.0	2.4
60	97.6	97.6	71.4	2.4
90	71.4	71.4	52.4	4.8

5.3 Sentinel-1 event-window evidence

The QA procedure selected ascending relative orbit 85: 289 scenes, 116 usable months, median valid-area fraction 95.6%, and a post-event acquisition. Descending orbit 19 had high historical coverage but no post-event scene within the configured 60-day event window; descending orbit 121 had only nine usable months and effectively zero median glacier coverage. The selected 16 and 28 August 2026 pair each retained 95.6% valid glacier coverage. Whole-glacier mean $\Delta VV = -0.0245$ dB and $\Delta VH = +0.1495$ dB. These whole-glacier mean changes do not establish a spatially validated failure footprint. The longer radar record is shown in Figure 3, while Table 4 records the candidate-track decision.

Table 4: Sentinel-1 candidate-track QA from frozen analysis outputs.

Pass	Orbit	Scenes	Usable months	Median valid area (%)	Selected
Ascending	85	289	116	95.6	Yes
Descending	19	326	120	98.0	No
Descending	121	301	9	0.0	No

5.4 Sentinel-2 event-window quality

The pre-event Sentinel-2 composite had 84.1% valid-area coverage, whereas the post-event composite had only 2.3%. Because the pre/post pair fails the configured coverage gate, the analysis sets `quantitative_change_estimated=False` and suppresses NDSI and snow-proxy change estimates. This negative result is retained rather than replaced by a visually plausible but quantitatively unsupported optical change.

5.5 Snow–climate association

The HAC regression of snow-fraction anomaly on temperature anomaly, adjusted PDD, and precipitation anomaly yielded a temperature coefficient of -0.05339 ($p = 3.17 \times 10^{-6}$), precipitation coefficient $+0.02448$ ($p = 0.0142$), and PDD coefficient -0.000717 ($p = 0.373$). Thus temperature and precipitation are associated with the optical snow proxy in this specification, while PDD does not add an independently significant coefficient after those covariates are included. The regression is descriptive and does not identify causal effects. Table 5 reports the full fitted coefficients.

Table 5: HAC snow–climate regression.

Term	Coefficient	HAC SE	p -value
Intercept	0.01452	0.01897	0.444
Temperature anomaly	-0.05339	0.01146	3.17×10^{-6}
Adjusted PDD	-0.00072	0.00081	0.373
Precipitation anomaly	0.02448	0.00998	0.0142

5.6 Unsupervised anomaly analysis

The primary climate+SAR-change model contains 106 complete pre-2026 training months. PCA retained four components explaining 96.29% of training variance. No scored 2026 month exceeded the training 95th-percentile ensemble threshold; the highest 2026 score was approximately 0.730 in June. August 2026 was not scored because the incomplete month failed the completeness requirements. Figure 4 shows the ensemble and component scores. K-means selected $K = 2$ with silhouette score approximately 0.355; the PCA regime visualization is relegated to the supplement because it is exploratory rather than central to the event inference.

5.7 Integrated evidence

Figure 5 juxtaposes temperature anomaly, PDD, optical snow proxy, and the monthly ensemble anomaly score. The strongest event-centered evidence is the multi-window thermal signal in Table 3; the monthly ensemble does not produce a unique precursor. These findings are not necessarily contradictory: short-duration extremes may be attenuated by monthly aggregation, while multivariate scoring may reduce the influence of an extreme signal in one subsystem when other features remain less unusual.

5.8 Research-question synthesis

Table 6 separates what each analysis supports from stronger claims that the analysis cannot make. RQ1 receives variable-dependent support: PDD and SAR trends are supported across the implemented trend tests, while temperature and optical trends weaken under autocorrelation-aware testing. RQ2 is supported by unusually high thermal conditions relative to matched historical windows, with the strongest short-window values reaching the 100th empirical percentile. RQ3 yields a negative result: no scored 2026 monthly state crosses the configured ensemble threshold. RQ4 is partial because the radar pair is usable but optical post-event coverage is insufficient and no validated spatial SAR footprint is available. RQ5 therefore supports environmental conditioning only.

H1 consequently receives variable-dependent support: short-window thermal and PDD conditions were unusually high relative to matched historical windows, and long-term PDD change and Sentinel-1 backscatter trends are supported, whereas long-term temperature and optical trends are not supported under the autocorrelation-aware tests. H2 is consistent with the implemented anomaly-model result, but this single-event retrospective analysis does not constitute predictive validation.

Table 6: Research-question evidence matrix.

RQ	Evidence	Outcome	Permitted conclusion
RQ1	Trend summary and multisensor series	Variable-dependent support	PDD and SAR trends supported across the implemented trend tests; temperature and optical trends weaker under autocorrelation-aware testing
RQ2	Matched 1–90 day windows	Supported	Unusually high thermal conditions, strongest over short windows; precipitation not comparably extreme
RQ3	Climate+SAR anomaly ensemble	No threshold-exceeding pre-event state identified	No threshold-exceeding scored 2026 monthly state
RQ4	Sentinel-1/Sentinel-2 event QA	Partial	Valid SAR pair; optical quantitative change rejected; no validated spatial failure footprint
RQ5	Integrated evidence	Conditioning only	Consistent with environmental conditioning, not causation or prediction

6 Discussion

6.1 Conditioning is not triggering

The event was preceded by unusually warm and melt-favourable conditions across several day-to-week windows, whereas precipitation was not similarly exceptional. This is consistent with a cautious interpretation of thermal and melt-favourable environmental conditioning. It does not identify why failure occurred at a particular time, because the pipeline does not directly observe fracture evolution, stress state, basal conditions, ice thickness, permafrost, or subsurface hydrology. Prior catastrophic glacier failures similarly motivate multi-process rather than single-trigger interpretations.

6.2 Why the negative AI result matters

A threshold-exceeding monthly multivariate anomaly is neither necessary nor sufficient for a short-duration hazardous transition. The model is intentionally unsupervised and is trained without event labels; its score therefore quantifies rarity, not hazard probability. The absence of a 2026 threshold exceedance should be read as a model result, not a failure to obtain a desired positive signal. It also highlights the importance of temporal representation, missingness, and physical observability; the present analysis does not test whether increasing algorithmic complexity would improve event discrimination.

6.3 Temporal scale and interpretation of the anomaly ensemble

The day-to-week matched analysis and the monthly anomaly ensemble answer different questions. The former asks how unusual a specific event-relative interval was compared with the same calendar interval in prior years. The latter asks whether a complete monthly vector spanning climate and SAR-change variables was rare relative to pre-2026 training months. Monthly aggregation can dilute a brief thermal excursion; moreover, percentile averaging across detectors and features can make a single-subsystem extreme less exceptional when other dimensions remain ordinary. Consequently, the thermal result and negative monthly AI result are complementary rather than contradictory.

The ensemble design also guards against over-interpreting a single unsupervised algorithm. Robust standardized distance emphasizes overall standardized departure, Isolation Forest captures isolation under recursive partitioning, and PCA reconstruction error emphasizes departure from the dominant covariance subspace. Mapping each to a training percentile places the three detectors on a common empirical scale before averaging. Nevertheless, the ensemble remains representation-dependent: feature selection, temporal aggregation, missingness, and training period define what “rare” means.

6.4 Radar interpretation and optical insufficiency

The long-term VV/VH trends demonstrate changing radar-observed surface characteristics but cannot uniquely diagnose mechanical instability. The valid 16–28 August pair motivates spatially explicit follow-up, but the current analysis output reports whole-glacier means and lacks a validated source-zone mask. Conversely, the optical event pair fails its quality gate. These two cases illustrate the value of sensor-specific evidence rules: “available” is not equivalent to “interpretable.”

6.5 Quality-aware scientific AI

A central result of this study is methodological restraint. The pipeline rejects the low-coverage Sentinel-2 post-event comparison, leaves incomplete August 2026 unscored in the monthly anomaly model, and permits the anomaly ensemble to return no unique precursor. These are not missing successes to be repaired after inspection; they are reproducibility properties. In hazard-oriented AI, preserving insufficient evidence is particularly important because post hoc relaxation of quality gates can create an appearance of predictive skill that was not available under the declared analysis.

6.6 Implications for future monitoring

A more complete monitoring architecture would link multiple physical layers: climate forcing → cryospheric state → terrain/slope state → physical deformation or change → event and runout. The present analysis primarily informs the first two layers and supplies partial radar surface-state evidence. It does not observe enough mechanical variables to close the chain. Future AI should therefore be evaluated as one component of a physical and quality-aware monitoring system rather than as a substitute for deformation measurements, field observations, or geotechnical interpretation.

7 Limitations

This is a single-event retrospective case study, so predictive sensitivity, specificity, calibration, and false-positive rates cannot be estimated. ERA5-Land is gridded reanalysis rather than an in-situ failure-site record. Sentinel-2 NDSI is incomplete for debris-covered glacier characterization. Sentinel-1 GRD backscatter is not InSAR displacement and does not directly measure deformation.

Validated source-zone and runout geometries are absent. The principal anomaly model operates at monthly resolution, and August 2026 is intentionally unscored because the month is incomplete. Consequently, the environmental anomaly score must not be interpreted as hazard probability.

8 Future work

Priority extensions are spatial Sentinel-1 change analysis using independently validated source and runout geometries; deformation/velocity products where technically valid; higher-frequency rolling anomaly representations; incorporation of glacier velocity, ice thickness, fracture indicators, permafrost and hydrological state; and systematic evaluation across multiple events and non-events. Multi-event validation is essential before estimating sensitivity, specificity, false-positive rates, or operational thresholds. Cross-glacier testing would also be required to determine whether the selected features represent generalizable environmental states or only the local Langtang record.

A future quality-aware glacier hazard digital twin could use the evidence architecture as a provenance layer, but an operational warning claim would require prospective validation and substantially stronger physical observability. The present study should therefore be understood as a retrospective environmental-analysis baseline.

9 Conclusions

The frozen analysis framework integrates ERA5-Land, Sentinel-2, Sentinel-1, trend analysis with autocorrelation-aware inference, and unsupervised anomaly detection under explicit quality controls. Matched historical windows show unusually warm, melt-favourable conditions over 1–60 days before the 26 August 2026 event, with the strongest short-window anomalies reaching the 100th empirical percentile. Long-term PDD and Sentinel-1 backscatter trends are statistically supported, whereas temperature and optical trends are not supported under the autocorrelation-aware tests. Sentinel-1 provides a high-coverage event pair but no validated spatial failure footprint; Sentinel-2 post-event coverage is insufficient for quantitative optical change. The monthly unsupervised model does not identify a unique extreme pre-event state. The defensible inference is therefore unusually warm and melt-favourable environmental conditions consistent with conditioning, not identification of a physical trigger, causal mechanism, predictive precursor, or early-warning system.

Code, data, and reproducibility

The manuscript is tied to a frozen experimental baseline so that reported numbers can be traced to machine-readable artifacts rather than manually transcribed analytical claims. The repository contains the analysis modules, notebooks, configuration, tests, QA outputs, tables, and generated figures. Automated local tests of the frozen repository completed with six passed tests, six Earth-Engine-dependent skips, and eleven passing subtests; the skips reflect external execution dependencies rather than failed numerical assertions.

Code and Data Availability

The analysis code, configuration, quality-assurance outputs, machine-readable result tables, anomaly-model diagnostics, and figure-generation workflow corresponding to this study are publicly available in the project repository. The V2 experimental baseline reported in this manuscript is pre-

served in the immutable GitHub release v2.0.0-arxiv: <https://github.com/debabratapruseth/Langtang-Glacier-AI-Multisensor-Analysis-of-the-2026-Langtang-Lirung-Event/releases/tag/v2.0.0-arxiv> Machine-readable supporting artifacts are mapped by repository-relative path in the supplementary artifact inventory. Subsequent development on the repository’s main branch is not part of the analysis reported in this manuscript.




Figure 1: Seasonally adjusted environmental indicators. Thin lines show monthly values and thick lines the centered 12-month mean. Statistical trend tests are reported in Table 2.




Figure 2: Pre-event observations compared with historical same-calendar-month distributions in the integrated analysis. This monthly-context figure complements the exact 1–90-day matched percentiles in Table 3.




Figure 3: Comparable-track Sentinel-1 VV and VH monthly backscatter for the selected geometry, with scene-count context.




Figure 4: Monthly unsupervised environmental anomaly score. The dashed line is the 95th percentile of the pre-2026 training ensemble distribution. A missing score denotes insufficient complete features rather than a zero anomaly.




Figure 5: Integrated analytical timeline. The shaded/event annotations provide temporal context only; the displayed associations are not causal evidence.

A Supplementary figures and diagnostics

Supplementary artifact provenance. The supplementary materials are drawn directly from the frozen analysis repository. Table 7 provides the repository-relative path for each artifact, allowing every supplementary result to be traced to the corresponding file under `outputs/charts`, `outputs/qa`, or `outputs/tables`. PNG artifacts are presented as supplementary figures where appropriate, while CSV and JSON artifacts are retained as machine-readable supporting evidence in the code repository. Artifact identifiers A01–A15 are repository inventory identifiers and are distinct from supplementary figure numbering. Figure 6 documents integrated feature availability; Figures 7–9 provide raw ERA5-Land context; Figure 10 shows post-monsoon Sentinel-2 proxy sensitivity; Figure 11 shows exploratory PCA regimes; Figure 12 shows cross-feature correlations; and Figure 13 shows the temperature STL decomposition. These are useful reproducibility and sensitivity diagnostics but are secondary to the principal event-centered inference. Supplementary machine-readable diagnostics retained in the repository include incidence-angle sensitivity, NDSI-threshold sensitivity, optical proxy trend sensitivity, feature and ROI provenance, change-point diagnostics, PCA loadings, lagged correlations, pair-count diagnostics, and a machine-readable research-question synthesis.

Table 7: Repository artifact and provenance inventory.

ID	Repository path	Purpose
A01	<code>outputs/charts/temperature_stl_decomposition.png</code>	Temperature seasonality/trend decomposition
A02	<code>outputs/charts/correlation_heatmap.png</code>	Cross-feature dependence
A03	<code>outputs/charts/environmental_regimes_pca.png</code>	Exploratory regime separation
A04	<code>outputs/charts/integrated_feature_availability.png</code>	Missingness and availability
A05	<code>outputs/charts/sentinel2_post_monsoon_proxy.png</code>	Optical proxy sensitivity
A06	<code>outputs/qa/sentinel1_angle_mask_sensitivity.png</code>	Incidence-angle QA
A07	<code>outputs/qa/sentinel2_ndsi_sensitivity.png</code>	NDSI threshold sensitivity
A08	<code>outputs/qa/sentinel2_proxy_trend_sensitivity.png</code>	Optical trend sensitivity
A09	<code>outputs/qa/feature_provenance.csv</code>	Feature provenance
A10	<code>outputs/qa/roi_provenance.json</code>	ROI provenance
A11	<code>outputs/tables/change_point_diagnostics.csv</code>	Change-point diagnostics
A12	<code>outputs/tables/pca_loadings.csv</code>	PCA interpretation
A13	<code>outputs/tables/lagged_correlations.csv</code>	Lagged relationship diagnostics
A14	<code>outputs/tables/correlation_pair_counts.csv</code>	Correlation sample-size diagnostics
A15	<code>outputs/tables/research_question_summary.md</code>	Machine-readable RQ synthesis




Figure 6: Integrated monthly feature availability.




Figure 7: ERA5-Land regional temperature context.




Figure 8: ERA5-Land positive degree-day context.




Figure 9: ERA5-Land regional precipitation context.




Figure 10: Post-monsoon Sentinel-2 snow/clean-ice proxy sensitivity.




Figure 11: Exploratory environmental regimes in a two-axis PCA representation. Cluster labels are arbitrary and are not hazard classes.




Figure 12: Pairwise seasonally adjusted feature correlations. Pair counts vary with sensor availability.




Figure 13: STL decomposition of the monthly temperature series.