ACCSCIENCE PUBLISHING

Asian Journal of Water, Environment and Pollution

ORIGINAL RESEARCH ARTICLE

Mapping chronic environmental pressure in Goa using machine learning and deep learning

Madhusudhan Das[1], Saurabh K. Gupta[1], Shruti Kanga[2]*, Divyansh Kumar Pandey[2], Bhartendu Sajan[1], Gowhar Meraj[3,4], and Suraj Kumar Singh[1]

[1]Centre for Climate Change and Water Research, Suresh Gyan Vihar University, Jaipur, Rajasthan, India

[2]Department of Geography, School of Environment and Earth Sciences, Central University of Punjab, Bathinda, Punjab, India

[3]Center of Geospatial Intelligence, American University of Sharjah, Sharjah, United Arab Emirates

[4]Department of Civil Engineering, College of Engineering, American University of Sharjah, Sharjah, United Arab Emirates

Abstract

Coastal landscapes experience overlapping environmental pressures that require integrated spatial assessment. This study combines satellite and geospatial indicators to map chronic environmental pressure in Goa, India. Fifteen predictors represent atmospheric composition, terrain, land-surface conditions, land use, and human activity. Time-filtered products cover available observations within a nominal 2018–2023 period, alongside static and differently dated products. Of 89 expert-informed reference records, 81 supplied predictor information for three tree-based algorithms and four neural-network architectures. Their probabilities were combined into an environmental-pressure classification. Low-, moderate-, and high-classes accounted for 56.75%, 30.54%, and 12.71% of valid mapped pixels, respectively. Higher scores characterised urban and industrial corridors, whereas forested eastern areas generally recorded lower scores. Night-time lights had the largest contribution to the high-class extreme gradient boosting Shapley additive explanations. Model scores align with the expert-informed reference classes. The integrated maps provide a relative screening framework in selecting locations for further environmental observation.

Keywords: Environmental pressure; Spatial screening; Remote sensing; Machine learning; Deep learning; Model interpretation; Goa

***Corresponding author:**
Shruti Kanga
(shruti.kanga@cup.edu.in)

Citation: Das M, Gupta SK, Kanga S, *et al*. Mapping chronic environmental pressure in Goa using machine learning and deep learning. *Asian J Water Environ Pollut.*
doi: 10.36922/AJWEP026340243

Received: August 22, 2026

Revised: October 16, 2026

Accepted: October 16, 2026

Published online: November 5, 2026

Copyright: © 2026 Author(s). This is an Open-Access article distributed under the terms of the Creative Commons Attribution License, permitting distribution, and reproduction in any medium, provided the original work is properly cited.

Publisher's Note: AccScience Publishing remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

1. Introduction

Human-induced warming has increased the frequency and intensity of hot extremes and heavy precipitation in many regions.[1] In coastal tropical environments, they are also governed by urbanisation, industrial activity, and land-use change.[2,3] Along the Indian coast, heavy monsoonal rainfall, storm surges, and sea-level rise contribute to compound flooding.[4-8] Single-hazard assessments may overlook interdependencies among these processes.[9,10] Heat, flooding, and poor air quality can occur simultaneously, with consequences dictated by their timing, intensity, and the exposure and vulnerability of affected populations.[11,12] Integrated assessment can identify overlapping pressures and

inform adaptation; however, spatial overlap alone does not illustrate an interaction that amplifies damage.[13,14]

Research on compound hazards has expanded considerably during the past decade. Coastal flood studies integrate rainfall, tides, storm surges, and sea-level rise,[15,16] while other assessments examine relationships among floods, landslides, and wildfires.[17-19] Studies of surface urban heat islands link urban expansion and vegetation loss to higher land-surface temperatures across climatic settings.[20-23] These findings motivate joint assessment while retaining distinctions among surface temperature, human heat stress, and measured hazard impacts.

Machine learning and deep learning can capture nonlinear relationships among climatic, topographic, hydrological, and land-use variables.[24-28] Their usefulness depends on the quality of reference observations and on validation that reflects the intended prediction task. Environmental applications include landslide susceptibility, spatiotemporal prediction, and urban thermal modelling.[29,30] However, improved fit to the available observations does not necessarily imply improved prediction at unobserved locations. Explainable artificial intelligence, including Shapley additive explanations (SHAP), helps describe how predictors contribute to model outputs.[31-33] Such attribution can reveal how a fitted model uses rainfall, terrain, vegetation, and land-use information; however, it does not establish causal effects. Predictive uncertainty must be assessed separately from model fit and the interpretation of individual predictors.[34]

Multi-hazard assessment increasingly integrates physical hazards with information on exposure and social vulnerability.[35-41] The distinction is important here: an environmental classification based on expert-informed labels and satellite proxies is a relative screening measure. Without independently observed consequences, it cannot estimate the probability of loss, expected damage, or health burden. Reviews also emphasise the need to connect data-driven methods with physical knowledge and clearly defined evaluation procedures.[42]

An important practical question is how openly available geospatial indicators can be combined to represent persistent environmental pressures across a small, heterogeneous coastal state. Published studies have documented mining-related degradation in Goa, village-level coastal vulnerability, and microplastic contamination in coastal and estuarine environments.[43-45] Studying these pressures requires different observations. The present framework focuses on flood susceptibility, surface thermal conditions, and atmospheric indicators that can be represented with spatially continuous geospatial data. Mining impacts and marine microplastics are not within the modelled scope. The flood component acts as a susceptibility screen; it does not simulate coastal water levels, storm surges, or inundation depths.

This study develops a framework for mapping chronic environmental pressure across Goa. Its objectives are to characterise the fitted tree-based and neural-network classifications, map composite and component scores, examine spatial differences among these scores, and describe predictor contributions using SHAP. The intended application is to identify candidate locations for additional environmental observation and assessment.

2. Materials and methods

2.1. Study area

Goa lies along the Konkan coast of western India, between approximately 14.90° N–15.80° N and 73.68° E–74.34° E (Figure 1). It is bounded by the Arabian Sea to the west, the Western Ghats to the east, Maharashtra to the north, and Karnataka to the south. The state covers about 3,702 km^2, with a population of approximately 1.46 million according to the 2011 Census.[46] Its tropical monsoon climate and marked coast-to-upland gradient influence environmental conditions. Coastal lowlands contain settlements, ports, and estuarine wetlands, while the lateritic midlands support agriculture and mixed land use. The eastern hills contain extensive forests and protected areas, including Bhagwan Mahavir Wildlife Sanctuary and Mollem National Park, Mhadei, Netravali, and Cotigao. Industrial and urban centres include Vasco da Gama, Ponda, Kundaim, Bicholim, Tuem, Honda, and Cuncolim. Forests, agricultural land, settlements, and industrial areas are in proximity. This heterogeneity provides a useful setting for comparing environmental classifications, although the reference sample must adequately represent these conditions.

2.2. Datasets

A multisource predictor stack was prepared in Google Earth Engine.[47] The supervised models utilised 15 variables: six atmospheric indicators, cross-polarised/vertical–horizontal (VH) backscatter variability, land-surface temperature, Shuttle Radar Topography Mission (SRTM) elevation and slope, population, night-time lights (NTLs), a MERIT Hydro elevation band, a vegetation-related band ratio, and a land-cover weight (Table 1). These variables correspond to export bands 1–14 and 16. The categorical land-cover band, separately calculated weighted indices, and unsupervised clusters were not inputs to these supervised models. Atmospheric columns and surface temperature were interpreted as environmental indicators rather than direct measures of surface exposure

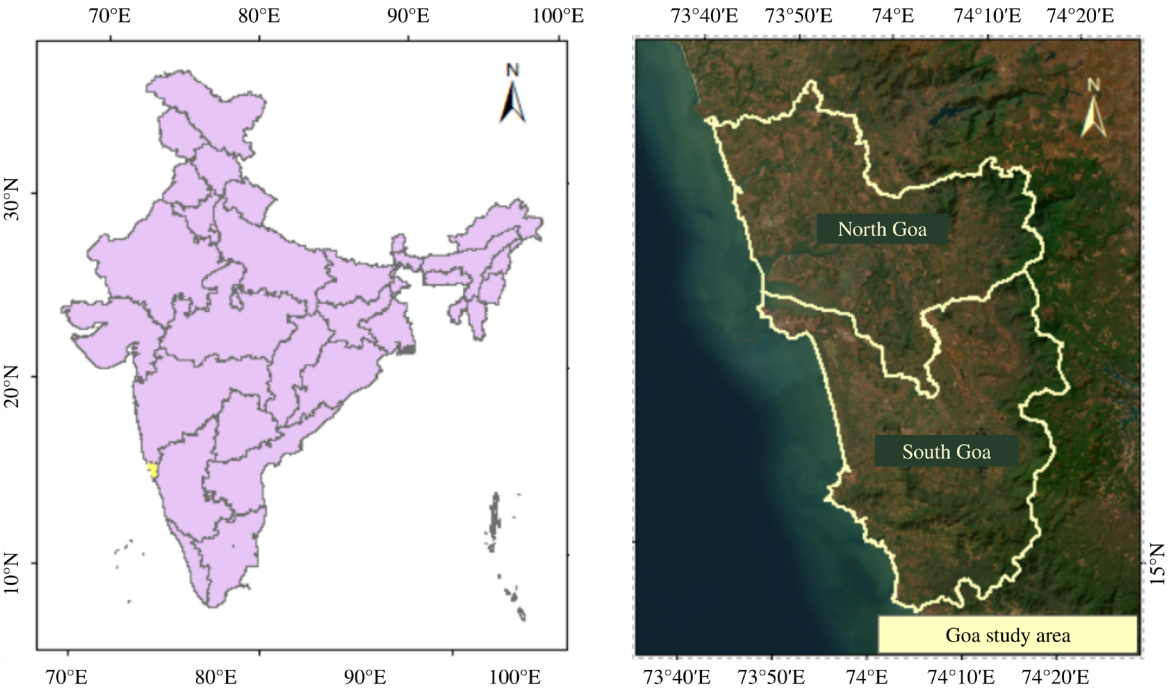

Figure 1. Location of Goa and its North Goa and South Goa administrative districts

or physiological heat stress.

The time-filtered collections used January 1, 2018, as the inclusive start and December 31, 2023, as the exclusive end. Sentinel-5P bands and VIIRS monthly radiance were averaged over available observations, while Sentinel-1 VH backscatter was summarised by its temporal standard deviation. Sentinel-5P availability begins in 2018.[48] Landsat 8 Collection 2 Level-2 scenes were selected using scene cloud cover below 15% and composited using the median. No additional pixel-level cloud, shadow, or saturation masks were applied. Meanwhile, WorldCover v100 represents 2020 data.[49] WorldPop was averaged across the images intersecting Goa without a date filter, and therefore does not represent a 2018–2023 population mean.[50] Terrain products are static. These different temporal supports preclude interpreting the stack as simultaneous observations of individual hazard events.

Land-surface temperature was obtained from ST_B10 using the scale and offset in Table 1. The vegetation predictor, labelled normalised difference vegetation index $(NDVI)_{inv}$, was calculated as one minus the normalised difference between the median SR_{B5} and SR_{B4} digital numbers. The reflectance scale factor of 0.0000275 and additive offset of −0.2 were not applied before this ratio.[51] Because the offset is not cancelled from the denominator, this variable is not equivalent to one minus reflectance-based NDVI. The hydrological predictor labelled TWI is the MERIT Hydro hydrologically adjusted elevation band (hnd), measured in metres; it is not a dimensionless topographic wetness index.[52] Land-cover classes 10, 20, 30, 40, 50, 60, 80, 90, and 95 received fixed weights of 0.2, 0.3, 0.3, 0.6, 1.0, 0.5, 0.1, 0.2, and 0.1, respectively. Unlisted classes were masked. These weights encode modelling assumptions rather than measured environmental effects.

The predictor stack was exported as floating-point bands in EPSG:4326 at a nominal scale of 30 m. No resampling kernel was explicitly specified. Computational grid spacing differs from source resolution: most TROPOMI ultraviolet-visible products have nadir footprints of approximately 3.5 × 7 km early in the mission and 3.5 × 5.5 km after August 2019, with coarser carbon monoxide footprints.[48] Landsat thermal observations are acquired at approximately 100 m and distributed on a 30 m grid.[51] Exporting these sources to a common grid does not create independent information below their native support. The 15 selected predictor bands were not percentile-normalised before export; imputation and standardisation for supervised fitting are described in Section 2.3.1.

Table 1. Geospatial predictors and source-product characteristics

Feature ID	Variable	Source	Source support	Processing/Units
NO_2	Tropospheric NO_2 column	S5P OFFL L3	3.5 × 7 km/3.5 × 5.5 km	Available-period mean; mol/m^2
AAI	Absorbing aerosol index (aerosol proxy, not direct PM2.5)	S5P NRTI AER AI	3.5 × 7 km/3.5 × 5.5 km	Available-period mean; dimensionless
SO_2	SO_2 column density	S5P OFFL L3	3.5 × 7 km/3.5 × 5.5 km	Available-period mean; mol/m^2
CO	CO column density	S5P OFFL L3	7 × 7 km/7 × 5.5 km	Available-period mean; mol/m^2
O_3	O_3 column density	S5P NRTI L3	3.5 × 7 km/3.5 × 5.5 km	Available-period mean; mol/m^2
HCHO	Formaldehyde column	S5P NRTI L3	3.5 × 7 km/3.5 × 5.5 km	Available-period mean; mol/m^2
VH_{stdDev}	SAR backscatter standard deviation	Sentinel-1 GRD	10 m	Temporal variability in VH backscatter; dB
LST	Land-surface temperature	Landsat-8 ST_B10	100 m thermal; 30 m grid	ST_B10 × 0.00341802 + 149 − 273.15; °C
Elevation	DEM	SRTM 1-arc	30 m	Metres AMSL
Slope	Terrain slope	Derived from SRTM	30 m	Degrees
Pop	Population indicator	WorldPop 100 m	100 m	Mean without date filtering; persons per native cell
NTL	Night-time light radiance	VIIRS DNB monthly	~ 500 m	Mean radiance; nW/cm^2·sr^1
TWI	MERIT Hydro hydrologically adjusted elevation	MERIT Hydro	~ 90 m	hnd band; m; legacy model identifier TWI
$NDVI_{inv}$	Inverted digital-number band ratio	Landsat-8 SR	30 m	1 − (SR_B5 − SR_B4)/(SR_B5 + SR_B4); unscaled digital numbers
$LULC_{Weight}$	ESA WorldCover risk weight	ESA WorldCover v100	10 m	2020 map; fixed class weights listed in Section 2.2

Note: Temporal support varies by dataset.
Abbreviations: AMSL: Above mean sea level; DEM: Digital elevation model; LULC: Land use/land cover; NDVI: Normalised difference vegetation index; SAR: Synthetic aperture radar; SRTM: Shuttle Radar Topography Mission; TWI: Topographic wetness index; VH: Cross-polarised/Vertical–horizontal; VIIRS DNB: Visible infrared imaging radiometer suite day/night band.

2.3. Methodology

The workflow comprised predictor preparation, reference classification, model fitting, ensemble integration, spatial mapping, and model interpretation. Reference records represented three expert-informed environmental-pressure classes. The fitting and evaluation procedures are described in Sections 2.3.1–2.3.4.

Three tree-based algorithms formed the machine-learning ensemble, while four neural-network architectures provided additional probability estimates. Weighted soft voting combined their outputs. Spatial exclusion reduced proximity between records used to fit and assess the tree classifiers.[53,54] Class-probability entropy and differences among component scores describe the mapped classifications, while SHAP characterises the fitted extreme gradient boosting (XGBoost) component.[55,56]

2.3.1. Reference dataset and class balancing

The reference dataset comprised 89 records, with 34 low, 29 moderate, and 26 high classifications. Field observations, expert assessment, Goa State Pollution Control Board information, and high-resolution imagery informed these classifications. They represent a common assessment of environmental pressure rather than independent measurements of a single physical outcome. The records contained 88 distinct coordinate pairs; two low-class records shared the same coordinates and labels. Therefore, sample counts referred to records rather than independent sampling locations.

Reference points were assigned to low, moderate, or high classes based on vegetation, land use, industrial activity, and proximity to pollution sources. Low-class sites included relatively undisturbed forests and protected areas; Moderate-class sites included agricultural, semi-urban, and coastal settlements; whereas high-class sites included industrial estates, ports, transport corridors, and dense urban centres. These classes describe relative chronic environmental pressure.

Predictor values were extracted as means within 3 × 3 pixel windows. Records were retained when at least one selected predictor contained a valid value; missing predictor values were subsequently replaced by the median. This procedure retained 81 records: 31 low, 25 moderate, and 25 high (Table 2). The median imputer and standardisation parameters were estimated from all retained records before the spatial loop. Within each tree-classifier fold, the test record and nearby records were excluded from classifier fitting, and ordinary synthetic minority over-sampling technique (SMOTE) was applied to the remaining training subset.[57] Nevertheless, the held-out record contributed to the preprocessing estimates. This information transfer can bias the generalisation assessment and is not removed by simply excluding the record from classifier fitting.[58,59]

For global model fitting, borderline-SMOTE balanced the standardised reference sample, with ordinary SMOTE as a fallback.[57,60] The balanced sample contained 93 records. An 80:20 stratified split of that sample supplied the neural-network training and internal validation subsets. Therefore, oversampling preceded the split: synthetic records could be derived from real observations in either subset, making the internal validation loss an unsuitable estimate of independent performance. For a minority-class vector x_i, interpolation takes the form of **Equation 1**:

$$x_{new} = x_i + \lambda\left(x_{nn} - x_i\right) \quad (1)$$

where x_{nn} is a selected minority-class neighbour in feature space, and λ is sampled between 0 and 1. The synthetic vector lies on the segment joining the two feature vectors. Spatial-loop evaluation used real held-out records, whereas the neural-network internal split was drawn from the globally balanced sample.

Land-cover weights, optical band ratio, population, and NTLs partly represent the characteristics used to assign reference classes. Therefore, agreement measures reproduction of the expert-informed classification. Geographic separation does not remove this dependence between labels and predictors; independent environmental validation requires separately observed outcomes.[59]

Table 2. Initial reference records and records retained after predictor extraction, by composite environmental-pressure class

Pressure class	Original (*n* = 89)	Valid (*n*= 81)	Site descriptors
Low	34	31	Forested/Protected
Moderate	29	25	Agricultural, semi-urban, and coastal sites
High	26	25	Urban and industrial sites
Total	89	81	-

2.3.2. Machine learning and deep learning models

The model set comprised LightGBM, XGBoost, and random forest, along with an attention multilayer perceptron (AttentionMLP), a one-dimensional convolutional neural network (CNN-1D), a residual multilayer perceptron (ResMLP), and TabNet. The tree algorithms formed one ensemble. Results are reported for that ensemble, the four neural networks, and the full ensemble, yielding six evaluated configurations rather than six individual algorithms. Their principal hyperparameters are summarised in Table 3.

(a) Machine-learning ensemble

LightGBM and XGBoost build boosted decision trees, whereas random forest aggregates trees fitted to bootstrap samples.[61-63] These methods can represent nonlinear relationships and interactions without imposing a linear response. Regularisation, restricted tree complexity, and aggregation can limit overfitting. Correlated predictors can also redistribute feature importance among related variables.

(b) Deep-learning architectures

AttentionMLP applies a trainable feature gate before classification through a fully connected network with hidden layers of 128, 64, and 32 units. Each hidden layer uses batch normalisation, Gaussian error linear unit activation, and a dropout of 0.30. The three PyTorch networks were fitted using AdamW, a learning rate of 0.001, weight decay of 0.0001, and cosine learning-rate scheduling. The feature gate comprises two linear transformations separated by a hyperbolic tangent activation:

$$A = \sigma\left(W_2 \tanh\left(W_1 x + b_1\right) + b_2\right) \quad (2)$$

where x is the predictor vector, W_1 and W_2 are weight matrices, b_1 and b_2 are bias vectors, and σ is the sigmoid function. The gated predictor vector is shown in **Equation 3**:

$$x' = A \odot x \quad (3)$$

where $\odot$ denotes element-wise multiplication. CNN-1D treats the predictor vector as a one-dimensional sequence, using 32 and 64 convolutional filters with a kernel size of 3 and adaptive pooling to 4 positions. Adjacency refers to the predictor order, not geographic proximity. ResMLP uses a 128-unit representation and 2 residual blocks with layer normalisation, Gaussian error linear unit activation, and a dropout of 0.35. These architectures model relationships among predictors rather than spatial neighbourhoods in image patches.

TabNet uses sequential attention for feature selection, with decision and attention dimensions of 16, 3 steps, Adam optimisation, and a learning rate of 0.002.[64] The batch size used was 128. AttentionMLP, CNN-1D, and ResMLP utilised class-weighted categorical cross-entropy:

$$L = -\sum_{i=1}^{C} w_i \, y_i \log\left(\widehat{y}_i\right) \quad (4)$$

where $C = 3$ is the number of classes, w_i is the inverse-frequency class weight, y_i is the one-hot reference indicator, and $\widehat{y}_i$ is the predicted probability. The 3 PyTorch networks were trained for up to 250 epochs, with patience of 30 epochs, restoring the weights associated with the minimum internal validation loss. TabNet monitored internal validation accuracy for up to 150 epochs, with patience of 25.

2.3.3. Ensemble integration

Class probabilities from the tree ensemble and four neural networks were combined by weighted soft voting:

$$P_{\text{ensemble}} = \frac{\sum_{i=1}^{n} w_i P_i}{\sum_{i=1}^{n} w_i} \quad (5)$$

where P_i is the probability vector from component i and w_i is its weight. The tree-ensemble probabilities combined XGBoost, random forest, and LightGBM with relative weights of 2, 1, and 1.5. The full-ensemble weights were 0.45 for the tree ensemble, 0.20 for AttentionMLP, 0.15 for ResMLP, 0.10 for CNN-1D, and 0.10 for TabNet. These weights were fixed before ensemble integration. The categorical prediction is the class with the largest ensemble probability. The continuous score is $R = 0.5P(\text{Moderate}) + P(\text{High})$, an expected ordinal class score on a 0–1 scale.

2.3.4. Spatial assessment and mapping

(a) Evaluation procedures and their scope

Spatial dependence can cause random cross-validation to overstate geographic transfer when nearby records share environmental information.[53,65,66] The implemented tree-classifier check excluded each retained reference location and other records within a nominal 2 km neighbourhood from classifier fitting. The exclusion was not extended to preprocessing or neural fitting. For each test location, eligible classifier-training records satisfied:

$$d\left(x_i, x_{\text{test}}\right) \geq 2\text{km} \quad (6)$$

where the implemented distance d is 111.32 times the Euclidean separation between coordinates in decimal degrees. This approximation does not account for the latitude dependence of longitude distance. Ordinary SMOTE was fitted to each eligible tree-classifier training subset, while the median imputer and scaler had previously been fitted globally. The neural networks were not retrained within those subsets. The full-ensemble score in each iteration combined predictions from the refitted tree ensemble with predictions from the globally fitted neural models. Thus, its diagnostic scores arise from a mixed evaluation procedure, rather than either a complete spatial holdout or a common in-sample fit. No independent estimate of full-pipeline spatial predictive performance is reported. Metrics were calculated relative to the assigned three-class labels.

For a confusion matrix C with rows denoting reference classes and columns denoting predictions, overall accuracy was calculated as:

$$OA = \frac{c}{s} \quad (7)$$

where $c = \Sigma_k C_{kk}$ is the number of correct predictions and $s = \Sigma_i \Sigma_j C_{ij}$ is the total number of predictions.[67]

Balanced accuracy is the unweighted mean of recall across the $K = 3$ classes, shown as **Equation 8**:

$$BA = \frac{1}{K} \sum_{k=1}^{K} \frac{C_{kk}}{t_k} \quad (8)$$

where $t_k = \Sigma_j C_{kj}$ is the number of reference observations in class k.[68] For each class, the F1-score was calculated from one-versus-rest precision and recall as:

$$F1 = 2\left(\frac{Precision \times Recall}{Precision + Recall}\right) \quad (9)$$

with

$$Precision = \frac{TP}{TP + FP} \quad (10)$$

and

$$Recall = \frac{TP}{TP + FN} \tag{11}$$

Macro-F1 is the unweighted mean of the three class F1-scores, whereas weighted F1 weights them by their reference support. The multiclass Matthews correlation coefficient was calculated as:

$$MCC = \frac{cs - \sum_k p_k t_k}{\sqrt{\left(s^2 - \sum_k p_k^2\right)\left(s^2 - \sum_k t_k^2\right)}} \tag{12}$$

where $p_k = \Sigma_i C_{ik}$ is the number predicted in class k, t_k is the corresponding reference count, c is the number correctly classified, and s is the total sample size. Cohen's kappa was calculated as $\kappa = (p_o - p_e)/(1 - p_e)$, where $p_o = c/s$ and $p_e = \Sigma_k p_k t_k / s^2$. These definitions distinguish the multiclass measures from their binary special cases.[69]

(b) Spatial mapping and score diagnostics

For spatial prediction, the fitted models were applied across pixels with valid values for all 15 predictors. The predictor raster was processed in blocks of 50 rows to limit memory use. Predictor values were transformed using the scaling parameters used for the final model fit, and component class probabilities were combined with the specified weights. A common valid-data mask determined the population of mapped pixels. Reported fractions describe this mapped pixel population.

The outputs comprised the ensemble classification, a continuous ensemble score, and four XGBoost score maps representing flood-related conditions, surface thermal conditions, atmospheric composition, and compound environmental pressure. All are relative outputs of models fitted to assigned reference classes. The compound component is a composite classification rather than an additional physical hazard.

The flood-related component used MERIT hnd, slope, SRTM elevation, VH backscatter variability, the optical digital-number ratio, and NO_2. The thermal component used land-surface temperature, NTLs, population, the land-cover weight, the optical ratio, and SRTM elevation. The atmospheric component used NO_2, absorbing aerosol index (AAI), SO_2, CO, O_3, HCHO, population, and NTLs, while the compound classifier used all 15 predictors. Each parameter used by the XGBoost component is listed in Table 3. Component fitting followed global imputation and scaling, with ordinary SMOTE applied to its labelled training records and an unresampled fallback if resampling failed. Agreement was assessed on those same real records.

Class-probability dispersion was summarised by Shannon entropy:

$$H = -\sum_{i=1}^{n} P_i \ln(P_i) \tag{13}$$

where P_i is the probability assigned to class i and $n = 3$.[70] Natural logarithms yield entropy in nats, from 0 to ln(3), approximately 1.099. High entropy indicates a more even probability distribution; low entropy indicates a concentrated prediction. Entropy does not include all sources of uncertainty, such as label error and model bias.[71] The continuous ensemble score was smoothed with a Gaussian standard deviation of 1.2 pixels; component scores and entropy used 1.0 pixel. Categorical predictions were derived from unsmoothed probabilities. The component-score contrast, denoted HSI, was defined as:

$$HSI = R_{combined} - \max_{i=1,\ldots,n}(R_i) \tag{14}$$

where $R_{combined}$ is the smoothed ensemble score and R_i denotes a smoothed component score. The maximum includes the flood, thermal, atmospheric, and compound components. Positive values indicate that the ensemble score exceeds every component score. The spatial-summary threshold was HSI > 0.05; it differs from the total number of positive values and does not define the categorical map. Neither the sign nor this threshold establishes physical synergy, which would require calibrated outcomes and an explicit non-interaction benchmark.[72,73] Because invalid pixels were zero-filled before smoothing, scores near masked boundaries may be affected by the neighbouring mask. Different smoothing widths also influence the contrast.

2.3.5. Model interpretation

Shapley additive explanations were used to describe predictor contributions to the fitted XGBoost classifier.[55,56] TreeExplainer provides explanations for this tree model rather than for the complete heterogeneous ensemble. For a model output $f(x)$, the attribution is expressed as **Equation 15**:

$$f(x) = \phi_0 + \sum_{i=1}^{M} \phi_i \tag{15}$$

where $f(x)$ is the output for sample x, φ_0 is its background expectation, M is the number of predictors, and φ_i is the contribution assigned to predictor i. The explanation is additive in the output space used by the explainer; it should not automatically be interpreted in probability units.

Shapley additive explanations values were evaluated for the 15 predictors at the 81 reference locations. Global importance for a class was summarised by the mean absolute attribution:

$$Importance_i = \frac{1}{N} \sum_{j=1}^{N} |\phi_{ij}| \tag{16}$$

where N is the number of reference records, and φ_{ij} is the contribution of predictor i at record j. Mean absolute attribution measures magnitude rather than direction or causality. Each displayed importance vector was rescaled by subtracting its minimum and dividing by its range.

Selected predictors are presented at two-decimal precision. Although within-method normalisation supports descriptive rank comparisons, it does not place different importance measures on a common inferential scale.

Table 3. Principal reported model hyperparameters

Model	Main hyperparameters
XGBoost	Trees = 600, depth = 6, LR = 0.03
RF	Trees = 500, depth = 10
LightGBM	Trees = 400, leaves = 20
AttentionMLP	Hidden layers = 128–64–32
CNN-1D	Filters = 32–64
ResMLP	Two residual blocks
TabNet	Nd = 16, steps = 3

Abbreviations: CNN-1D: One-dimensional convolutional neural network; LR: learning rate; MLP: Multilayer perceptron; Nd: TabNet decision-feature dimension; RF: Random forest.

3. Results

3.1. Diagnostic agreement with the reference classification

Table 4 and Figure 2 report diagnostic agreement under the three evaluation procedures described in Section 2.3.4. The tree ensemble achieved an accuracy of 0.7531 and macro-F1 of 0.7444 after spatially excluded classifier fitting with global preprocessing. AttentionMLP recorded an apparent accuracy of 0.8272 and macro-F1 of 0.8231 after global fitting. The full ensemble yielded an accuracy of 0.7654 and macro-F1 of 0.7594 when refitted tree probabilities were combined with global neural predictions.

One-dimensional convolutional neural network and ResMLP yielded apparent macro-F1 values of 0.7496 and 0.7298, respectively. TabNet used a batch size of 128 for a 74-record internal training subset. Its score is treated descriptively because effective optimisation depends on the handling of incomplete batches.

The AttentionMLP confusion matrix showed 67 correct classifications out of 81 reference points (Figure 3), corresponding to an accuracy of 0.8272. Class-specific F1-scores were 0.90 for low, 0.77 for moderate, and 0.80 for high. Five high-class points were predicted as moderate, whereas three moderate-class points were predicted as high.

The receiver operating characteristic curves summarise score discrimination at the reference points (Figure 4). AttentionMLP achieved apparent one-versus-rest area under the curve values of 0.961, 0.935, and 0.946 for low, moderate, and high, respectively.

3.2. Diagnostic agreement of component classifications

The component XGBoost classifiers closely reproduced

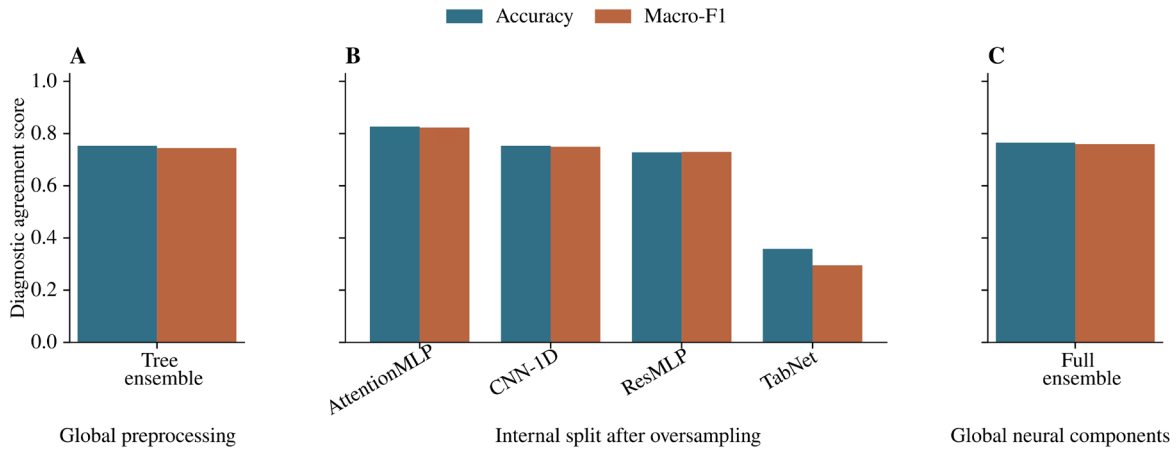

Figure 2. Diagnostic agreement scores grouped by evaluation procedure. (A) Tree-classifier fitting with nominal 2 km exclusion, following global preprocessing. (B) Globally fitted neural networks, whose internal splits followed oversampling. (C) Mixed full-ensemble evaluation combining the refitted tree ensemble and global neural components. Bars show accuracy and macro-F1; other metrics are given in Table 4.
Abbreviations: CNN-1D: One-dimensional convolutional neural network; MLP: Multilayer perceptron.

Table 4. Diagnostic agreement relative to the assigned reference classification

Configuration	Accuracy	Balanced accuracy	Macro-F1	Weighted F1	Cohen's kappa	MCC
Tree ensemble (E)	0.7531	0.7437	0.7444	0.7538	0.6276	0.6299
AttentionMLP	0.8272	0.8237	0.8231	0.8288	0.7400	0.7407
CNN-1D	0.7531	0.7488	0.7496	0.7548	0.6281	0.6291
ResMLP	0.7284	0.7325	0.7298	0.7334	0.5948	0.6026
TabNet	0.3580	0.3325	0.2953	0.3100	0.0000	0.0000
Full ensemble (M)	0.7654	0.7596	0.7594	0.7663	0.6467	0.6468

Notes: Tree ensemble (E) represents classifier refitting after nominal 2 km exclusion, with global preprocessing. Neural-network rows represent global fitting after an internal split of the oversampled sample. Full ensemble (M) represents excluded tree fitting combined with global neural predictions. TabNet is subject to the batch-configuration limitation discussed in Section 3.1.
Abbreviations: CNN-1D: One-dimensional convolutional neural network; MCC: Matthews correlation coefficient; MLP: Multilayer perceptron.

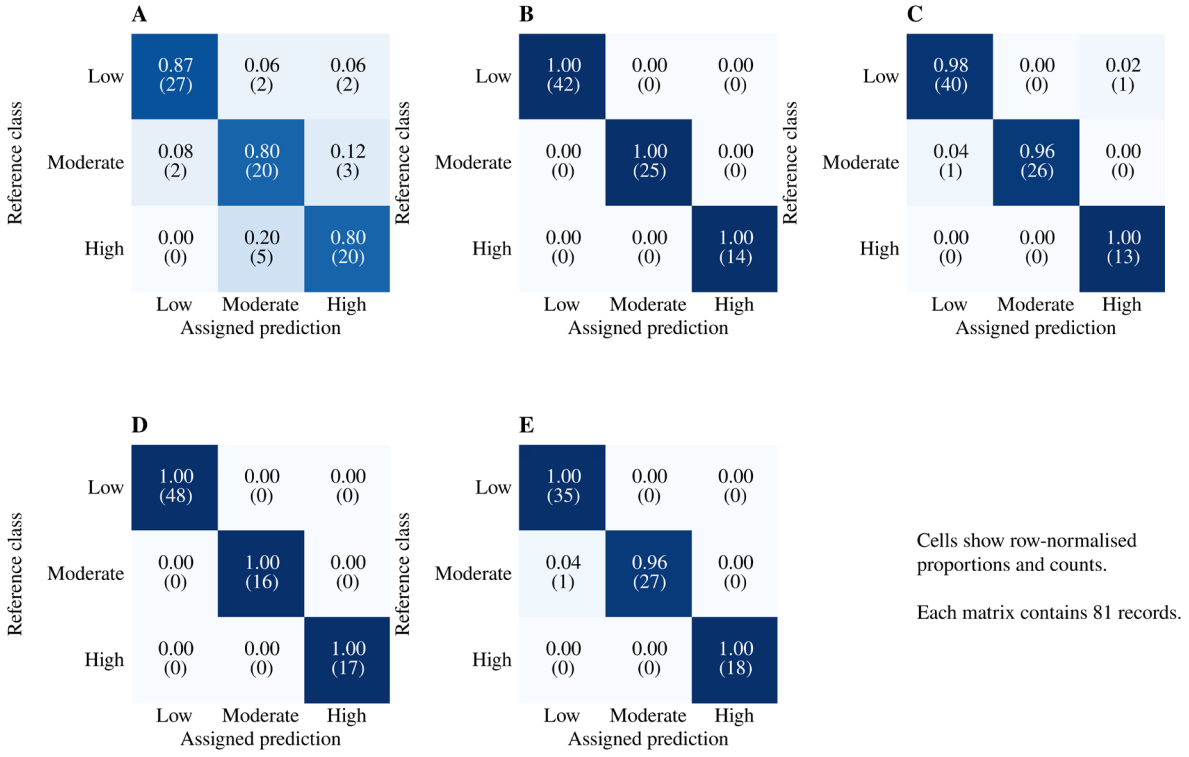

Figure 3. Agreement with assigned reference classes for AttentionMLP and four extreme gradient boosting components. (A) Attention multilayer perceptron. (B) Flood component. (C) Thermal component. (D) Atmospheric component. (E) Compound component. Cells contain row-normalised proportions and counts. Each matrix totals 81 records. Neural-model agreement is apparent, and component agreement is assessed on records used in fitting. Component-specific class counts differ from the composite classes in Table 2.

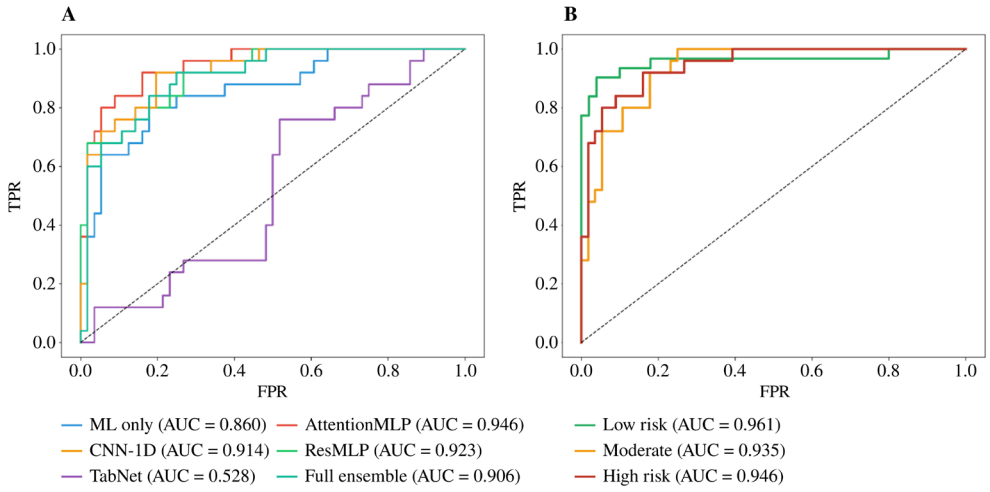

Figure 4. One-versus-rest receiver operating characteristic diagnostics at reference records. (A) combines configurations assessed under different procedures; (B) shows apparent class discrimination by AttentionMLP. Area under the curve values are descriptive.
Abbreviations: AUC: Area under the curve; FPR: False positive rate; TPR: True positive rate.

their assigned component labels (Table 5; Figure 3). Flood and atmospheric components correctly classified all 81 evaluated records, while thermal and compound components classified 79 and 80 records accurately, respectively. Because these records were used in fitting, the values represent resubstitution diagnostics.

The component labels exhibited different class distributions. Figure 3 yields low/moderate/high counts of 42/25/14 for flood, 41/27/13 for thermal, 48/16/17 for atmospheric, and 35/28/18 for compound pressure. These counts represent the reference-class totals obtained by summing all cells in each row of Figure 3, whereas the diagonal entries indicate only the correctly classified records. These totals distinguish the component classifications from the composite labels.

Table 5. Resubstitution diagnostics for extreme gradient boosting component classifications on 81 reference records

Component	Accuracy	Macro-F1	Cohen's kappa	MCC
Flood	1.000	1.000	1.000	1.000
Thermal	0.975	0.973	0.960	0.960
Atmospheric	1.000	1.000	1.000	1.000
Compound	0.988	0.989	0.981	0.981

Note: Accuracy values agree with the counts in Figure 3.
Abbreviation: MCC: Matthews correlation coefficient.

3.3. Spatial distribution of environmental pressure

The valid predictor mask contained 4,065,551 pixels (Table 6). Low, moderate, and high classes represented 56.75%, 30.54%, and 12.71% of these pixels, respectively, and the mean continuous score was 0.3369. The spatial summary reported 30.56% for the component-score contrast, using an HSI > 0.05 threshold. The continuous map shows higher scores along parts of the urbanised coastal belt and lower scores across much of the forested eastern uplands (Figure 5).

Table 6. Descriptive summaries of valid mapped pixels

Spatial metric	Value
Total valid pixels	4,065,551
Low pressure (class 0)	56.75%
Moderate pressure (class 1)	30.54%
High pressure (class 2)	12.71%
Mean ensemble score	0.3369
Pixels above the score-contrast threshold (HSI > 0.05)	30.56%

Note: Class fractions utilise argmax; the score-contrast fraction uses HSI > 0.05.
Note: HSI denotes the component-score contrast defined in Section 2.3.4.

Higher scores occurred around the Mormugao–Vasco–Zuari corridor, Panaji and its surroundings, and parts of the Margao–Navelim–Cuncolim corridor. These patterns

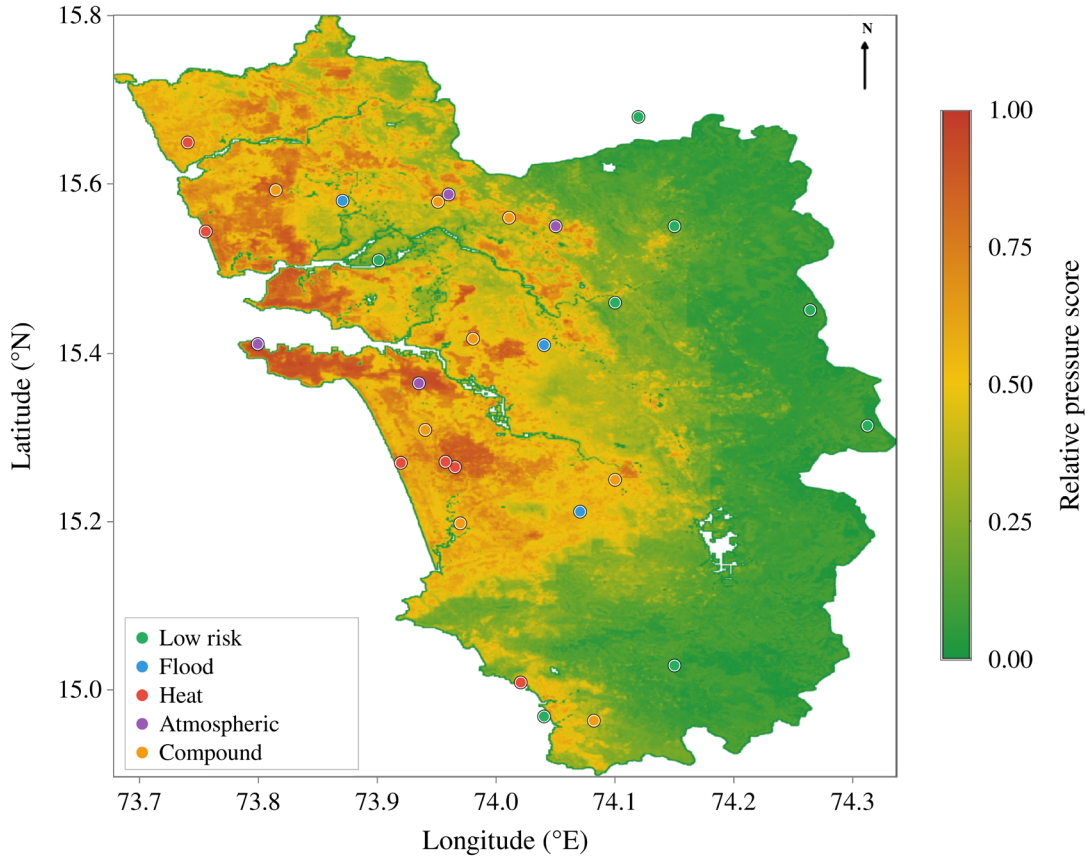

Figure 5. Relative continuous environmental-pressure score from the fitted ensemble. Reference symbols show assigned site categories.

are consistent with the urban and industrial descriptors used in reference-label assignment. Much of the eastern forest belt is mapped as low. This is consistent with its vegetation cover and the reference-label definition. The categorical map summarises the low, moderate, and high classes assigned from the largest ensemble probability (Figure 6). The continuous score is retained separately.

3.4. Combined and individual component patterns

The component maps show differing spatial patterns (Figures 7 and 8). The flood-related component yielded higher scores in parts of the coastal lowlands and estuarine settings, represented by terrain and radar-backscatter indicators. The thermal component recorded higher scores in several built-up areas, including Panaji, Vasco da Gama, Margao, Mapusa, and the Ponda–Marcel belt. Land-surface temperature and vegetation cover describe surface thermal conditions.

The atmospheric-indicator component was higher around several urban and industrial centres. Its predictors included atmospheric columns and AAI, alongside land-use and human-activity indicators. The combined-minus-maximum score contrast varied across Goa (Figure 7). At the reference points, the mean individual score and combined ensemble score yielded a Pearson correlation of approximately 0.937, indicating consistency among these fitted outputs.

3.5. Regional reference-point summaries

The regional graphic summarises reference points assigned to North, Central, and South zones, with mean scores of 0.474, 0.523, and 0.183, respectively (Table 7; Figure 9). The reported high-class shares were 28.6% in the North and 36.0% in the Central zone, while 80.0% of South-zone reference points were low. Higher mean scores in the Central and North reference groups are consistent with their sampled urban and industrial locations.

The three zones were grouped into latitude bands: North included records at or north of 15.55° N; Central included records from 15.20° N to below 15.55° N; and

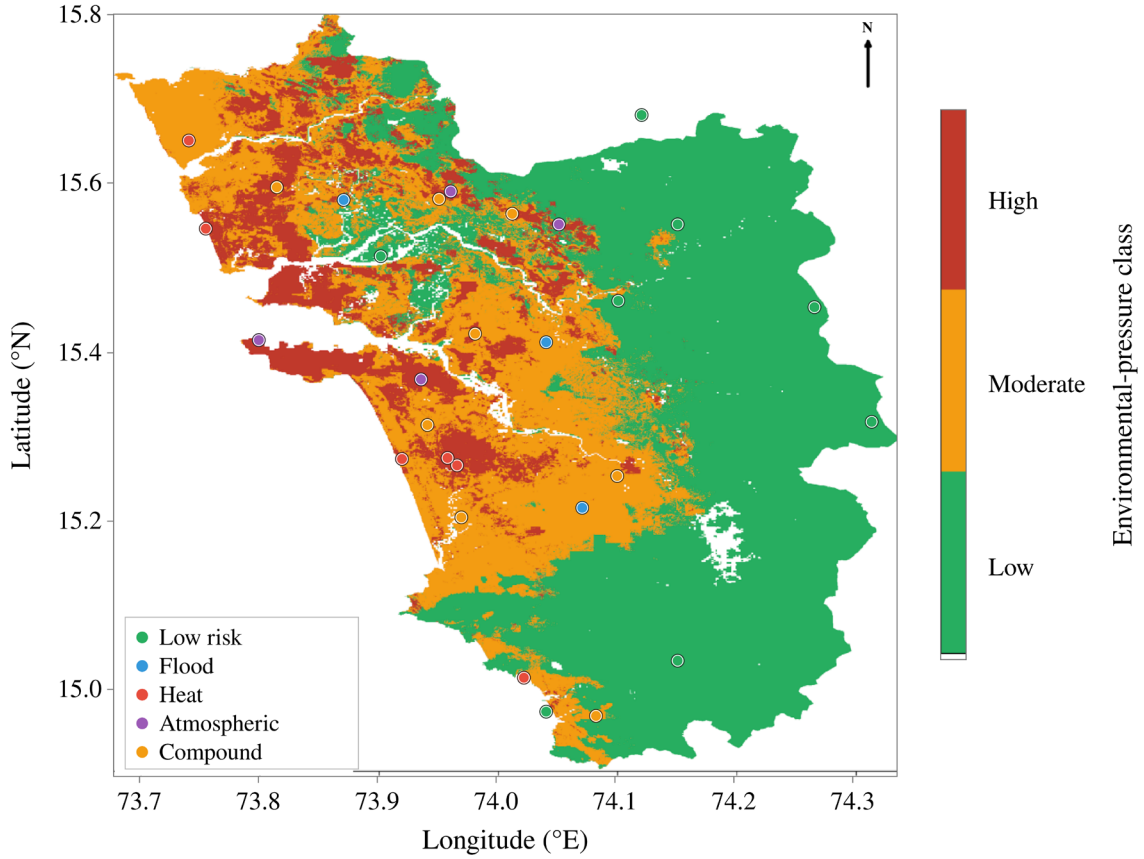

Figure 6. Low, moderate, and high environmental-pressure classes assigned by the largest full-ensemble probability. The model fitting and evaluation procedures are described in Sections 2.3.3 and 2.3.4.

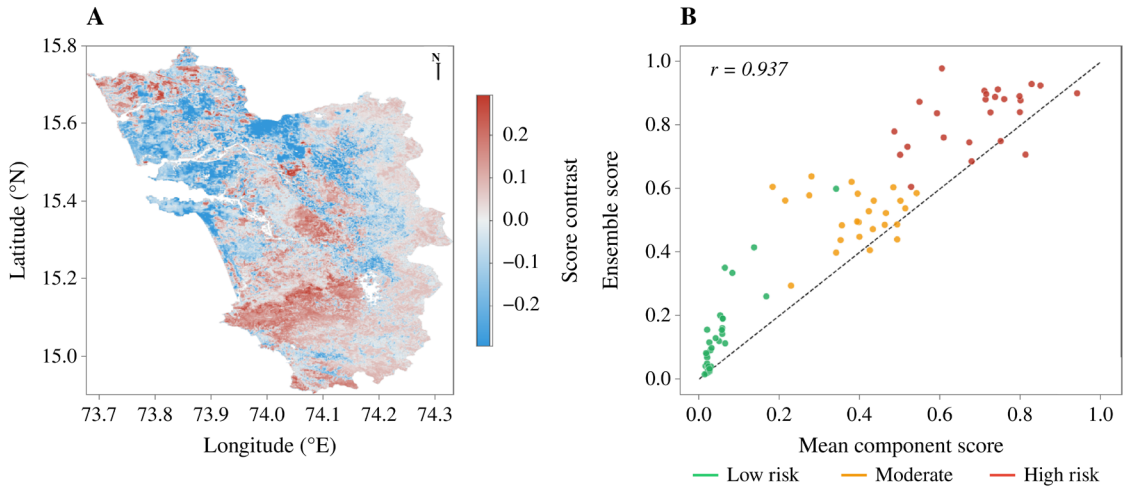

Figure 7. Descriptive comparison among fitted score surfaces. (A) Ensemble score minus the maximum of the four component scores. (B) Mean component score versus ensemble score at reference records. Shared predictors and related labels influence their association.

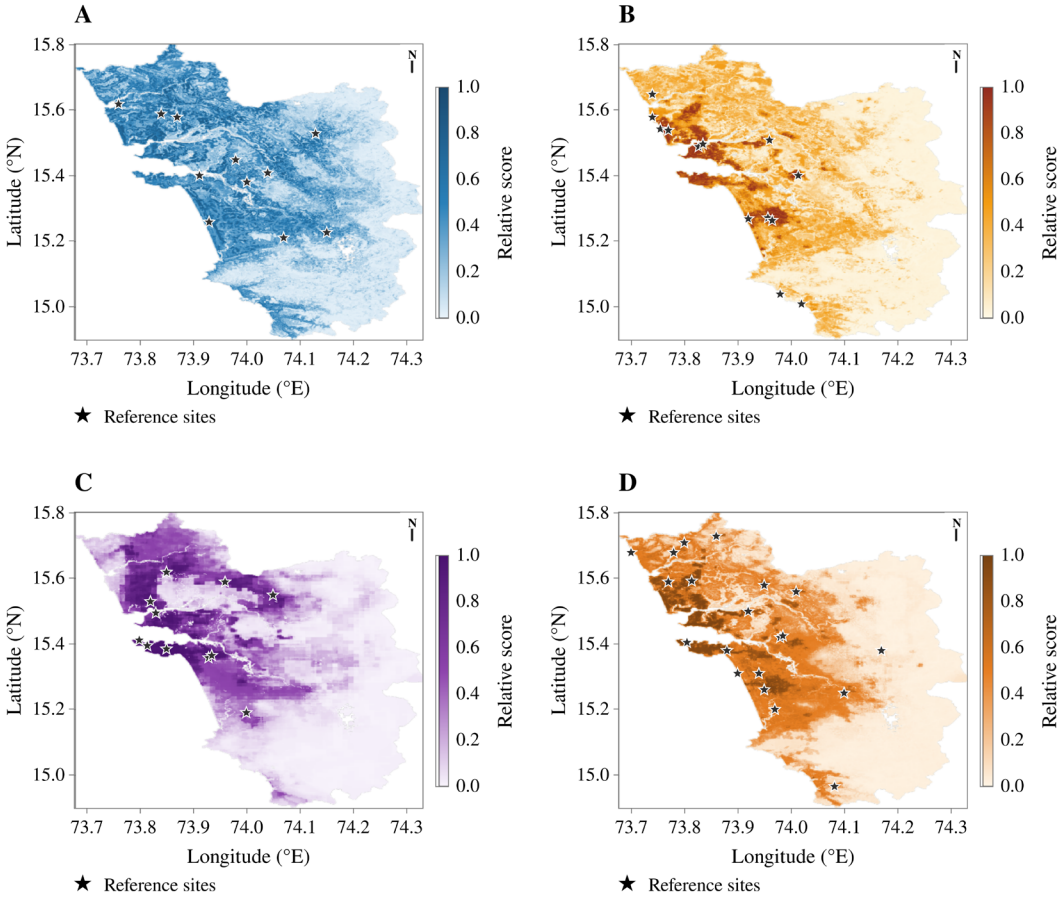

Figure 8. Relative extreme gradient boosting component scores for (A) flood-related conditions, (B) surface thermal conditions, (C) atmospheric composition, and (D) compound environmental pressure. Stars indicate assigned reference-site categories.

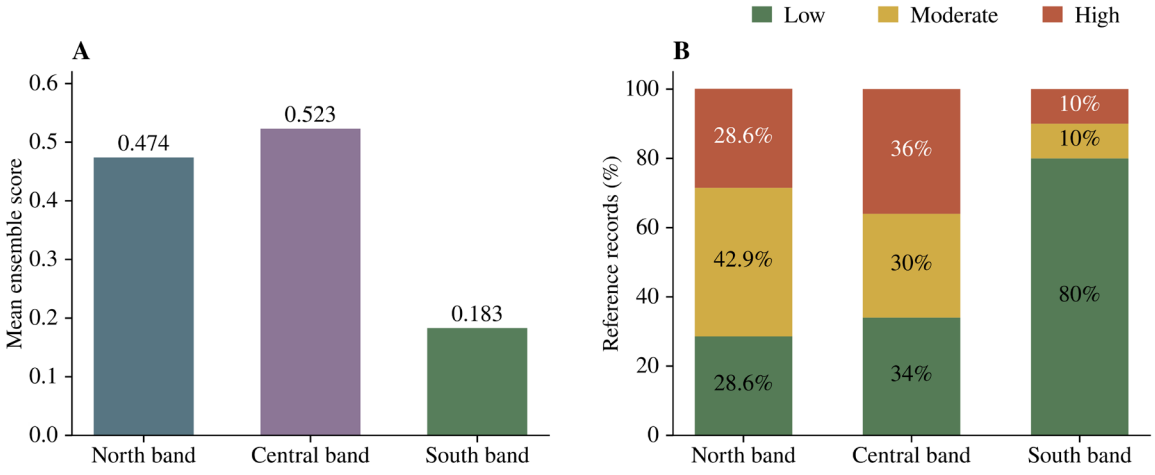

Figure 9. Descriptive summaries of reference records in three latitude bands. (A) Mean fitted ensemble score. (B) Assigned proportions for low, moderate, and high classes. North, Central, and South are analytical bands rather than administrative districts. Percentages were rounded to describe the sampled records.

South included records below 15.20° N. These analytical groups differ from Goa's two administrative districts (Figure 1).

Table 7. Environmental-pressure summaries for reference records in three latitude bands

Reference-point group	Mean ensemble score	High points (%)	Moderate points (%)	Low points (%)
North band	0.474	28.6	**42.9**	28.6
Central band	**0.523**	**36.0**	30.0	34.0
South band	0.183	10.0	10.0	**80.0**

Notes: North represents latitude ≥ 15.55° N; Central represents 15.20° N ≤ latitude < 15.55° N; South represents latitude < 15.20° N. Percentages describe assigned reference classes; rounding can produce totals slightly different from 100%. Bold values indicate column maxima among the three latitude bands.

3.6. Feature attribution

Night-time lights yielded the highest normalised XGBoost and random forest importance and the largest high-class SHAP magnitude (Figures 10 and 11; Table 8). AttentionMLP assigned relatively high weights to AAI, population, and NTLs. The TabNet importance vector was not used for environmental interpretation because of the training-batch limitation.

Although population exhibited high tree-based importance, its high-class SHAP magnitude was smaller than those of land-surface temperature and AAI (Table 8). The distinction reflects differences between tree importance and class-specific attribution. Mean absolute SHAP values do not indicate whether an increase in a predictor raises or lowers the output.

Land-surface temperature and AAI contributed to high-class predictions in the graphical explanations. The optical digital-number ratio also contributed, with its calculation defined in Section 2.2. The beeswarm plots show class-specific feature contributions, rather than a single predictor ranking shared by all classes (Figure 10). NTL was prominent for low and high classes, whereas atmospheric and radar variables were prominent in the moderate-class plot. Correlated predictors and their overlap with label construction can influence both rankings and attribution signs.

The prominence of NTL across the tree and SHAP summaries indicates that several fitted methods recover a similar settlement-related signal. This agreement is expected to some extent, since the methods share the reference records and predictors.

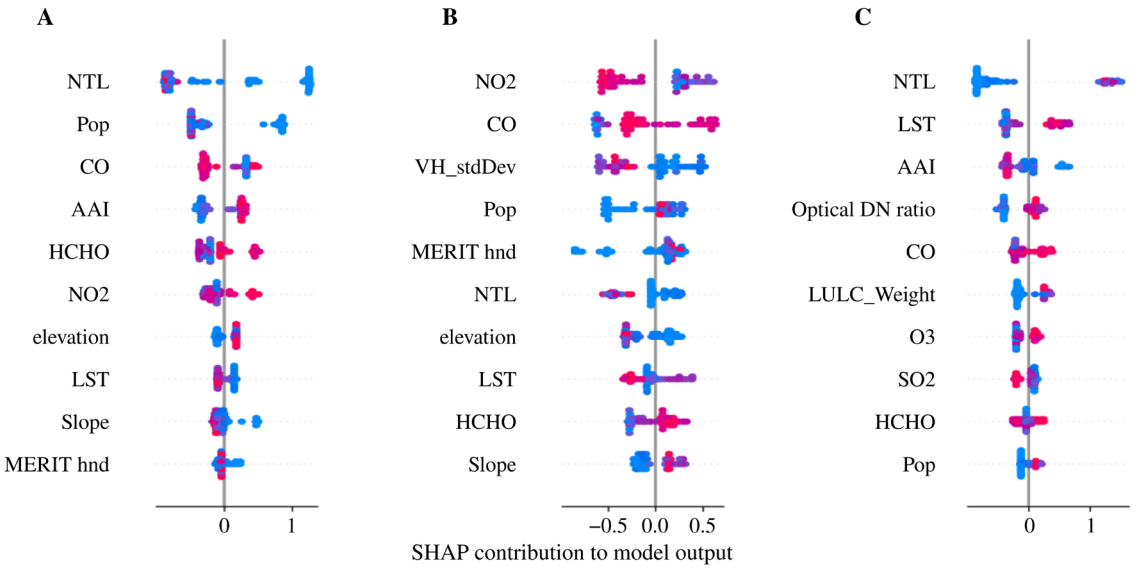

Figure 10. Class-specific SHAP attribution for the globally fitted XGBoost component. (A) Low class. (B) Moderate class. (C) High class. Point positions and colours describe contributions and predictor values at reference records. AAI is the absorbing aerosol index; MERIT hnd is hydrologically adjusted elevation; the optical DN ratio is the unscaled band ratio described in Section 2.2. All variables are described in Table 1.
Abbreviations: LST: Land-surface temperature; NTL: Night-time light; Pop: Population indicator.

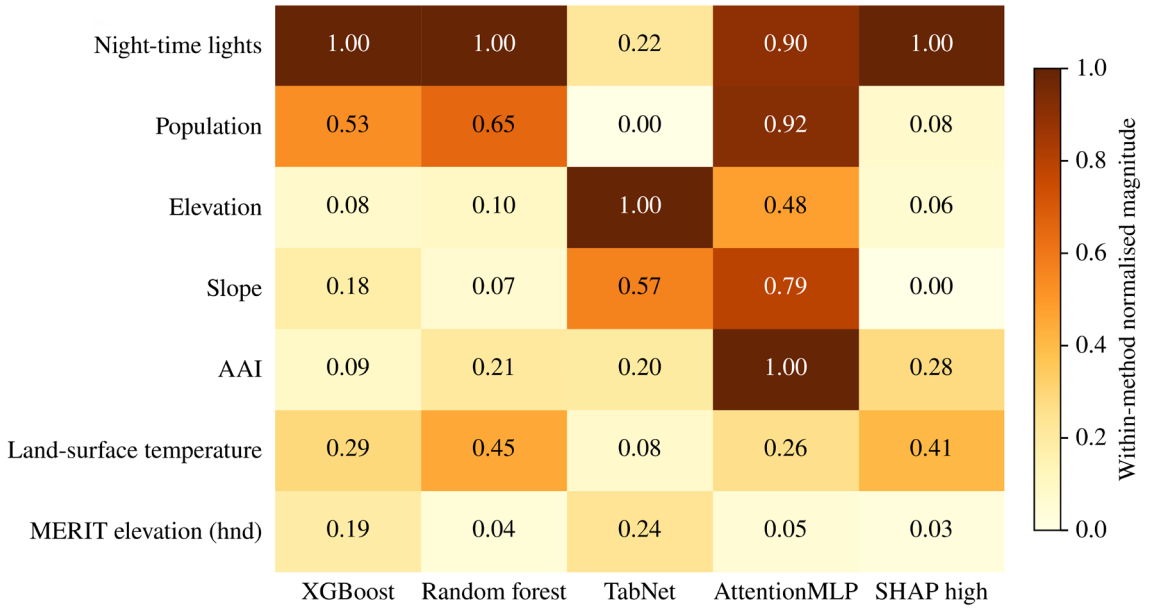

Figure 11. Within-method normalised attribution for the selected predictors in Table 8, displayed at two-decimal precision. SHAP high denotes mean absolute XGBoost attribution for the High class. Magnitudes were scaled separately by method and were not on a common inferential scale. TabNet values are retained as diagnostics subject to the batch limitation.
Abbreviations: AAI: Absorbing aerosol index; SHAP: Shapley additive explanations; XGBoost: Extreme gradient boosting.

Table 8. Within-method normalised attribution for selected predictors, displayed at two-decimal precision in Figure 11

Feature	XGBoost	Random forest	TabNet	AttentionMLP	SHAP high
Night-time lights	**1.00**	**1.00**	0.22	0.90	**1.00**
Population	0.53	0.65	0.00	0.92	0.08
Elevation	0.08	0.10	**1.00**	0.48	0.06
Slope	0.18	0.07	0.57	0.79	0.00
AAI	0.09	0.21	0.20	**1.00**	0.28
LST	0.29	0.45	0.08	0.26	0.41
NO_2	0.08	0.14	0.15	0.22	0.00
MERIT hnd	0.19	0.04	0.24	0.05	0.03

Notes: Each method was minimum–maximum scaled over its full predictor vector. The displayed subset is not renormalised. TabNet values are retained descriptively and are not used for environmental interpretation. Bold values indicate the maximum normalised attribution within each method.
Abbreviations: AAI: Absorbing aerosol index; LST: Land-surface temperature; SHAP: Shapley additive explanations; XGBoost: Extreme gradient boosting.

3.7. Model diagnostics and environmental implications

Training losses generally declined, whereas internal validation losses flattened or increased as fitting proceeded (Figure 12). Divergence was most pronounced for ResMLP, and was also visible for AttentionMLP and CNN-1D. The three networks restored the weights from the minimum-validation-loss epoch.

Entropy was highest at intermediate continuous scores and generally declined towards the score extremes (Figure 13). This pattern describes dispersion among the fitted class probabilities.

The reliability diagrams depart from the identity line (Figure 14). Low predicted probabilities frequently correspond to observed frequencies near zero, while several higher-probability bins have observed frequencies near one. Points above the diagonal indicate underprediction and points below it indicate overprediction. The curves use eight quantile-based bins and full-ensemble predictions at the reference records. They describe apparent calibration within this sample.

Pearson correlations among the assigned ordinal component labels were 0.27 for flood–thermal, 0.23 for flood–atmospheric, and 0.69 for thermal–atmospheric classifications (Figure 15). Correlations with the compound label were 0.67 for flood and 0.77 for thermal and atmospheric labels. Twenty-six records had all four labels at moderate or high, while 31 had all four labels at low.

The hotspot display ranks selected reference locations by predicted high-class probability (Figure 16). Urban and industrial sites appear among those with high scores, consistent with the label definitions. The accompanying elevation–temperature scatter shows that many higher-score observations occur at low elevations and higher surface temperatures. These spatial associations identify locations for further field assessment.

4. Discussion

The framework provides a common spatial representation of atmospheric, surface, and human-activity indicators across Goa. It identifies contrasts between urban and industrial corridors and the forested eastern landscape, and relates those contrasts to the fitted environmental-pressure classification. The mapped contrasts provide a spatial basis for examining several environmental pressures together. Their practical value lies in directing observations towards locations where urban activity, thermal conditions, and flood-related terrain coincide.

Comparisons among tabular learning methods depend on reference-sample size, predictor structure, and evaluation design. Recent work shows that different tabular approaches can be useful under different conditions.[64,74-76] In this study, AttentionMLP showed high agreement with the assigned classes, while the different fitting procedures require the model scores to be interpreted separately.

The atmospheric predictors describe column properties.

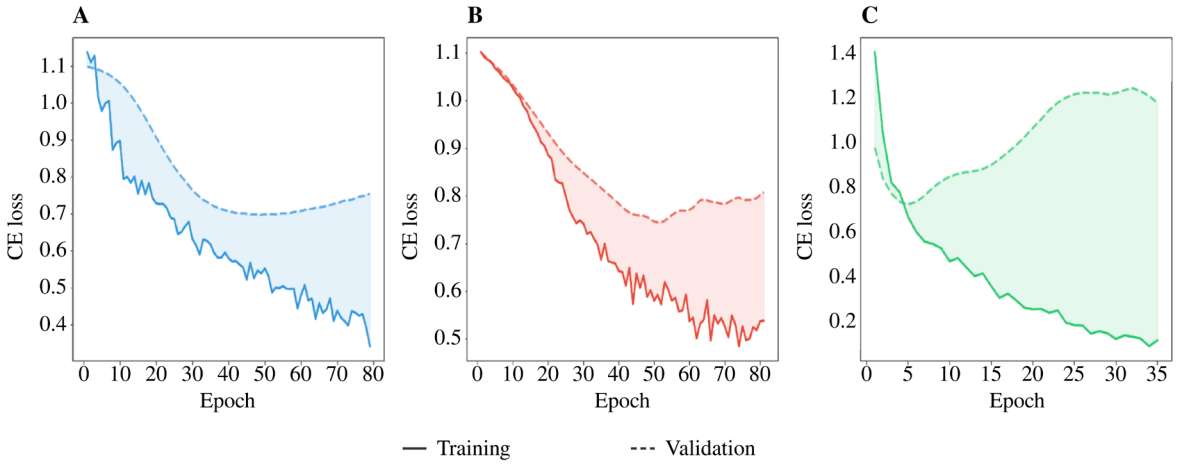

Figure 12. Training and internal validation losses for (A) attention multilayer perceptron, (B) one-dimensional convolutional neural network, and (C) residual multilayer perceptron. The internal split followed global preprocessing and oversampling. The curves describe optimisation and checkpoint selection.
Abbreviation: CE: Cross-entropy.

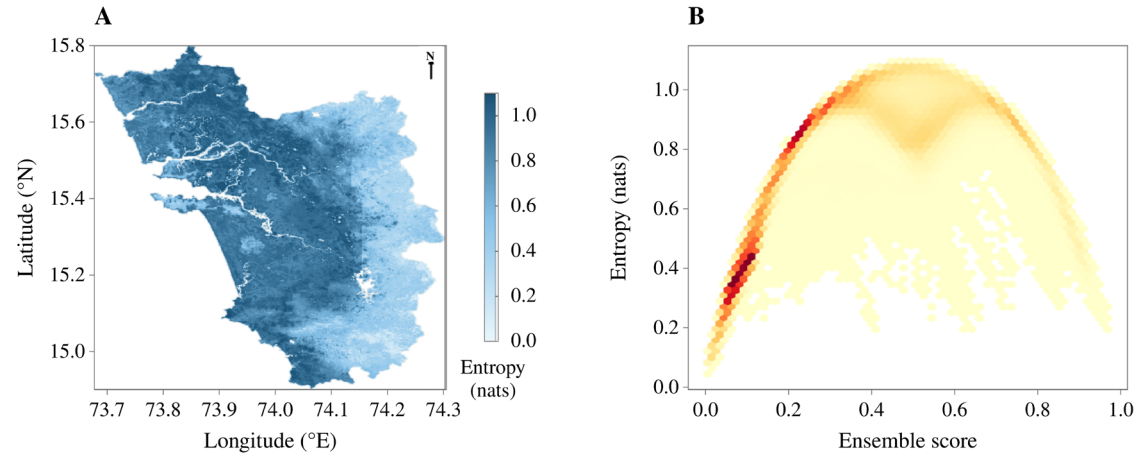

Figure 13. Class-probability dispersion on the fitted map. (A) Entropy of fitted class probabilities on the predictor grid and (B) its relationship with the continuous ensemble score. Entropy is expressed in nats; its theoretical maximum for three classes is ln(3).

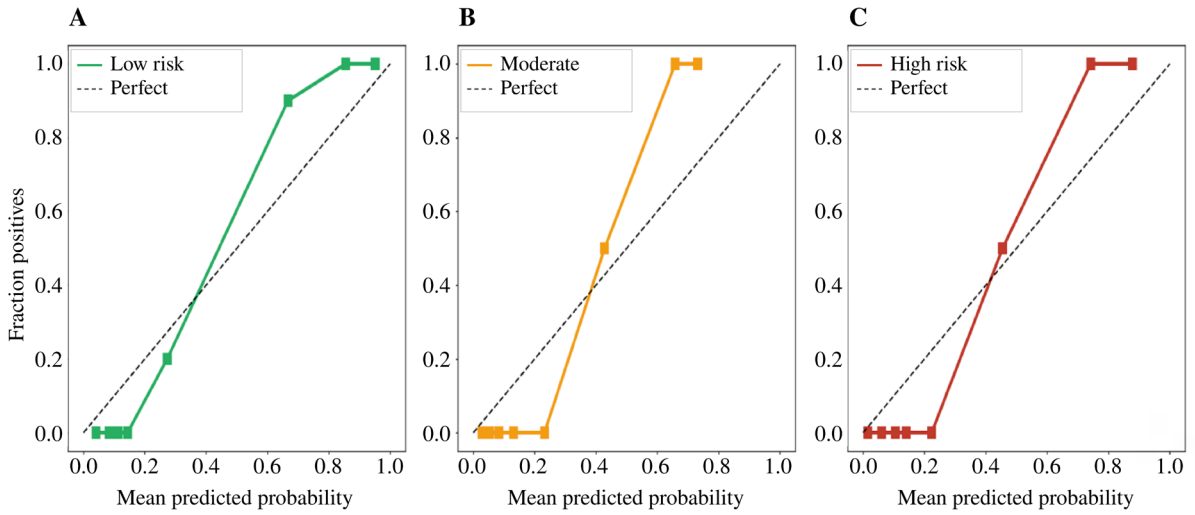

Figure 14. Apparent reliability diagrams for the globally fitted full ensemble at reference records. (A) Low class. (B) Moderate class. (C) High class. The identity line denotes equality between predicted probabilities and reference frequencies.

In particular, AAI is derived from contrasts in ultraviolet reflectance and is sensitive to absorbing aerosols such as smoke and dust.[48] It does not represent surface PM2.5, and total-column ozone does not reflect ground-level ozone exposure. The thermal component similarly describes land-surface conditions rather than physiological heat stress. Although these distinctions limit direct health and regulatory interpretations, they do not prevent the predictors from contributing to a relative environmental classification.

Night-time light was the strongest consistent feature across the tree and SHAP summaries. Previous work has utilised night-time radiance to describe settlement, economic activity, and electricity use.[77-80] In this study, it captures settlement-related information also represented in the reference classes. Class-specific SHAP attribution assigns greater high-class importance to land-surface temperature and AAI than to population.

Night-time light is affected by lighting practices, non-residential light sources, blooming, and the processing of dim or background signals.[77-80] The coastal and urban corridors identified here are candidates for targeted monitoring of surface air quality, thermal conditions, and flood history. The forested eastern areas provide a

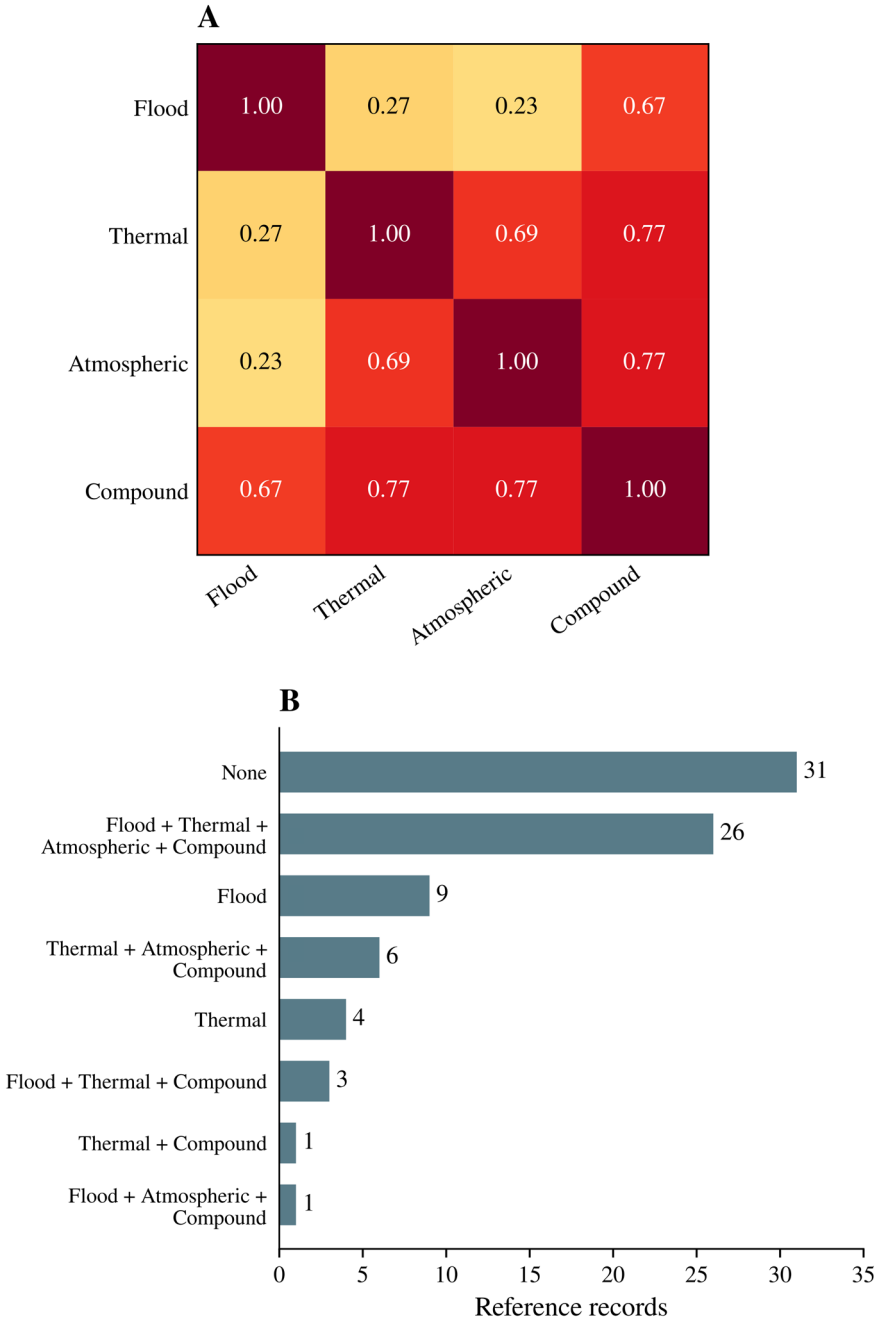

Figure 15. Relationships among assigned component labels at reference records. (A) Pearson correlations of the ordinal labels. (B) Counts of category combinations for which labels are moderate or high; None denotes all four low labels. Shared label construction influences both panels.

contrasting setting for continued ecological assessment.

The combined-minus-maximum score contrast describes differences among model surfaces. It complements the spatial comparison of individual components, while physical multi-hazard studies examine interactions among processes and their consequences.[73,81,82] The threshold and smoothing procedures used here are specified in Section 2.3.4.

The evaluation has three principal limitations. Global preprocessing and oversampling before the neural internal

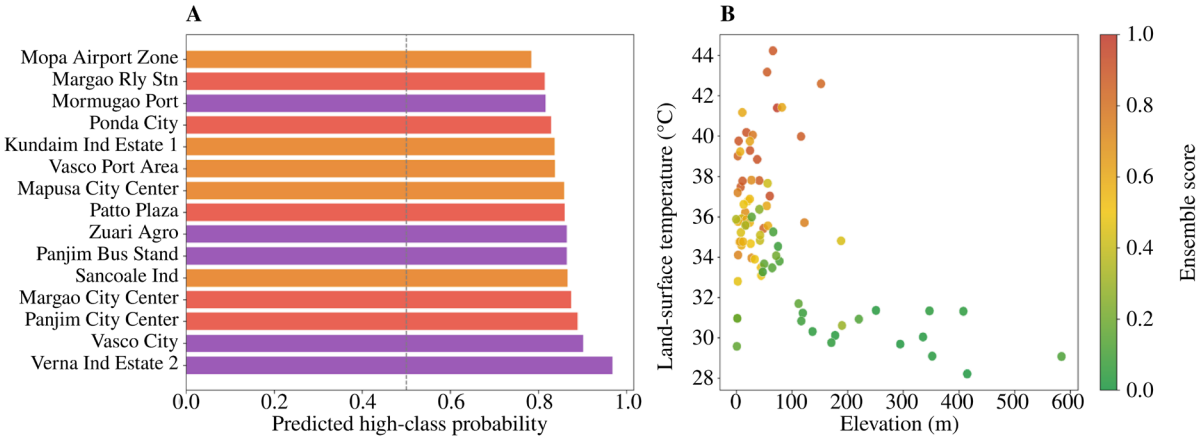

Figure 16. Descriptive characterisation of selected reference records. (A) Fitted high-class probabilities. (B) Elevation and land-surface temperature, coloured by fitted ensemble score.

split allow information leakage, and the neural models and full ensemble were not evaluated by complete spatial cross-validation. In addition, the reference classes partly share information with the predictors, and independent outcome validation was not undertaken. Consequently, the reported scores describe agreement with the assigned classification. Predictive evaluation would require preprocessing, resampling, and model fitting within spatial training folds, along with independently defined outcomes.[58,59,66] The present maps are intended for relative environmental screening.

5. Conclusion

This study combines three tree-based algorithms and four neural-network architectures to map chronic environmental pressure in Goa. The fitted classification distinguishes parts of the urban and industrial corridors from much of the forested eastern landscape. The high class accounts for 12.71% of valid mapped pixels. NTLs have the largest high-class SHAP contribution, followed by land-surface temperature and AAI among the predictors considered. The ensemble and component maps provide complementary descriptions of these spatial patterns. The integrated representation of atmospheric, thermal, and land-surface indicators identifies candidate locations for further environmental observation. It provides a basis for directing field assessment towards the coastal and urban corridors where several indicators coincide.

Acknowledgments

The authors acknowledge Kaggle for GPU computing resources used in neural-network training. Geospatial products were accessed through Google Earth Engine from Copernicus/ESA Sentinel-5P and Sentinel-1, USGS Landsat and SRTM, NOAA/NASA VIIRS, WorldPop, MERIT Hydro, and ESA WorldCover.

Funding

None.

Conflict of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author contributions

Conceptualization: Madhusudhan Das, Shruti Kanga
Formal analysis: Divyansh Kumar Pandey, Madhusudhan Das, Gowhar Meraj
Investigation: Saurabh K Gupta, Madhusudhan Das, Divyansh Kumar Pandey
Methodology: Shruti Kanga, Gowhar Meraj, Suraj Kumar Singh
Writing–original draft: Bhartendu Sajan, Madhusudhan Das
Writing–review & editing: Gowhar Meraj, Suraj Kumar Singh

Availability of data

The datasets generated and/or analysed during the current study are available from the corresponding author upon reasonable request.

References

1. Intergovernmental Panel on Climate Change (IPCC). *Climate Change 2021 – The Physical Science Basis: Working Group I Contribution to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge, England: Cambridge University Press; 2021.

   doi: 10.1017/9781009157896

2. Patel JB, Raval Z. The impacts of urbanization on ecological systems: a comprehensive study of the complex challenges arising from rapid urban growth. *Res Rev J Indian Knowl Syst*. 2024;1(1):1-10.

   doi: 10.31305/rrjiks.2024.v1.n1.001

3. Rimal B, Baral H, Stork N, Paudyal K, Rijal S. Growing city and rapid land use transition: assessing multiple hazards and risks in the Pokhara Valley, Nepal. *Land*. 2015;4(4):957-978.

   doi: 10.3390/land4040957

4. Dutta D, Srinivas VV, Bala G. Disparity in spatiotemporal variability and risk of compound coastal extremes between India's East-West coasts. *NPJ Clim Atmos Sci*. 2025;8(1):168.

   doi: 10.1038/s41612-025-01045-5

5. Kundzewicz ZW, Kanae S, Seneviratne SI, *et al*. Flood risk and climate change: global and regional perspectives. *Hydrol Sci J*. 2014;59(1):1-28.

   doi: 10.1080/02626667.2013.857411

6. Almar R, Ranasinghe R, Bergsma EWJ, *et al*. A global analysis of extreme coastal water levels with implications for potential coastal overtopping. *Nat Commun*. 2021;12(1):3775.

   doi: 10.1038/s41467-021-24008-9

7. Balica SF, Wright NG, van der Meulen F. A flood vulnerability index for coastal cities and its use in assessing climate change impacts. *Nat Hazards*. 2012;64(1):73-105.

   doi: 10.1007/s11069-012-0234-1

8. Cazenave A, Cozannet GL. Sea level rise and its coastal impacts. *Earths Future*. 2014;2(2):15-34.

   doi: 10.1002/2013EF000188

9. Wang Q, Hou J, Li J, Guo H. A multihazard assessment framework integrating single and coupled scenario analysis for droughts, floods, and landslides. *Geomatics Nat Hazards Risk*. 2026;17(1):2620070.

   doi: 10.1080/19475705.2026.2620070

10. Wenzel T, van Westen C, Sunil M, *et al*. Towards a practical multi-hazard interrelation classification: implications for assessing their impacts. *Nat Hazards*. 2026;122(2):82.

    doi: 10.1007/s11069-025-07901-0

11. Marlier ME, Connolly R, Ma Y, Jerrett M, Benmarhnia T. Compound climate exposures and public health: a synthesis of research and future directions. *Annu Rev Public Health*. 2026;47(1):305-323.

    doi: 10.1146/annurev-publhealth-090924-113926

12. Jäger WS, de Ruiter MC, Tiggeloven T, Ward PJ. What can we learn about multi-hazard impacts from global disaster records? *Nat Hazards Earth Syst Sci*. 2025;25(8):2751-2769.

    doi: 10.5194/nhess-25-2751-2025

13. Hafezi M, Sahin O, Stewart R, Mackey B. Creating a novel multi-layered integrative climate change adaptation planning approach using a systematic literature review. *Sustainability (Basel)*. 2018;10(11):4100.

    doi: 10.3390/su10114100

14. Nghia Nguyen H, Quan Le Q, Dung Nguyen V, *et al*. Assessment of coastal inundation triggered by multiple drivers in Ca Mau Peninsula, Vietnam. *Nat Hazards Earth Syst Sci*. 2025;25(10):4227-4246.

    doi: 10.5194/nhess-25-4227-2025

15. Barnard PL, Erikson LH, Foxgrover AC, *et al*. Dynamic flood modeling essential to assess the coastal impacts of climate change. *Sci Rep*. 2019;9(1):4309.

    doi: 10.1038/s41598-019-40742-z

16. Sun H, Zhang X, Ruan X, Jiang H, Shou W. Mapping compound flooding risks for urban resilience in coastal zones: a comprehensive methodological review. *Remote Sens (Basel)*. 2024;16(2):350.

    doi: 10.3390/rs16020350

17. Akosah S, Gratchev I. Systematic review of post-wildfire landslides. *GeoHazards*. 2025;6(1):12.

    doi: 10.3390/geohazards6010012

18. Moftakhari H, AghaKouchak A. Increasing exposure of energy infrastructure to compound hazards: cascading wildfires and extreme rainfall. *Environ Res Lett*. 2019;14(10):104018.

    doi: 10.1088/1748-9326/ab41a6

19. Razavi-Termeh SV, Sadeghi-Niaraki A, Ali F, Pradhan B, Choi SM. Optimizing ensemble learning for satellite-based multi-hazard monitoring and susceptibility assessment of landslides, land subsidence, floods, and wildfires. *Sci Rep*. 2025;15(1):30968.

    doi: 10.1038/s41598-025-15381-2

20. Arshad S, Ahmad SR, Abbas S, Asharf A, Siddiqui NA, Islam Zu. Quantifying the contribution of diminishing green spaces and urban sprawl to urban heat island effect in a rapidly urbanizing metropolitan city of Pakistan. *Land Use Policy*. 2022;113:105874.

    doi: 10.1016/j.landusepol.2021.105874

21. Rahaman ZA, Kafy AA, Saha M, *et al*. Assessing the impacts of vegetation cover loss on surface temperature, urban heat island and carbon emission in Penang city, Malaysia. *Build Environ*. 2022;222:109335.

doi: 10.1016/j.buildenv.2022.109335

22. Grigoraș G, Urițescu B. Land use/land cover changes dynamics and their effects on surface urban heat island in Bucharest, Romania. *Int J Appl Earth Obs Geoinf*. 2019;80:115-126.

doi: 10.1016/j.jag.2019.03.009

23. He JF, Liu JY, Zhuang DF, Zhang W, Liu ML. Assessing the effect of land use/land cover change on the change of urban heat island intensity. *Theor Appl Climatol*. 2007;90(3-4):217-226.

doi: 10.1007/s00704-006-0273-1

24. Afandi GE, Moustafa A, Ibrahim S, Irfan M. An in-depth analysis of approaches for forecasting extreme rainfall events leveraging artificial intelligence, machine learning, and deep learning techniques. *Earth Syst Environ*. 2026;10(4):5073-5104.

doi: 10.1007/s41748-025-00813-w

25. Pugliese Viloria AdJ, Folini A, Carrion D, Brovelli MA. Hazard susceptibility mapping with machine and deep learning: a literature review. *Remote Sens (Basel)*. 2024;16(18):3374.

doi: 10.3390/rs16183374

26. Kumar R, Bhanu M, Mendes-Moreira J, Chandra J. Spatio-temporal predictive modeling techniques for different domains: a survey. *ACM Comput Surv*. 2025;57(2):1-42.

doi: 10.1145/3696661

27. Li W, Hsu CY. GeoAI for large-scale image analysis and machine vision: recent progress of artificial intelligence in geography. *ISPRS Int J Geoinf*. 2022;11(7):385.

doi: 10.3390/ijgi11070385

28. Feng L, Zhang M, Mao Y, *et al*. Convolutional neural network-based deep learning for landslide susceptibility mapping in the Bakhtegan watershed. *Sci Rep*. 2025;15(1):13250.

doi: 10.1038/s41598-025-96748-3

29. Amato F, Guignard F, Robert S, Kanevski M. A novel framework for spatio-temporal prediction of environmental data using deep learning. *Sci Rep*. 2020;10(1):22243.

doi: 10.1038/s41598-020-79148-7

30. Lynda D, Logeswari G, Tamilarasi K, Rakesh S. Hybrid Bayesian deep learning model for predicting urban heat island intensity in African cities. *Sci Rep*. 2025;15(1):31280.

doi: 10.1038/s41598-025-13492-4

31. Das A, Rad P. Opportunities and challenges in explainable artificial intelligence (XAI): a survey. arXiv. Preprint posted online 2020. Updated June 23, 2020. Accessed September 21, 2026.

doi: 10.48550/arXiv.2006.11371

32. Geng H, Wang W, Liu J, Benson D. Landslide susceptibility modeling based on SHAP interpretability and ensemble learning: a case study in Fuyuan County, Southwest China. *Front Earth Sci*. 2025;13:1731872.

doi: 10.3389/feart.2025.1731872

33. Zhang J, Ma X, Zhang J, *et al*. Insights into geospatial heterogeneity of landslide susceptibility based on the SHAP-XGBoost model. *J Environ Manage*. 2023;332:117357.

doi: 10.1016/j.jenvman.2023.117357

34. Zhu H, Zhang X, Wei G, *et al*. A method for better mapping of susceptibility to thaw hazards in data-scarce cold regions. *Remote Sens Environ*. 2026;337:115338.

doi: 10.1016/j.rse.2026.115338

35. Gomez Rave DV, Scolobig A, del Jesus M. Review article: rethinking preparedness for coastal compound flooding: insights from a systematic review. *Nat Hazards Earth Syst Sci*. 2025;25(10):3977-4002.

doi: 10.5194/nhess-25-3977-2025

36. Haer T, de Ruiter M. A multi-hazard perspective on equitable adaptation and how to assess it. *PLOS Clim*. 2024;3(12):e0000521.

doi: 10.1371/journal.pclm.0000521

37. Antofie TE, Luoni S, Tilloy A, *et al*. Spatial identification of regions exposed to multi-hazards at the pan-European level. *Nat Hazards Earth Syst Sci*. 2025;25(1):287-304.

doi: 10.5194/nhess-25-287-2025

38. Hochrainer-Stigler S, Šakić Trogrlić R, Reiter K, *et al*. Toward a framework for systemic multi-hazard and multi-risk assessment and management. *iScience*. 2023;26(5):106736.

doi: 10.1016/j.isci.2023.106736

39. Nalmpant-Sarikaki DM, Theocharis AI, Koukouzas NC, Zevgolis IE. A framework for effective multi-hazard risk assessment in post-mining areas. *Safety (Basel)*. 2025;11(1):18.

doi: 10.3390/safety11010018

40. Damalas A, Mettas C, Evagorou E, *et al*. Development and implementation of a DECATASTROPHIZE platform and tool for the management of disasters or multiple hazards. *Int J Disaster Risk Reduct*. 2018;31:589-601.

doi: 10.1016/j.ijdrr.2018.05.011

41. Gacu J, Kantoush S, Candelario R, *et al*. Integrated multi-hazard risk assessment under compound disasters using analytical hierarchy process (AHP). *Heliyon*. 2025;11(11):e43173.

doi: 10.1016/j.heliyon.2025.e43173

42. Houssein EH, Dirar M, Ali AA, Mohamed WM, Hamed K. Advances in artificial intelligence for climate science methods applications and evaluation challenges. *Discov Artif Intell*. 2026;6(1):441.

doi: 10.1007/s44163-026-01262-8

43. Llavina Pascual X, Pablos Cascallar M, Planell i Calle M, Puigdollers Rodés E, Solé Carbonell C, Vila i Casau A. *Iron Mining in Goa (India): An Interdisciplinary Study* [bachelor's thesis]. Barcelona, Spain: Universitat Autònoma de Barcelona; 2013. Accessed September 21, 2026. https://ddd.uab.cat/pub/trerecpro/2013/hdl_2072_223218/PFC_IronMiningGoa.pdf

44. Nigam R, Luis AJ, Gagnon AS, Vaz E, Damásio B, Kotha M. Assessing coastal vulnerability at the village level using a robust framework, the example of Canacona in South Goa, India. *iScience*. 2024;27(4):109129.

doi: 10.1016/j.isci.2024.109129

45. Saha M, Naik A, Desai A, *et al*. Microplastics in seafood as an emerging threat to marine environment: a case study in Goa, west coast of India. *Chemosphere*. 2021;270:129359.

doi: 10.1016/j.chemosphere.2020.129359

46. Department of Information and Publicity. About Goa. Government of Goa. Accessed September 21, 2026. https://dip.goa.gov.in/about-goa/

47. Gorelick N, Hancher M, Dixon M, Ilyushchenko S, Thau D, Moore R. Google Earth Engine: planetary-scale geospatial analysis for everyone. *Remote Sens Environ*. 2017;202:18-27.

doi: 10.1016/j.rse.2017.06.031

48. TROPOMI. Aerosol Index (L2__AER_AI) and Carbon Monoxide (L2__CO____) product documentation. Data Products. Accessed September 21, 2026. https://www.tropomi.eu/data-products

49. ESA WorldCover Consortium. WorldCover 10 m 2020 v100. Earth Engine Data Catalog. Accessed September 21, 2026. https://developers.google.com/earth-engine/datasets/catalog/ESA_WorldCover_v100

50. WorldPop. WorldPop Global Project Population Data (R2025A): estimated residential population per 100x100m grid square. Earth Engine Data Catalog. Accessed September 21, 2026. https://developers.google.com/earth-engine/datasets/catalog/WorldPop_GP_100m_R2025A_v1_pop

51. US Geological Survey. Landsat 8 Level 2 Collection 2 Tier 1. Earth Engine Data Catalog. Accessed September 21, 2026. https://developers.google.com/earth-engine/datasets/catalog/LANDSAT_LC08_C02_T1_L2

52. Yamazaki D, Ikeshima D, Sosa J, Bates PD, Allen GH, Pavelsky TM. MERIT Hydro: a high-resolution global hydrography map based on latest topography dataset. *Water Resour Res*. 2019;55(6):5053-5073.

doi: 10.1029/2019WR024873

53. Ploton P, Mortier F, Réjou-Méchain M, *et al*. Spatial validation reveals poor predictive performance of large-scale ecological mapping models. *Nat Commun*. 2020;11(1):4540.

doi: 10.1038/s41467-020-18321-y

54. Wang M, Beelen R, Eeftens M, Meliefste K, Hoek G, Brunekreef B. Systematic evaluation of land use regression models for NO2. *Environ Sci Technol*. 2012;46(8):4481-4489.

doi: 10.1021/es204183v

55. Lundberg SM, Lee SI. A unified approach to interpreting model predictions. In: von Luxburg U, Guyon I, Bengio S, Wallach H, Fergus R, eds. *Proceedings of the 31st International Conference on Neural Information Processing Systems (NIPS'17)*; December 4-9, 2017; Long Beach, CA. Red Hook, NY: Curran Associates Inc.; 2017:4768-4777.

doi: 10.5555/3295222.3295230

56. Lundberg SM, Erion G, Chen H, *et al*. From local explanations to global understanding with explainable AI for trees. *Nat Mach Intell*. 2020;2(1):56-67.

doi: 10.1038/s42256-019-0138-9

57. Chawla NV, Bowyer KW, Hall LO, Kegelmeyer WP. SMOTE: synthetic minority over-sampling technique. *J Artif Intell Res*. 2002;16:321-357.

doi: 10.1613/jair.953

58. Kaufman S, Rosset S, Perlich C, Stitelman O. Leakage in data mining: formulation, detection, and avoidance. *ACM Trans Knowl Discov Data*. 2012;6(4):15:1-15:21.

doi: 10.1145/2382577.2382579

59. Kapoor S, Narayanan A. Leakage and the reproducibility crisis in machine-learning-based science. *Patterns*. 2023;4(9):100804.

doi: 10.1016/j.patter.2023.100804

60. Han H, Wang WY, Mao BH. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In: Huang DS, Zhang XP, Huang GB, eds. *Advances in Intelligent Computing (Lecture Notes in Computer Science)*. Vol 3644. Berlin, Germany: Springer; 2005:878-887.

doi: 10.1007/11538059_91

61. Ke G, Meng Q, Finley T, *et al*. LightGBM: a highly efficient gradient boosting decision tree. In: Guyon I, von Luxburg U, Bengio S, Wallach H, Fergus R, Vishwanathan S, Garnett R, eds. *Advances in Neural Information Processing Systems 30*. Red Hook, NY: Neural Information Processing Systems Foundation, Inc.; 2017:3146-3154.

62. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. In: *KDD '16: Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*; August 13-17, 2016; San Francisco, CA. New York, NY: Association for Computing Machinery; 2016:785-794.

doi: 10.1145/2939672.2939785

63. Breiman L. Random forests. *Mach Learn*. 2001;45(1):5-32.

doi: 10.1023/A:1010933404324

64. Arik SÖ, Pfister T. TabNet: attentive interpretable tabular learning. *Proc AAAI Conf Artif Intell*. 2021;35(8):6679-6687.

doi: 10.1609/aaai.v35i8.16826

65. Le Rest K, Pinaud D, Monestiez P, Chadoeuf J, Bretagnolle V. Spatial leave-one-out cross-validation for variable selection in the presence of spatial autocorrelation. *Glob Ecol Biogeogr*. 2014;23(7):811-820.

doi: 10.1111/geb.12161

66. Roberts DR, Bahn V, Ciuti S, *et al*. Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*. 2017;40(8):913-929.

doi: 10.1111/ecog.02881

67. Powers DMW. Evaluation: from precision, recall and F-measure to ROC, informedness, markedness & correlation. *J Mach Learn Technol*. 2011;2(1):37-63.

68. Brodersen KH, Ong CS, Stephan KE, Buhmann JM. The balanced accuracy and its posterior distribution. In: *Proceedings of the 20th International Conference on Pattern Recognition*; August 23-26, 2010; Istanbul, Turkey. Piscataway, NJ: IEEE Computer Society; 2010:3121-3124.

doi: 10.1109/ICPR.2010.764

69. Gorodkin J. Comparing two K-category assignments by a K-category correlation coefficient. *Comput Biol Chem*. 2004;28(5-6):367-374.

doi: 10.1016/j.compbiolchem.2004.09.006

70. Shannon CE. A mathematical theory of communication. *Bell Syst Tech J*. 1948;27(3):379-423.

doi: 10.1002/j.1538-7305.1948.tb01338.x

71. Tornetta GN. Entropy methods for the confidence assessment of probabilistic classification models. arXiv. Preprint posted online March 28, 2021. Accessed September 21, 2026.

doi: 10.48550/arXiv.2103.15157

72. Cedergreen N. Quantifying synergy: a systematic review of mixture toxicity studies within environmental toxicology. *PLoS One*. 2014;9(5):e96580.

doi: 10.1371/journal.pone.0096580

73. Gill JC, Malamud BD. Reviewing and visualizing the interactions of natural hazards. *Rev Geophys*. 2014;52(4):680-722.

doi: 10.1002/2013RG000445

74. Hollmann N, Müller S, Purucker L, *et al*. Accurate predictions on small data with a tabular foundation model. *Nature*. 2025;637(8045):319-326.

doi: 10.1038/s41586-024-08328-6

75. Si JYH, Cheng WY, Cooper M, Krishnan R. InterpreTabNet: distilling predictive signals from tabular data by salient feature interpretation. *Proc Mach Learn Res*. 2024;235:45353-45405. Accessed September 21, 2026. https://proceedings.mlr.press/v235/si24a.html

76. Grinsztajn L, Oyallon E, Varoquaux G. Why do tree-based models still outperform deep learning on tabular data? arXiv. Preprint posted online July 18, 2022. Accessed September 21, 2026.

doi: 10.48550/arXiv.2207.08815

77. Bennett MM, Smith LC. Advances in using multitemporal night-time lights satellite imagery to detect, estimate, and monitor socioeconomic dynamics. *Remote Sens Environ*. 2017;192:176-197.

doi: 10.1016/j.rse.2017.01.005

78. Bhattarai D, Lucieer A, Lovell H, Aryal J. Remote sensing of night-time lights and electricity consumption: a systematic literature review and meta-analysis. *Geogr Compass*. 2023;17(4):e12684.

doi: 10.1111/gec3.12684

79. Mellander C, Lobo J, Stolarick K, Matheson Z. Night-time light data: a good proxy measure for economic activity? *PLoS One*. 2015;10(10):e0139779.

doi: 10.1371/journal.pone.0139779

80. Zheng Q, Seto KC, Zhou Y, You S, Weng Q. Nighttime light remote sensing for urban applications: progress, challenges, and prospects. *ISPRS J Photogramm Remote Sens*. 2023;202:125-141.

doi: 10.1016/j.isprsjprs.2023.05.028

81. Gill JC, Malamud BD. Hazard interactions and interaction networks (cascades) within multi-hazard methodologies. *Earth Syst Dyn*. 2016;7(3):659-679.

doi: 10.5194/esd-7-659-2016

82. White CJ, Adnan MSG, Arosio M, *et al*. Review article: towards multi-hazard and multi-risk indicators—a review and recommendations for development and implementation. *Nat Hazards Earth Syst Sci*. 2025;25(11):4263-4281.

doi: 10.5194/nhess-25-4263-2025