A Plot-Level Hazard and Living Suitability Screen for Hill Land

HillScan Reports

October 2026

Abstract

National and state-level hazard maps are built to guide planning decisions at the scale of a district or a watershed. They have no demonstrated skill at the scale that matters to a buyer, lender, or developer: one plot of land. This paper describes a screening model that scores individual hill-land parcels in Uttarakhand and Himachal Pradesh, India, for landslide, flood, and glacial-lake-outburst hazard, and for living suitability, at 30-metre ground resolution inside each of 136,935 scored patches. The model is built from satellite elevation, open hazard and climate records, and road-network access data, scored against fixed, pre-declared thresholds. It is tested against landslide, flood, and outburst events the model had not seen, and separately against live property-listing prices, a test the model was not built to pass. Results, data sources, and stated limits follow.

1. Introduction

Hill land plots that look alike on a brochure behave very differently on the ground. Two parcels a few hundred metres apart can sit on stable terrain and on an active slide path respectively, differ by tens of metres in distance from a seasonal torrent, or fall on opposite sides of a glacial-lake-outburst corridor — none of which shows up in a listing photo or a plot boundary survey.

Hazard information that does exist is built for a different scale. National landslide susceptibility maps and state disaster atlases are produced to guide district- or watershed-level planning. Tested directly against documented landslide events recorded after the model was built, a national susceptibility layer of this kind showed no usable skill at the scale of an individual parcel. The gap is not a flaw in those maps — they were not built for plot-level decisions — but it leaves buyers, lenders, developers, and the lawyers and brokers who advise them without a number that applies to the specific plot in front of them.

This paper describes a model built to close that gap: one score, and a set of underlying measurements, per plot, each traceable to a named source, each tested against events the model did not see during development.

2. Coverage

The screen covers, starting with Uttarakhand and Himachal Pradesh:

- 136,935 scored patches, each approximately 0.8 km² — an H3 hexagonal cell at resolution 8.
- 2 states, 20 districts, 200 tehsils (sub-districts).
- 30-metre ground resolution for terrain measurements inside each patch.
- 80 or more measured values, plus 27 climate and hazard rows (each given a fixed 1–5 score), per plot report, organized across 9 report sections.

The report structure was cross-checked line by line against two published international sample report editions for hill land. All 72 line items of those editions are accounted for in this product: each has either a matching data source and licence, or an explicit stated reason it is not, or not yet, included.

3. Method

The score for each patch is built from four layers of measurement. The exact thresholds, factor weights, and internal rule structure used to combine them are not published (Section 7); what follows describes what is measured and how it is checked, not the rules themselves.

3.1 Terrain geometry

Slope, elevation, aspect, and drainage are derived from satellite elevation data at 30-metre resolution. A row/column duplication artifact was found in the raw elevation export during development — a resampling effect of the export process, not a feature of the terrain — and corrected before any terrain-derived layer was computed from it.

3.2 Hazard exposure

An independent landslide-susceptibility estimate, river and torrent proximity, and traced glacial-lake and glacial-lake-outburst-flood corridors are computed for every patch. Where two or more of these layers cover the same ground, they are cross-checked against each other.

3.3 Climate and access

Heat, frost, fire weather, drought, heavy rainfall, and wind are drawn from open climate records. Access is measured as road-network travel distance — not straight-line distance — to the district capital, the nearest highway, the nearest hospital, the nearest school, and the nearest market town.

3.4 Fixed, pre-declared thresholds

Every hazard band and flag is set against a physical cutoff decided before scoring begins. None of these cutoffs are rescaled relative to whatever else happens to be on the map on a given run. The practical consequence: the same ground reads the same way regardless of what else is nearby or what else has been scored that day.

The remainder of this paper — validation results, market and quality-control checks, limitations and data sources — is available on request.
Contact hello@hillscanreport.com.

4. Validation

The model is tested against hazard events it had not seen during development. Three separate backtests are reported below, followed by a market-based check in Section 5 that is independent of both.

4.1 Held-out backtest

The landslide layer was trained on one landslide inventory and tested on [illegible] recorded afterward across Himachal Pradesh and Uttarakhand — events outside the training data by construction.

Table 1 — Held-out backtest, [illegible] landslide events ([illegible] and other reports, in government records)

[illegible]	Result
Events in the highest-ranked fifth of hill pincher	[illegible]
Same-district accuracy (AUC, [illegible] district)	[illegible]
Road-matched accuracy (AUC)	[illegible]
Same-district AUC — [illegible] for government records	[illegible]

4.2 Watched-ground backtest

A second backtest compares [illegible] previously located events from a national landslide catalogue, [illegible] independently documented local incidents, and [illegible] glacial-lake outburst

events, each against matched ground nearby — the same state block, the same terrain class, the same elevation band, and in some comparisons the same road-distance band.

Table 2. Matched-ground backtest results

Flag or band	Comparison	Lift over matched ground
Tornado – storm flood flag	National landslide analogue	[illegible]
Tornado – storm flood flag	Environmental-band landslide	[illegible]
Elevated lake sediment corridor	[illegible] of its documented sediment events, vs. [illegible] have rain on matched ground	[illegible]

The observed-occurrence flood grade was checked for accuracy (AUC) against an independent satellite-derived flood map, [illegible] overall, and [illegible] for locations with more than one recorded flood.

[illegible] events filtered scores against [illegible] set aside events

The combined hazard score is a model fitted to recorded events. Its landslide and flood terms were fitted (shallow gradient-boosted trees, flood settings) on [illegible] recorded landslide, flood and sediment events with exact or village-level locations, using terrain, channels and height above them, rainfall, soil, vegetation and built-up ground; the flood term is averaged half-and-half with the physical flood score, and the earthquake term is unchanged. Another [illegible] events were set aside before fitting, split by whole map cells so that neighbouring events may not sit one side, and balanced by state block and event type. Every design choice was made on the [illegible] events, which were inspected over three rounds of changes; the set-aside events were scored once for the decision. A robustness check was afterwards [illegible] ground beside the set-aside cells, which the model had not seen in training; given the figures in Table 3, with comparison ground the model had seen, the first read gave an AUC of [illegible]. Each set-aside event was compared with ground of the same type in the same region that had no record: same state block, terrain class, river setting, and road-distance band, using comparison ground the model had not seen.

Table 3. Event-fitted scores against matched ground, [illegible] set-aside events

Measure	Result
Accuracy (AUC) [illegible] — chance, [illegible] of the [illegible] events [illegible] had no comparison ground)	[illegible] to [illegible], [illegible] by map cell, as [illegible]

Gain over the earlier physical event	[illegible]
Inventory (AUC), earlier physical event or the same events	[illegible]
Events in the highest-ranked fifth of matched ground	[illegible]
Events in the highest-ranked tenth	[illegible]
Events in the lowest-ranked fifth	[illegible]
Events on the top half of the land by score (50% of land)	[illegible]
Events in the lowest-scored 1–5% of land	[illegible]
Landslide terms on landslide events / flood terms on flood events	[illegible]
Primary / runners / Himachal block	[illegible]

Checks on this result: a random same-type comparison patch given to each event scores as pre-dropping any one map cell leaves the AUC between 0.91 and 0.93; the gate holds for landslides and floods, for events before and after 2018, and for snow and government records. It is smallest for the Geological Survey of India mapped slides in Himachal cohort inventory sliders, where the change is not distinguishable from noise. The earlier version (weights fixed from physical reasoning) scored 0.84 on 487 events and 0.88 on 2019 alone, the new sliders, and that earlier set-aside group held five western-block events (5% of 2019, as the table's 0.71 to the like-for-like comparison. The share-of-land cross describes the same events and are not a separate test.

5. Market validation

A separate check was run against live market data rather than against disaster records. Asking prices were collected from several property portals (photo omit). The largest portal was used as the guide. Its listings were dropped by robust outlier removal within each town and logic checks, with no fixed price range, and other portals' listings were kept only where they fell within 15 factor of four of that guide. Prices are expressed as price per square metre relative to the median of each town. Listings that name a local land unit (kanal, bigha, kanal, marla, sqft) were not priced. Prices were compared at the level of the locality, not the individual listing, since many listings in one locality share the same geocoded point. Each locality is scored as the mean of its hexagon and the six neighbours.

- across 48k places (74k listings, after merging duplicate place names and dropping placenames more than 25 km from their town), the suitability score (shown to reports as Living Suitability, with an honest dedicated) tracks relative price. Spearman $\rho = +0.74$ (p<0.01 to +0.68 to +0.78, bootstrap resample whole-pl but' map-valley).
- On hill terrain (220 places) $\rho = +0.31$ (p<0.01 to +0.28, fit a fitted set of 50 hill places to map-wills set aside before the method was checked), $\rho = +0.29$ (p<0.01 to +0.26). On plains and foothills there is no clear link ($\rho = +0.06$, n.s.) to +0.03). Outside Ombwallon, hill places give $\rho = +0.29$ (p<0.01 to +0.26), so the hill result depends on Ombwallon, where within-town ordering is cleanest ($\rho = +0.38$ in the fitted set).
- Two variants were tried on three fitted rounds and the sealed set: adding an urban-intensity group (hill $\rho = +0.38$, a difference of +0.0100 with interval −0.0042 to +0.0162 against the published score) and weights fitted on the three rounds (hill $\rho = +0.39$, difference +0.0148, interval −0.0028 to +0.0301). Neither is shown to be better, so the published score is unchanged.
- The physical trend hazard shown on prior discussion (hill $\rho = +0.29$, n.s.; 0.31 to +0.29, fitted set +0.30, n.s.; 0.40 to +0.02), and neither does the event-fitted score. Settled, accessible land has both higher prices and a higher hazard score, because the hazard terms include built-up exposure. An earlier discussion of $\rho = +0.29$ came from a table in which one of pair rows were duplicate names for the places, and it does not hold on the cleaned table. Risk is not presented as a price driver.
- In 48 Himachal places (mostly a single listing, each) suitability ($\rho = +0.03$) shows no positive link with price; the sample is too small to say more.
- In a check on the listing data itself, plot rates published by a property portal for 19 Ombwallon localities (price per square yard, plots only, apartment rates were kept apart) agree with the median of our own listings for the same localities ($\rho = +0.86$).

Earlier figures from a smaller sample (28 localities) were stronger and did not hold when the sample grew; the figures above replace them.

6. Quality control

Every report is checked before it is scored, and the scoring is checked before it is sent.

- 4,200 of 4,200 automated checks gave across 14 test plots covering every report section, with zero failures.
- 48 scored patches reproduce bit-for-bit on an independent re-run of the scoring pipeline — 4 given plot's score does not depend on when or how it was requested.

- A targeted quality pass drew a no-zero-patch sample from a by the radius around a named towns with active property interest including Shawnee, Stillwater, Yamhill, Shiloh, Hawk, and Pharmacokinetics instrument geometry and statistical consistency checks were run — recomputing terrain slope and the final score directly from stored inputs and drilling against the record result, checking hazard flags against physical ranges, ranking patches by statistical deviation from same-terrain, same-hazard based peers. The pass never statistically unusual patches, plus a separate zero-patch random control group, were then independently reviewed for geographic plausibility. Zero confirmed anomalies in either group.

9. Limitations

- This is a relative-risk screening tool, built on a model trained against historical disaster records and compared with asking prices. It does not replace official certifications or the qualified engineering, legal, and administrative requirements a purchase, loan, or build will needs.
- The hardest event samples are modest in size (n=, p, by, to and spit events respectively). The AP and accuracy figures in Section 4 should be read as directions, not as an exact probability for any one plot.
- The event-blind score uses built-up ground, as part of by enabling reflects where events get recorded (settled ground and roads). On the 412 wedding events, matching comparison ground as built-up shows large more than half of the gain over the physical score, not all of it. The landslide term also takes a published susceptibility map as one input, and that map's training events may overlap the test events. The test covers 10 map cells, 412 of the 438 events were in hill or high-mountain terrain and 26 on plains, so the event-blind score is reported for sloping hill land only and the physical drivers stand alone on flat ground, plains and foothills. The Geological Survey of India mapped slides in Himachal improves least. The earlier version dilutions (landslide driver without the susceptibility map (=3), built-up shows alone (=2) without matching describe the earlier physical score.
- Event records show toward locations near roads and towns — a known reporting bias in disaster records generally.
- The market check in Section 5 uses asking prices, not sale prices. Over half of the localities have a single listing, and Silverados supplies many of the hill localities without it the hill bids to weaken (=+0.20, interval includes zero). The score is a ranking, not a probability.
- The exact scoring thresholds, factor weights, and internal rule set are not published. Method, hill data sources, and complete test results are available on request.

8. Conclusion

The model scores calculated across Uttarakhand and Himachal Pradesh for hazard and living suitability, at 30-metre ground resolution, from measurements that trace to named open sources. A separate events it had not seen in development from similar nearby ground at rates above chance, as reported in Section 5, and its living suitability score shows a modest link with asking prices on full hand, mostly through elevation, as reported in Section 5. Both checks are reported with their sample sizes and confidence intervals, and the limits of each are stated in Section 7.

9. Data sources

- Copernicus DEM (elevation)
- OpenStreetMap contributors (ODbL) — road network and access
- Bhuvan Land — climate records
- JRC Global Surface Water
- Bi Atlas glacial lakes
- Randolph Glacier Inventory
- Geological Survey of India landslide inventory
- WorldPop