https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

Geochemical predictability quantification and spatial transferability of soil organic carbon in Ganges Delta

Md Rakib Hasan[1,2], Fazla Zawadul Arabi[1], Mst Anika Khatun Rupa[1], Shoumik Zubyer[1], Mohammad Fazle Alam Rabbi[1], and Md Jashim Uddin[1]

[1]Department of Soil, Water and Environment, University of Dhaka, Dhaka-1000, Bangladesh
[2]Fermium Systems, Dhaka-1207, Bangladesh

Correspondence: Md Rakib Hasan (rakibhhridoy@fermium.systems) and Md Jashim Uddin (juswe@du.ac.bd)

Abstract. Deltas are among the least studied but most dynamically active reservoirs that store a significant amount of the world's soil organic carbon (SOC). This study evaluates the complex dynamics of deltas by combining soil profiles from four geomorphological zones in the Ganges Delta of Bangladesh with 6,652 soil profiles from fourteen major deltas worldwide. These combined profiles test the limits of regional SOC models, showing where local frameworks hold and where they break. Statistical analysis shows that salinity, cation exchange capacity, and moisture correlate most strongly with SOC ($r = 0.66$, 0.64, and 0.56, respectively), with most saline profiles storing approximately 6.5 times more carbon than freshwater sites. Salinity also exhibits a strict scale-dependent effect of strong correlations across geomorphological zones but effectively vanishing within them. This spatial paradox introduces major implications for digital soil mapping protocols in coastal deltas. Projections from a process model show that mangrove restoration efforts on half of the most saline tracts would increase SOC levels by 26.6 % in 2050, far exceeding baseline model of business-as-usual projections that show a 5.9–11.6 % decline across climate pathways. Model simulations indicate that local management has 4.6 to 7.2 times more influence on SOC storage than the gap between RCP 4.5 and RCP 8.5. To prevent data overlap from multi-depth sampling, the cross-validation had to hold back entire locations to avoid spatial leakage. As a result gradient boosting reached an R^2 of 0.64, against an inflated 0.91, and deeper neural networks provided no reliable improvement at this sample size. Analyzing every combination of the fifteen deltas exposed a sharp limit: models generalize within a single delta (median $R^2 = +0.25$) but fail between deltas, yielding positive scores in only 6 % of 210 pairs (median $R^2 = -0.55$). Transfer failed in both directions, even when fourteen deltas were used to predict the fifteenth delta. Unlike statistical weights, the structural framework of the process model transferred well, showing that the form of the model is broadly applicable when parameter tuning remains local site specific. Consequently, gathering data from foreign deltas cannot substitute for local profile measurements. These findings demonstrate that global datasets can guide structural frameworks but effective coastal carbon management ultimately depends on local site-specific observation.

1 Introduction

Deltas occupy less than one percent of Earth's land surface yet store carbon far out of proportion to their footprint. A combination of continuous sediment supply, waterlogging, brackish flooding, and reduced microbial decomposition builds and preserves

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

organic-rich soils in environments that would otherwise be mineral-dominated (Rogers et al., 2019; Lovelock et al., 2017; Alongi, 2012). Deltas are also among the most dynamic environments, increasingly threatened by embankments, aquaculture conversion, upstream damming, and altered freshwater delivery on timescales far shorter than pedogenesis. Furthermore, relative sea-level rise and saltwater intrusion are redrawing the geochemical envelope within which soils develop (Ericson et al., 2006; Allison and Meselhe, 2014; Murdiyarso et al., 2015; Sarker et al., 2011), making these rapidly changing SOC reservoirs among the least understood in terms of the physical and chemical mechanisms driving carbon storage.

The lack of understanding is not due to a scarcity of predictive models. As a common answer to spatial data scarcity, digital soil mapping increasingly relies on machine learning, displacing the traditional geostatistical baselines for SOC (Wadoux et al., 2021; Padarian et al., 2019; McBratney et al., 2003).

Although global models like SoilGrids provide routine continental-scale SOC mapping (Hengl et al., 2017), they largely derive their predictive controls, features, and validation frameworks from upland or temperate systems. Unlike upland or temperate domains, deltaic SOC dynamics are driven by unique physical and biogeochemical processes, including tidal inundation, salinity gradients, aquaculture-induced redox changes, and localized mineralogy. Whether relationships calibrated in other systems can be successfully transferred to deltas remains an open and consequential question.

Four fundamental questions guide the deltaic SOC literature. First, what regulates carbon accumulation and persistence? While clay mineralogy, redox state, mineral-organic binding, and salinity drive general SOC stabilization (Sollins et al., 1996; Schmidt et al., 2011; Six et al., 2002; Kleber et al., 2015; Lehmann and Kleber, 2015), observations across mangroves and tidal wetlands single out salinity-suppressed microbial activity as a unique deltaic mechanism (Setia et al., 2013; Wong et al., 2010; Rasmussen et al., 2018). Despite this, their relative importance across tropical deltas featuring steep salinity and texture gradients remains largely unquantified.

The second question explores the extent to which climate change against management actions control mid-century (2050) SOC trajectories. While warming drives faster decomposition through Q_{10} kinetics (Davidson and Janssens, 2006; Conant et al., 2011; Todd-Brown et al., 2013), restoration practices like mangrove planting can simultaneously lift carbon inputs and curb losses (Donato et al., 2011; Lovelock et al., 2017). Despite this dual action, the dominant lever governing deltaic carbon dynamics has not been rigorously tested.

The third is whether SOC can be predicted reliably from currently accessible covariates in the data-scarce areas common to coastal studies (Grinsztajn et al., 2022; Gorishniy et al., 2021; Shwartz-Ziv and Armon, 2022). The fourth, and arguably the most consequential for scaling, is whether models and process mechanisms developed for one delta transfer to others (Padarian et al., 2019; Arrouays et al., 2014).

Most studies examine these components separately, focusing on geochemical mapping, temporal process simulation, or machine-learning-driven prediction, while leaving cross-domain transferability unverified. In practice, these four distinct strands function as a tightly coupled system. Determining the primary environmental controls directly guides feature engineering for predictive models. Furthermore, process-based simulations rely on how those mechanisms react to external forcing, and the ultimate value of any model depends on its ability to generalize and transfer beyond its original training site.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

This study addresses these four strands within a single analytical pipeline (Fig. 1) focused on a major, climate-vulnerable delta of the Ganges. Twenty 1-meter soil pits across five depths in four geomorphological zones of the Ganges Delta drive both the geochemical evaluation and the calibration of a first-order decomposition model. Machine-learning models are trained on this feature matrix using depth-aware spatial cross-validation, while their transferability is tested across a fourteen-delta global panel from WoSIS 2023 (Batjes et al., 2024), with Rhine-Meuse data supplemented by the Dutch BIS Nederland dataset (Wageningen Environmental Research, 2020). The panel spans tropical to subarctic climates.

The core insight of this study is termed *layered transferability*. Predictive frameworks display dual-tier scaling in which fine-tuned weights and direct relationships stay local, but structural decisions of which predictors matter, how features are engineered from them, and the functional form of the process model travel successfully to new regions. Recognizing this hierarchy offers a more honest approach to scaling deltaic SOC science than the extremes of absolute non-transferability or universal global models.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

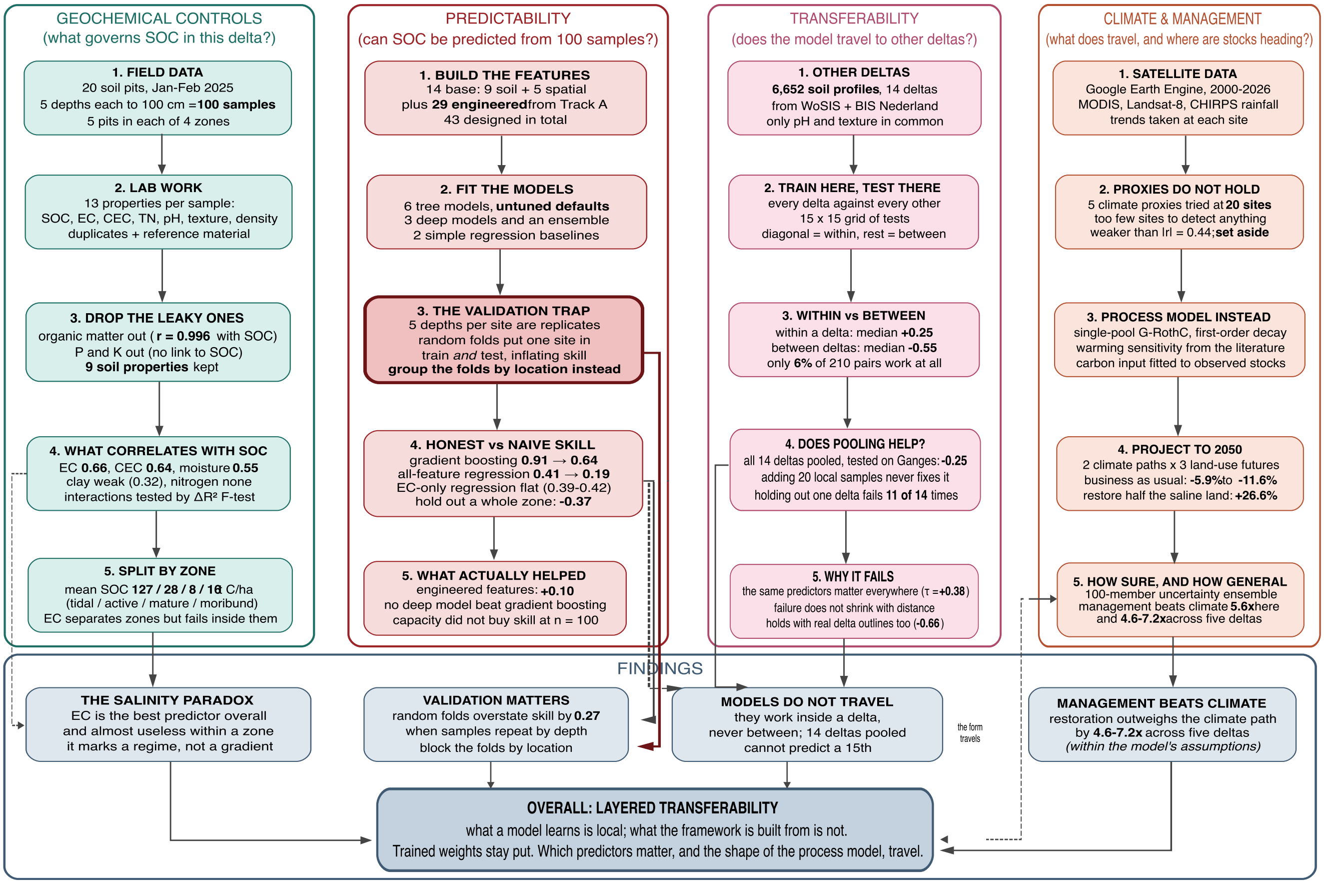

Figure 1. Analytical workflow. Four evidence tracks run from data to result and converge on the findings. Track A establishes which soil properties govern SOC; Track B tests how well SOC can be predicted from them at $n = 100$, and what validation design that requires; Track C asks whether the resulting model travels to other deltas; and Track D asks what does travel, projecting stocks to 2050 with a process model. The highlighted box in Track B marks the methodological trap that shapes every predictive result reported here: because five depths are sampled at each location, random cross-validation places the same location in both training and test folds. Values shown are those reported in the Results.

2 Study area and data

2.1 Ganges Delta setting

The Ganges–Brahmaputra–Meghna system, with a footprint of roughly 100,000 km^2, is the world's largest deltaic complex and supports more than 160 million people (Allison and Meselhe, 2014; Ericson et al., 2006). Sampling covered the western Ganges Delta in southwestern Bangladesh (21.5–24.5°N, 88.0–90.0°E), where mean annual temperatures average 26 °C and rainfall

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

1,600–2,000 mm yr^{-1} (Fig. 2). Land cover ranges from paddy rice and brackish aquaculture ponds to the Sundarbans mangrove forest, the world's largest contiguous mangrove ecosystem.

The study area was stratified into four geomorphological zones following Hossain et al. (2015), delineated by delta landform, sediment age and fluvial–tidal position rather than by salinity: *Tidal Active* (tidally influenced mangrove and estuarine areas), *Active Delta* (recently deposited alluvial sediments), *Mature Delta* (stable floodplain under established agriculture) and *Moribund Delta* (tectonically uplifted terrain with reduced fluvial input). Each sampling location was assigned to a zone by spatial overlay with the corresponding polygons. Electrical conductivity (EC), later shown to be a strong SOC correlate, was deliberately not used in the classification, so that the observed between-zone EC contrasts (mean EC 4.7, 0.39, 0.12 and 0.13 dS m^{-1} in Tidal Active, Active, Mature and Moribund respectively; Supplementary Table S2) are an emergent property of the zones rather than a defining criterion. This matters for interpretation: it means the salinity–SOC relationship examined below is not circular with respect to the zone assignment.

2.2 Soil sampling and analysis

During the dry season of January–February 2025, twenty $1 \times 1 \times 1$ m soil pits were excavated, five per zone, and each was sampled at five fixed depths (0–20, 20–40, 40–60, 60–80 and 80–100 cm) using cores and an Edelman auger, yielding 100 samples. Sampling to 100 cm followed IPCC Tier 1 conventions for national SOC inventories; pilot observations indicated that below this depth the soils are dominated by alluvial parent material with negligible biological activity. The 1 m^3 pit face permitted full horizon description and undisturbed core extraction at each depth.

Thirteen physicochemical properties were determined for each sample using standard procedures: SOC (Walkley–Black wet oxidation; Walkley and Black, 1934), organic matter (Walkley–Black and Van Bemmelen), EC (1:5 soil:water), cation exchange capacity (CEC; ammonium acetate at pH 7; Chapman, 1965), total nitrogen (TN; Kjeldahl; Bremner, 1996), pH (1:2.5 soil:water), bulk density (core method), gravimetric moisture, Olsen P, ammonium-acetate/AAS K, and sand, silt and clay fractions (hydrometer; Bouyoucos, 1962). Quality assurance included duplicate analysis on every tenth sample and certified reference material (ISE 921, Wageningen); analytical precision for SOC was ± 5 % CV. P and K showed no meaningful univariate association with SOC and were not retained, leaving eleven measured properties. After excluding OM, which correlates $r = 0.996$ with SOC through the Van Bemmelen factor and so constitutes data leakage, and SOC itself as the target, nine soil properties entered the modeling pipeline.

SOC stocks were calculated as $\mathrm{SOC_{stock}} = \mathrm{BD} \times \mathrm{OC} \times d \times 10^{-1}$ (t C ha^{-1}). Values ranged from 1.1 to 275.5 t C ha^{-1} (mean 45.1 ± 62.7), with a strongly right-skewed distribution (skewness 2.2; Shapiro–Wilk $p < 0.001$). All univariate correlations therefore report bootstrap 95 % confidence intervals from 2,000 resamples alongside Spearman rank correlations as a distributional robustness check. For temporal and predictive analyses, locations were also classified by salinity into Freshwater ($n = 10$ locations, EC < 0.5 dS m^{-1}), Brackish ($n = 6$) and Coastal ($n = 4$, EC > 1.5).

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

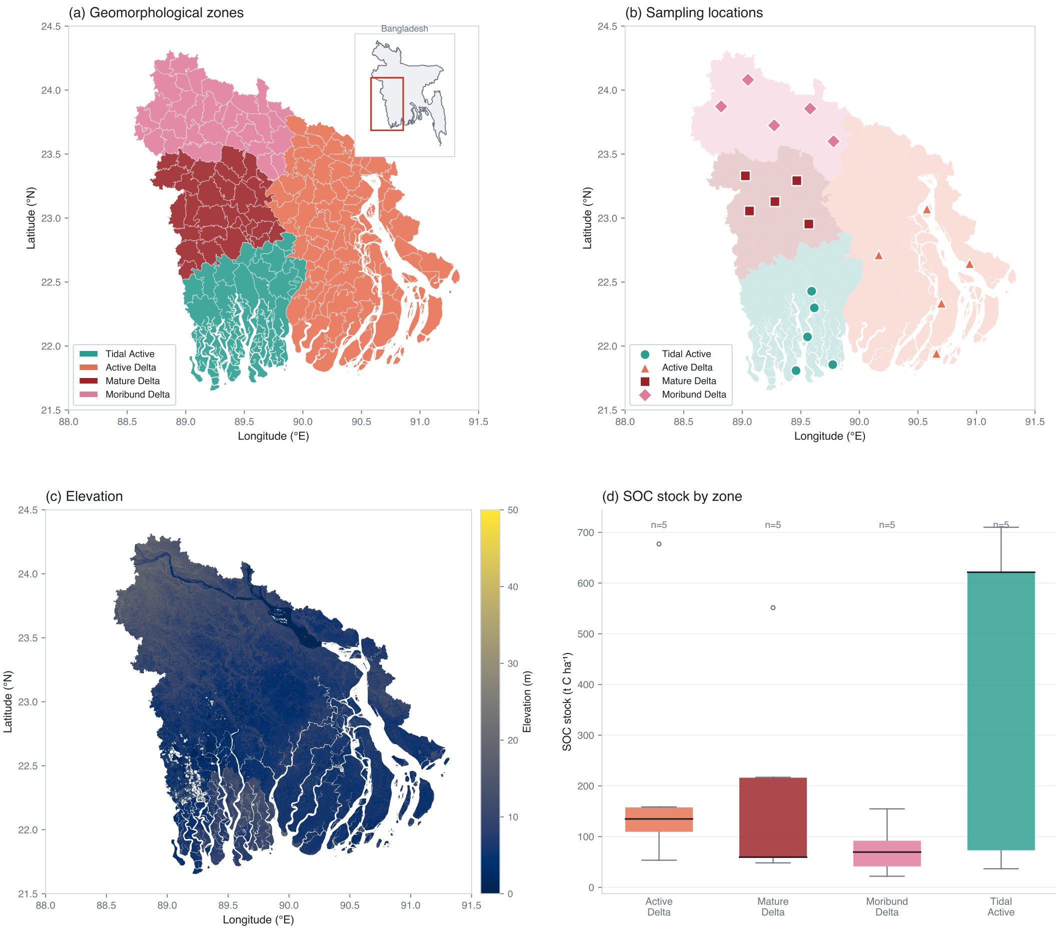

Figure 2. Sampling locations across the four geomorphological zones of the Ganges Delta. Each of the 20 locations was sampled at five depths (0–100 cm at 20 cm intervals), yielding 100 samples.

2.3 Remote sensing

Multi-temporal satellite imagery was acquired via Google Earth Engine covering 2000–2026: MODIS at 250 m for NDVI and land-surface temperature, Landsat 8 at 30 m for NDVI and NDWI, and CHIRPS for precipitation. Time series were extracted at each sampling location and annual trends computed by linear regression.

2.4 Cross-delta panel

The transferability panel drew primarily from the World Soil Information Service snapshot WoSIS 2023 December (Batjes et al., 2024), so that variable definitions, units and depth conventions were consistent across deltas without ad-hoc reconciliation. From WoSIS, all top 1 m (0–100 cm) layers falling within the bounding box of each major river delta were extracted, retaining

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

the predictors co-measured across deltas (pH and the sand, silt and clay fractions) with organic carbon (g kg^{-1}) as the target. The well-sampled Rhine-Meuse cell was additionally augmented with the BIS Nederland national soil survey (Wageningen Environmental Research, 2020) using standard conversions: soil organic matter percentages were converted to SOC via the Van Bemmelen factor (SOC $= 0.58 \times$ SOM), rescaled to g kg^{-1} to match WoSIS units, and Dutch RD-New coordinates (EPSG:28992) transformed to WGS84. BIS profiles within about 200 m of an existing WoSIS Rhine-Meuse profile were removed to avoid double-counting. The augmentation added 2,148 complete-case profiles to Rhine-Meuse, raising its sample size from 504 to 2,652, and provides an independent robustness check on this data-rich delta.

The full panel comprises 6,652 complete-case profiles across fourteen deltas: Mississippi ($n = 2{,}043$), Rhine-Meuse (2,652), Niger (327), Yellow River (315), Limpopo (268), Chao Phraya (204), Paraná–Plata (179), Sacramento (178), Indus (167), Mackenzie (77), Danube (68), Amazon (62), Senegal (57) and Pearl (55), spanning tropical, temperate, arid, Mediterranean and subarctic climates. Bulk density, CEC and EC were too sparsely co-measured across deltas to enter the common predictor set, restricting the feature-based transfer test to pH and texture. The tropical Asian monsoon deltas most physiographically analogous to the Ganges (Mekong, Irrawaddy, Godavari and Krishna) hold WoSIS organic-carbon values but lack co-measured texture or pH, and so could not enter the feature-based transfer test. That absence is itself a marker of the data scarcity motivating this analysis.

Bounding boxes are a coarse delineation, and they can admit non-deltaic upland soils. A second, geomorphologically stricter panel was therefore built from published active-delta polygons (Nienhuis et al., 2020) buffered by 0.5°, which retains 1,449 complete-case profiles across a more balanced set of deltas. This panel is utilized primarily for robustness validation rather than primary analysis because global soil archives are situated largely inland relative to active depositional plains. As a result, implementing strict spatial criteria enhances geomorphological exactness while substantially compromising statistical power through reduced sample sizes.

3 Methods

3.1 Geochemical analysis

Seven soil properties (clay, moisture, pH, TN, CEC, EC, bulk density) and three pairwise interactions (clay $\times$ moisture, CEC $\times$ pH, EC $\times$ moisture) were related to SOC using Pearson correlations, quadratic regression for pH, and hierarchical regression (ΔR^2, F-test) for the interactions (Sollins et al., 1996). These were treated as a single coordinated set of analyses, with family-wise error controlled jointly (Sect. 3.5). Since SOC stock is the product of concentration, bulk density and depth, the bulk density was correlated with SOC concentration to prevent algebraic dependency. Because the remaining predictors impact SOC stock strictly through concentration, replacing stock with concentration resulted in negligible shifts (≤ 0.005) across every Pearson coefficient. All correlations report Fisher z-transformed and bootstrap 95 % confidence intervals (2,000 resamples), variance inflation factors were computed to check multicollinearity, and zone-level analyses tested whether controls act uniformly across the delta.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

3.2 Process modeling

To investigate how soil organic carbon (SOC) responds to future climate and management scenarios over coming decades, a simplified single-pool decomposition model, designated G-RothC, was constructed. While drawing inspiration from the classical RothC framework (Coleman et al., 1997), G-RothC consolidates the original five distinct pools into a single, homogeneous soil organic carbon compartment:

$$\frac{d\mathrm{SOC}}{dt} = \underbrace{C_{\mathrm{in}}\, f_{\mathrm{moisture}}(\mathrm{NDWI})\, f_{\mathrm{EC}}(\mathrm{zone})\, f^{\mathrm{in}}_{\mathrm{mgmt}}}_{\text{carbon input}} - \underbrace{k_0\, f^{\mathrm{dec}}_{\mathrm{mgmt}}\, Q_{10}^{(T-T_{\mathrm{ref}})/10}\,\mathrm{SOC}}_{\text{first-order decomposition}} \tag{1}$$

Decomposition is first-order in the current stock, so losses scale with SOC. Here $k_0 = 0.05\,\mathrm{yr}^{-1}$ is the base decomposition rate, $Q_{10} = 2.4$ the temperature sensitivity (Lloyd and Taylor, 1994) and $T_{\mathrm{ref}} = 25\,°\mathrm{C}$. Land management acts on both terms: $f^{\mathrm{in}}_{\mathrm{mgmt}}$ modifies inputs and $f^{\mathrm{dec}}_{\mathrm{mgmt}}$ modifies decomposition (mangrove: 1.15 and 0.80 respectively, representing enhanced litter inputs and anoxic suppression of decomposition; aquaculture: 1.15 and 0.85 inverted; both bounded within a 0.7–1.2 envelope). Rather than prescribing C_{in} from the dimensionless NDVI directly, the input flux ($\mathrm{t\,C\,ha^{-1}\,yr^{-1}}$) was calibrated per location from the steady-state condition $C_{\mathrm{in}} f_{\mathrm{moisture}} f_{\mathrm{EC}} = k_0 Q_{10}^{(T_{\mathrm{base}}-T_{\mathrm{ref}})/10}\mathrm{SOC}_{\mathrm{obs}}$: that is, the input required to sustain the observed stock under baseline climate ($T_{\mathrm{base}} = 26\,°\mathrm{C}$). This yielded a mean input of $1.9\,\mathrm{t\,C\,ha^{-1}\,yr^{-1}}$ and a mean turnover time of about 18 yr.

To execute projections from 2000 to 2050, each location was initialized at its observed stock under baseline steady-state conditions and advanced using annual Euler steps. The experimental design evaluated six scenarios created by crossing two climate pathways (RCP 4.5 with $+1.0\,°\mathrm{C}$ warming and RCP 8.5 with $+2.0\,°\mathrm{C}$ warming, aligned with IPCC AR6 South Asia projections) with three management pathways: business-as-usual and 10 % or 50 % mangrove restoration of the most saline land. Because the available SOC data come from a single survey, the model is not validated against an observed time series; projected scenario differences and parameter sensitivity are therefore reported rather than a calibrated hindcast skill score. Sensitivity was assessed by varying Q_{10} (1.5–3.5) and k_0 (0.01–0.05), and a 100-member parameter ensemble ($k_0 \sim U(0.01, 0.05)$, $Q_{10} \sim U(1.5, 3.5)$) propagated parameter uncertainty into 90 % prediction intervals.

3.3 Predictive modeling

The full designed feature space comprised 43 predictors: 9 soil properties, 5 spatial and topographic features (latitude, longitude, elevation, slope, aspect), and 29 domain-informed engineered features. The 14 soil and spatial features constitute the base matrix, and the 29 engineered features were evaluated incrementally in ablation. The engineered features comprised interaction terms drawn from the geochemical analysis (EC × CEC, clay × moisture, pH × EC), ratio features (TN/EC, EC/CEC, sand/clay) and polynomial expansions (EC^2, EC^3, TN^2). All features were standardized to zero mean and unit variance prior to model training, with the standardizer fitted only on training data within each fold to prevent information leakage.

Six tree-based models were evaluated: gradient boosting regressor (GBR), Random Forest, XGBoost, LightGBM, CatBoost and a stacking ensemble (Pedregosa et al., 2011), against two ordinary least-squares baselines (all features; EC only). To test

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

whether additional model capacity improves prediction at this sample size, three deep-learning architectures were also explored: a multi-relational graph attention network with three edge types (spatial proximity ≤ 50 km, same geomorphological zone, same location at adjacent depths) (Veličković et al., 2018), a mixture-of-experts with top-2 gating (Shazeer et al., 2017), and a tabular vision transformer with soil and spatial modality tokens (Gorishniy et al., 2021), together with a weight-constrained ensemble combining these with gradient boosting. As reported below, none surpassed gradient boosting under leak-free cross-validation, so gradient boosting is retained as the predictive model. All experiments used a fixed random seed (random_state = 42; NumPy 1.26.4, PyTorch 2.x). Gradient boosting was run at library defaults, without hyperparameter search; all settings are given in Supplementary Table S3.

Because the 100 samples are depth replicates nested within 20 field locations, random k-fold cross-validation places samples from the same location in both training and test folds, leaking spatial information and inflating apparent skill. Location-grouped cross-validation is therefore treated as the authoritative metric. Four schemes are reported: random 5-fold (for comparability with the wider literature), location-grouped 5-fold (GroupKFold, holding out whole locations), Leave-Location-Out (LOLO, 20-fold) and Leave-Zone-Out (LZO, 4-fold). Ablation isolated the contributions of the geochemistry-informed engineered features and of soil-only inputs under location-grouped cross-validation.

3.4 Cross-delta transferability

Transferability was quantified with a full train-on-delta-i, test-on-delta-j matrix across fifteen deltas (fourteen external plus the Ganges). The diagonal is within-delta predictive skill (5-fold cross-validation, asking whether pH and texture recover SOC within a delta at all), and off-diagonal cells are zero-shot transfer (the R^2 of a model trained on one delta and applied without retraining to another). Also reported are: forward transfer, in which the Ganges-trained model is applied to each external delta; reverse transfer, in which a model pooled over all fourteen external deltas (6,652 samples) is applied to the Ganges, with few-shot calibration adding $n = 0, 5, 10, 15, 20$ local Ganges samples (20 repeats); Leave-One-Delta-Out cross-validation, which trains on all deltas but one and predicts the held-out delta; a generalized few-shot recovery analysis applying the reverse-transfer design to every delta; per-delta feature-importance rankings (Kendall τ against the Ganges); and distance-decay tests correlating transfer R^2 with geographic and feature-space distance between delta pairs. All models were gradient-boosting regressors with in-fold standardization. The full analysis was then repeated on the polygon-delineated panel of Sect. 2.4 as a robustness check.

The G-RothC form was then applied to five representative deltas (Ganges, Mississippi, Mekong, Rhine-Meuse and Yellow River) with region-specific baseline temperatures and RCP 4.5 / RCP 8.5 mid-century warming (IPCC AR6 regional projections), under business-as-usual and mangrove-restoration scenarios. Because the ratio of the management to the climate signal is independent of the absolute initial stock (both effects scale with the stock at risk), this comparison is robust to each delta's differing SOC baseline. Per-delta management/climate ratios are therefore reported rather than absolute projected stocks.

3.5 Statistical framework

All hypothesis tests used $\alpha = 0.05$. Bootstrap confidence intervals from 2,000 resamples were computed for correlations and model performance. Family-wise error across the 25 sub-analyses was controlled with Bonferroni-adjusted significance

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

thresholds ($\alpha_{\mathrm{adj}} = 0.05/25 = 0.002$), and Benjamini–Hochberg FDR-corrected q-values are reported alongside raw p-values (Supplementary Table S7). Simple baselines (OLS with all features, and EC-only OLS) quantified the marginal gain of the machine-learning approaches. Post-hoc power analysis indicated that at $n = 20$ locations the minimum detectable correlation at $\alpha = 0.05$ with 80 % power was $|r| \geq 0.44$; effects below this threshold are flagged as underpowered. Residual diagnostics included Shapiro–Wilk (normality) and Breusch–Pagan (heteroscedasticity) tests. Effect sizes are interpreted using Cohen's d (small ≥ 0.2, medium ≥ 0.5, large ≥ 0.8) and ΔR^2 (small ≥ 0.02, medium ≥ 0.13, large ≥ 0.26).

4 Results

4.1 Geochemical controls on SOC

Three critical findings emerge from the geochemical data. A handful of salinity- and cation-exchange-related properties control SOC variation, but their influence shifts with scale, dominating broad inter-zone macro comparisons while disappearing at the local micro level. At micro scales, accurately capturing within-zone dynamics requires modeling pairwise interactions instead of relying on single-variable controls.

Three properties correlated strongly with SOC: EC ($r = 0.664$, $p < 0.0001$), CEC ($r = 0.636$, $p < 0.0001$) and moisture ($r = 0.552$, $p < 0.0001$). Clay showed a weak but significant association ($r = 0.323$, $p = 0.001$), TN none at the univariate level ($r = 0.045$, $p = 0.657$), and bulk density a moderate negative correlation with SOC concentration ($r = -0.511$, $p < 0.0001$). Seven of the ten sub-analyses survived Bonferroni correction ($\alpha_{\mathrm{adj}} = 0.002$; Table 1). Spearman rank correlations, insensitive to the right skew, were consistent with or stronger than the Pearson estimates (EC $\rho = 0.78$, CEC $\rho = 0.78$, moisture $\rho = 0.53$, clay $\rho = 0.47$), verifying that the observed patterns are robust against data skewness.

Interactions carried much of the finer structure. The CEC × pH term was the strongest ($\Delta R^2 = 0.095$, $F = 20.75$, $p < 0.0001$), consistent with pH-dependent, exchange-mediated stabilization, and the clay × moisture term ($\Delta R^2 = 0.068$) confirmed that clay protects SOC more effectively under waterlogged conditions (Baldock and Skjemstad, 2000). When samples were regrouped by a simple salinity threshold, brackish samples (EC $> 0.5\,\mathrm{dS\,m^{-1}}$) stored 6.5 times more SOC than freshwater samples (means 104.1 versus $16.0\,\mathrm{t\,C\,ha^{-1}}$; Cohen's $d = 1.87$; Kruskal–Wallis $H = 50.2$, $p < 0.001$). Because that grouping is defined on EC itself it merely illustrates the magnitude of the salinity–SOC association; the independent evidence for that association comes from the geomorphologically defined zones.

The scale-dependence is the more striking observation. Within each of the four zones, the strongest pooled predictor, EC, largely disappeared (Fig. 3; Supplementary Table S10). It correlated weakly and non-significantly with SOC within the two saline zones (Tidal Active $r = -0.20$, Active Delta $r = +0.34$) and persisted only in the freshwater ones (Mature $r = +0.65$, Moribund $r = +0.55$). Salinity therefore acts primarily as a between-zone discriminator, the switch between the salinity-suppressed decomposition regime described for other coastal systems (Setia et al., 2013; Wong et al., 2010) and the freshwater regime, rather than as a within-zone gradient. This behavior is termed the EC paradox here: EC is the single most informative predictor when the delta is treated as one system, and among the least informative predictors within the settings where it is highest. Digital soil mapping in coastal contexts that relies on salinity as a covariate should expect exactly this behavior.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

Table 1. Geochemical controls on SOC: individual properties and theory-motivated pairwise interactions.

ID	Mechanism	r / R^2 / ΔR^2	p	Status
H1.1	Clay mineralogy	$r = 0.323$	0.001	Significant (weak effect)
H1.2	Anoxia/Moisture	$r = 0.552$	<0.0001	Strong
H1.3	pH nonlinearity	$R^2 = 0.108$	0.031	Weak
H1.4	Nutrient (TN)	$r = 0.045$	0.657	Not confirmed
H1.5	CEC aggregation	$r = 0.636$	<0.0001	Strong
H1.6	Salinity (EC)	$r = 0.664$	<0.0001	Strong
H1.7	Bulk density†	$r = -0.511$	<0.0001	Moderate
H1.8	Clay×Moisture	$\Delta R^2 = 0.068$	<0.05	Confirmed
H1.9	CEC×pH	$\Delta R^2 = 0.095$	<0.0001	Confirmed
H1.10	EC×Moisture	$\Delta R^2 = 0.028$	0.021	Confirmed

Note: Most rows report the Pearson correlation coefficient r; H1.3 (pH) reports the coefficient of determination R^2 from a quadratic regression, reflecting the hypothesized non-linear pH–SOC relationship; H1.8–H1.10 report the incremental ΔR^2 contributed by the interaction term in hierarchical regression. For H1.3, the incremental pH^2 term is significant at $\alpha = 0.05$ ($p = 0.031$) but does not survive Bonferroni correction ($\alpha_{\text{adj}} = 0.002$), hence its "Weak" status. †H1.7 (bulk density) is correlated against SOC *concentration* rather than stock, to avoid the algebraic dependency of stock on bulk density.

Zone-level analysis also exposed sign reversals hidden by pooling. TN correlated strongly but oppositely in Tidal Active ($r = -0.79$) versus Active Delta ($r = +0.73$), and pH switched from $+0.61$ (Tidal Active) to -0.67 (Mature). Averaged across zones, both variables looked weak; disaggregated, they are neither weak nor consistent. TN also ranked second in the tree-based model's feature importance (30.9 %) despite its null univariate correlation, largely through the TN/EC ratio: the sort of interaction gradient boosting can exploit through sequential splits but that a linear analysis conceals (Kleber et al., 2015; Wagai and Mayer, 2009).

4.2 Predictive skill and the cost of validation leakage

Preliminary baselines set the floor: a straightforward EC-based OLS produced steady performance ($R^2 \approx 0.39$–0.42) regardless of validation strategy. A full-feature OLS initially looked promising under random cross-validation ($R^2 = 0.41$), but dropped sharply to $R^2 = 0.19$ when whole locations were withheld, demonstrating that depth-replicated designs foster data leakage and that complex models suffer the steepest penalties. Gradient boosting was the strongest predictor, achieving a location-grouped R^2 of 0.64 (Table 2). Three deep learning architectures (multi-relational GNN, mixture-of-experts, tabular ViT) and their ensemble were also tested. Under leak-free validation, none outperformed gradient boosting at this sample size, supporting its selection as the operational model. The greater capacity of these three architectures drove overfitting to local fold quirks, a known issue in small-sample tabular tasks (Grinsztajn et al., 2022; Gorishniy et al., 2021; Chen et al., 2024). At this scale, pushing model capacity beyond gradient boosting did not improve spatial prediction.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

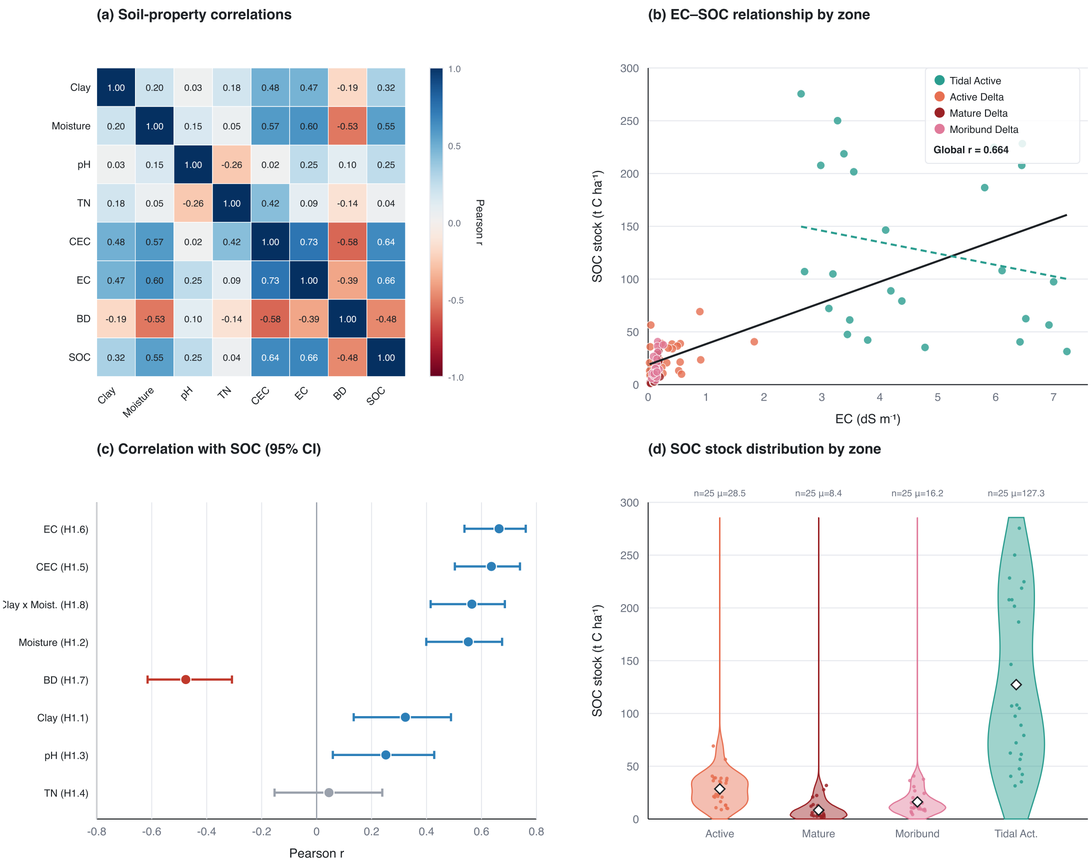

Figure 3. Geochemical controls on SOC: (a) correlation matrix; (b) EC against SOC by zone, illustrating the EC paradox; (c) correlation forest plot with 95 % confidence intervals; (d) SOC distribution by zone. Abbreviations: SOC, soil organic carbon; EC, electrical conductivity; CEC, cation exchange capacity; TN, total nitrogen; BD, bulk density.

Table 2. Predictive skill (R^2) under progressively more conservative cross-validation. Because the 100 samples are depth replicates within 20 field locations, random k-fold leaks within-location structure and is optimistic; location-grouped cross-validation is the authoritative metric for this design, and leave-zone-out (LZO) is the most demanding. The low-capacity EC-only baseline is essentially flat across schemes, whereas gradient boosting's apparent skill declines once whole locations, and especially whole zones, are withheld. All models exclude OM.

Model	Random 5-fold	Location-grouped	LOLO	LZO
EC-only OLS	*0.39*	*0.42*	*0.40*	*0.38*
GBR	0.91	0.64	0.87	−0.37

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

The gap between naive and leak-free validation is the operationally important number. Random 5-fold cross-validation gave $R^2 = 0.91$; location-grouped 5-fold gave $R^2 = 0.64$, a drop of 0.27 (Fig. 4a). Leave-Location-Out cross-validation, which holds out a single location and is therefore less demanding, gave $R^2 = 0.87$. Leave-Zone-Out cross-validation failed outright ($R^2 = -0.37$): a model trained on three of the four geomorphological zones could not predict the fourth (Supplementary Fig. S1). EC-only OLS, by contrast, was essentially flat across schemes because it had no within-location structure to exploit. The location-grouped values are therefore treated as authoritative, and the random k-fold numbers as optimistic upper bounds retained for cross-study comparability.

Reconstructed geochemistry-informed features, tested under location-grouped validation, added steady gains to gradient boosting ($\Delta R^2 = +0.11$ over soil-only models and $+0.10$ over base features; Fig. 4d). The top-performing interaction terms, TN/EC and EC $\times$ CEC, capture SOC variance that standalone raw covariates miss.

Residual diagnostics on the location-grouped out-of-fold residuals showed non-normality (Shapiro–Wilk $W = 0.79$, $p < 0.001$) and heteroscedasticity (Breusch–Pagan $p < 0.001$), with systematic under-prediction of the highest-SOC samples. Hold-out R^2 values are inflated by the small test set ($n = 20$); the leak-free cross-validation estimates should be treated as the authoritative performance metric.

4.3 Layered transferability across fifteen deltas

The transfer matrix (Fig. 5) shows a sharp separation between within-delta and cross-delta skill. Within a delta, the pH and texture predictors carried genuine SOC signal: diagonal 5-fold cross-validation R^2 was positive in most deltas with adequate data, with a median of $+0.25$ and values reaching $+0.99$ for Senegal, $+0.84$ for Sacramento, $+0.79$ for Danube and $+0.73$ for Indus (Supplementary Table S6). Augmenting Rhine-Meuse with the BIS Nederland national survey raised its within-delta R^2 from $+0.15$ (504 WoSIS profiles) to $+0.19$ (2,652 combined profiles), confirming that the within-delta signal is real and strengthens with denser, higher-quality data.

Across deltas, that predictive power disappeared. Out of 210 train-test pairs, only 6.2 % scored an $R^2 > 0$, and the median fell to -0.55, with numerous strongly negative results. Transfer failed symmetrically. Forward transfer from the Ganges-trained model was negative for all fourteen external deltas. Reverse transfer from a model pooled over all fourteen external deltas (6,652 samples) applied to the Ganges reached $R^2 = -0.25$, and adding up to 20 local Ganges samples did not lift it above zero. Applying the few-shot test across every delta showed that data pooling offers no universal remedy: combining external data with local samples failed to achieve 50 % of the intra-delta cross-validation ceiling anywhere, and scored positive held-out R^2 values in just two deltas. In high-performing local systems like Senegal ($+0.99$), Sacramento ($+0.84$), and Danube ($+0.79$), pooling data was actively harmful, performing orders of magnitude worse than local-only models.

Three converging diagnostics explain why. First, feature-importance rankings were broadly concordant across deltas (mean Kendall τ against the Ganges $= +0.38$; pH was the top predictor in six of the fourteen external deltas). The same handful of covariates matter everywhere, so it is not that deltas are governed by entirely different soil-forming controls. Second, cross-delta transfer R^2 was uncorrelated with geographic, feature-space and latitudinal distance between deltas (all Spearman $|\rho| \leq 0.11$, $p \geq 0.10$). Nearby and environmentally similar deltas transferred no better than distant ones, ruling out a simple

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

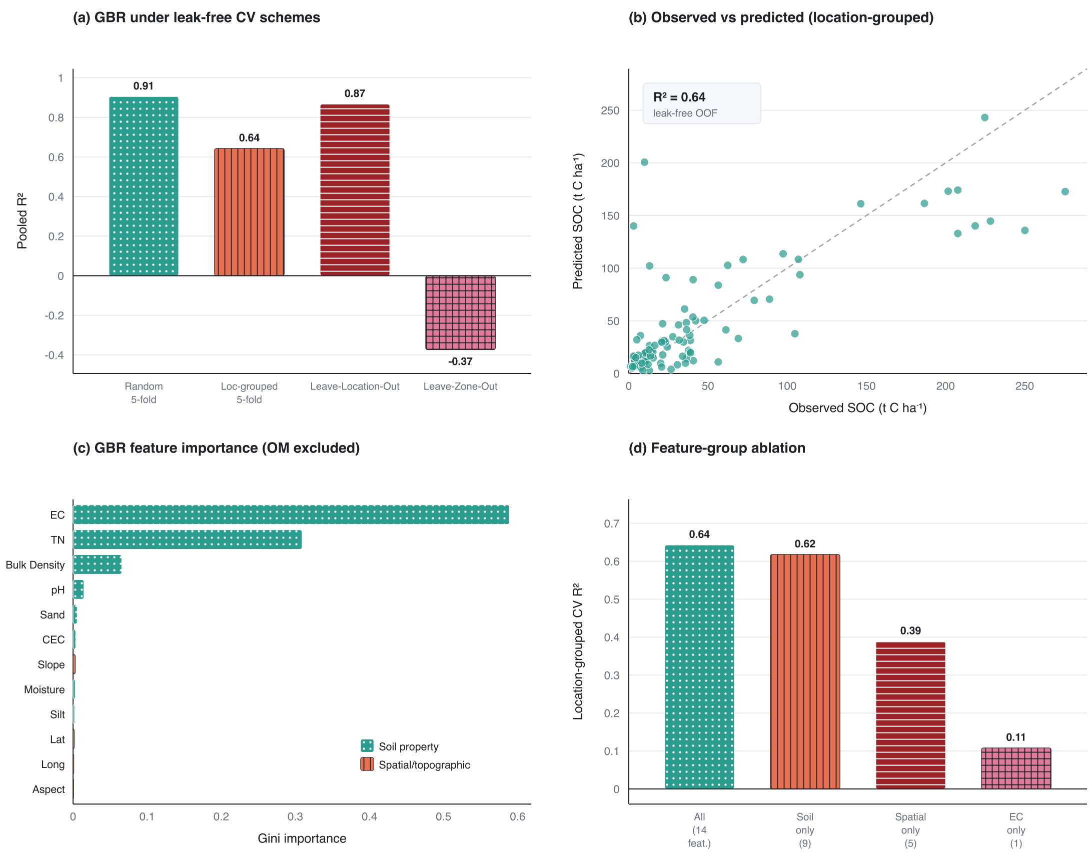

Figure 4. Machine-learning prediction under leak-free validation: (a) gradient-boosting R^2 under random versus location-grouped cross-validation, illustrating the optimism of random k-fold in this depth-replicated design; (b) observed versus predicted SOC from location-grouped out-of-fold predictions; (c) gradient-boosting feature importance (OM excluded), top 12 predictors colored by feature group; (d) feature-group ablation under location-grouped cross-validation. Abbreviations: CV, cross-validation; GBR, gradient boosting regressor; SOC, soil organic carbon; OM, organic matter.

ecological-distance interpretation. Third, Leave-One-Delta-Out cross-validation confirmed this at the pooled level: holding out an entire delta yielded negative R^2 for eleven of fourteen deltas, so a fourteen-delta, 6,652-sample pooled model still could not predict a delta absent from training.

To address whether transfer failures are an artifact of non-deltaic upland soils captured within rectangular bounding boxes, the entire analysis was replicated using strict polygon delineations (Sect. 2.4). Among 1,449 profiles falling inside active-delta polygons, the cross-delta median R^2 was -0.66, only 5.2 % of 210 pairs exceeded zero, reverse pooled-to-Ganges transfer scored $R^2 = -0.74$, and leave-one-delta-out cross-validation remained negative for 13 of 14 deltas (Supplementary Table S11). Applying a stricter delineation process ultimately solidifies the core conclusions. Although within-delta median skill drops from

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

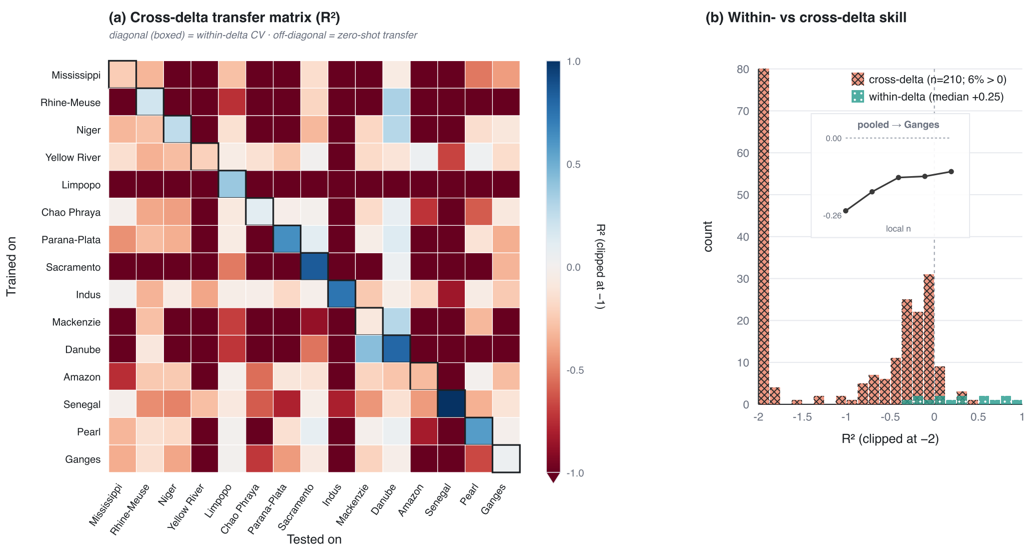

Figure 5. Cross-delta transferability of SOC prediction across 15 deltas (14 from harmonized WoSIS profiles, Batjes et al., 2024, with the BIS Nederland national survey augmenting Rhine-Meuse, Wageningen Environmental Research, 2020; plus the Ganges; predictors pH and sand/silt/clay). (a) Transfer matrix: rows are the training delta, columns the test delta; the boxed diagonal is within-delta 5-fold cross-validation and off-diagonal cells are zero-shot transfer R^2 (color clipped at -1; deep red indicates no transferable skill). (b) Distribution of within-delta versus cross-delta R^2 (clipped at -2): within-delta skill is positive (median $+0.25$) whereas only 6 % of 210 cross-delta pairs exceed zero (median -0.55); the inset shows the reverse-transfer few-shot curve (pooled 14-delta model applied to the Ganges), which does not become positive even with 20 local Ganges samples added. Abbreviations: SOC, soil organic carbon.

$+0.25$ to $+0.05$, this attenuation is expected given a fourfold reduction in sample sizes and the greater internal uniformity of the active delta plain relative to its surroundings.

The picture is therefore clear: predictor-to-SOC relationships are delta-specific in their quantitative form, even though the predictors themselves are universally relevant. Models generalize within a delta at the micro scale but not between deltas at the macro scale, in either direction. Also pooling many deltas does not substitute for local sampling in a new delta, and failure is not a smooth function of environmental separation between deltas.

4.4 Climate, management, and what the process model carries

Limited by a small location-level sample size ($n = 20$), the remote-sensing proxy screening can detect only correlations exceeding $|r| \geq 0.44$ at 80 % power. Five proxies were examined, namely NDWI (flooding), EC (salinity), precipitation CV (monsoon intensity), NDVI trend (vegetation change) and $T_{\max}$ trend (temperature). At this scale and sample size, no association was robust enough to carry a claim. Results shifted depending on the specific product and trend window selected, and warming trends across twenty adjacent regional sites proved too spatially uniform to resolve an independent SOC gradient. The climate

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

assessment therefore relies on the process model, with the temperature sensitivity of decomposition ($Q_{10} = 2.4$) taken from the literature (Lloyd and Taylor, 1994; Bond-Lamberty and Thomson, 2010) rather than estimated from the proxy data.

From a baseline mean of 45.1 t C ha^{-1}, G-RothC projected a modest decline under business-as-usual warming and net gains under restoration (Fig. 6). Under RCP 4.5 business-as-usual, mean SOC fell to 42.4 t C ha^{-1} by 2050 (-5.9 %), with higher-stock brackish samples losing carbon faster (-0.12 t C ha^{-1} yr^{-1}) than freshwater samples (-0.02), consistent with first-order kinetics in which absolute losses scale with the stock at risk. Restoring 50 % of the most saline land to mangrove raised projected 2050 SOC to 57.1 t C ha^{-1} ($+26.6$ % relative to baseline; $+14.6$ t C ha^{-1} relative to business-as-usual), whereas the gap between climate pathways was smaller (-11.6 % under RCP 8.5 versus -5.9 % under RCP 4.5, a 2.6 t C ha^{-1} difference). Within the model, the management intervention thus outweighed the climate-pathway effect by a factor of 5.6.

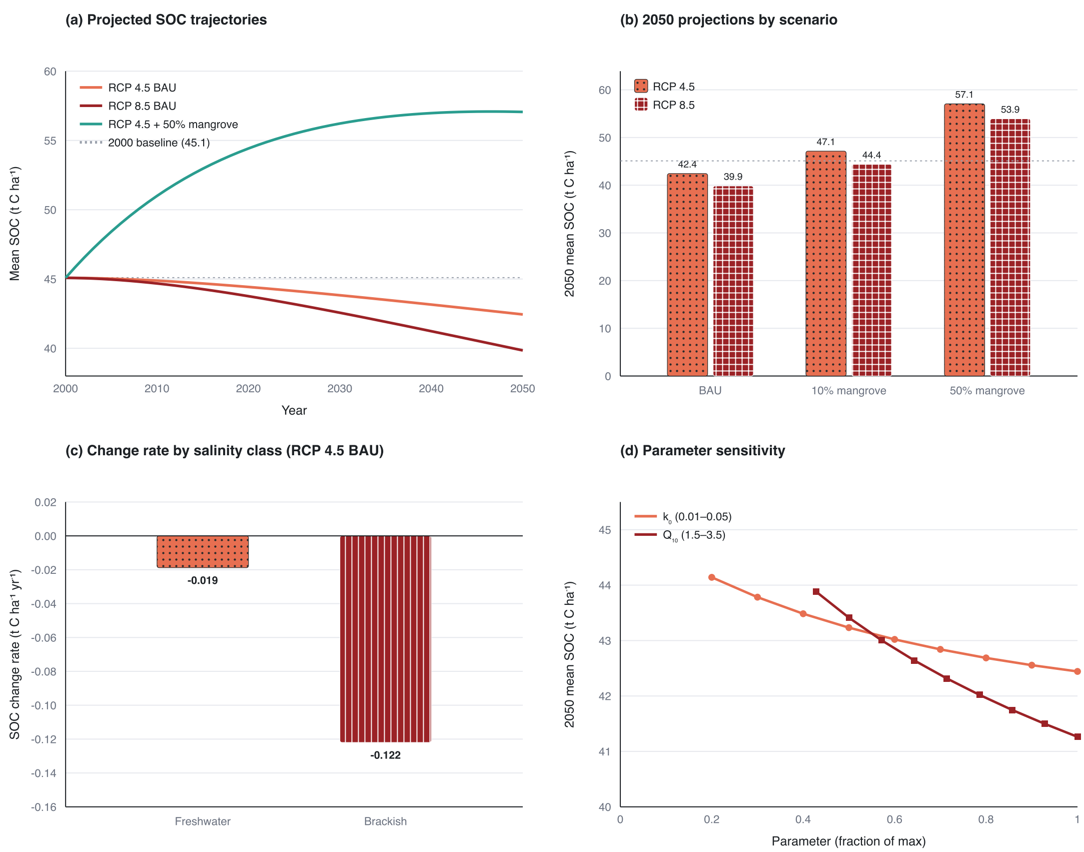

Figure 6. Climate dynamics from the first-order G-RothC: (a) projected mean SOC trajectories (2000–2050) under business-as-usual and mangrove-restoration scenarios; (b) projected 2050 mean SOC under the six climate × land-use scenarios; (c) SOC change rate by salinity class under RCP 4.5 business-as-usual; (d) sensitivity of 2050 mean SOC to k_0 and Q_{10}. Abbreviations: SOC, soil organic carbon; RCP, Representative Concentration Pathway; Q_{10}, temperature sensitivity of decomposition; k_0, base decomposition rate.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

A 100-member parameter ensemble (k_0 and Q_{10} sampled over their plausible ranges) placed the 90 % prediction interval for 2050 RCP 4.5 business-as-usual SOC at 41.8–44.2 t C ha^{-1}, and the restoration signal remained positive in 100 % of ensemble members. Because the input flux was calibrated to the observed steady state, projection uncertainty was governed more by Q_{10} (6.2 % of the 2050 mean) than by k_0 (3.9 %) (Sierra et al., 2015): the calibrated input co-varies with k_0, so the net response to forcing is dominated by the temperature sensitivity of decomposition rather than by the baseline rate itself.

The climate-scenario portability of the process model behaved quite differently from the learned weights discussed previously, as it relies on a mechanistic ratio rather than a fitted statistical mapping. Applied to the five representative deltas with region-specific climate forcing, the per-unit mangrove-restoration contrast outweighed the RCP 4.5 to RCP 8.5 climate-pathway effect in every case, by a factor of 4.6 to 7.2. Because this ratio is independent of the absolute initial stock, it is robust to the differing SOC baselines across deltas, and it extends the Ganges finding that land management is the dominant near-term SOC lever to a multi-delta context spanning tropical, temperate, arid, Mediterranean and subarctic regimes (Lovelock et al., 2017; Rogers et al., 2019). The underlying structure of the model successfully transfers to new regions even where its previously calibrated weights failed to do so.

5 Discussion

5.1 Layered transferability

The four analytical processes unite into a single workflow in which geochemistry identifies the key properties, process models predict how these properties respond to climate and management, then machine learning operationalizes them with custom features, and finally cross-delta tests map their transferability. The result is clear: the framework's components scale in distinct layers.

Calibrated model parameters and quantitative covariate-to-SOC functions are fundamentally restricted to local deltaic systems. Combining data from multiple deltas, even with local adjustments, cannot replace the need for local field sampling. However, the higher-level structures of the framework are considerably more portable to other regions.

Cross-delta comparisons reveal strong similarity on key feature importance, yielding a mean Kendall $\tau = +0.38$ against the Ganges, and identifying pH as the top predictor in six of fourteen external sites. The process model's mechanics translate equally well, with the management-to-climate signal ratio varying by less than a factor of 1.6 across five unique climate regimes. The only component that fails to travel is the numerical mapping of predictors into SOC estimates.

The concordance of feature rankings alongside the failure of learned mappings is a critical insight. It refutes the intuitive assumption that cross-delta transfer fails because different regions are governed by entirely distinct soil-forming controls. They are not; the same concise set of predictors matters everywhere. The breakdown occurs instead in the specific numerical coupling between those predictors and SOC: in how a given increase in pH or clay translates into a corresponding change in carbon concentration. This coupling depends on delta-specific factors (mineralogy, hydroperiod, litter chemistry, historical land use) that texture and pH alone fail to capture and that pooled global data cannot average out. The practical implication is direct: large-scale

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

SOC mapping efforts that extrapolate models from data-rich temperate deltas to data-poor tropical ones risk systematic bias, and local calibration data remain indispensable no matter how much global data are available.

The distance-decay diagnostic sharpens the point. If differences were driven by smooth ecological gradients, transfer success would drop steadily with distance. It does not. Nearby deltas transfer just as poorly as distant ones. In short, each delta is a unique biogeochemical system with its own quantitative signature, not just a point further away from a reference state.

5.2 The EC paradox and its implications for digital soil mapping

In the Ganges, the salinity paradox goes beyond mere statistical anomaly: while EC is the top pooled predictor of SOC overall, it provides no predictive power within the saline zones where SOC concentrations peak. This dynamic highlights a critical ecological mechanism. Salinity does not act as a continuous and graduated control; instead, it operates as a threshold switch separating the suppressed decomposition typical of tidal-brackish environments from normal freshwater conditions. Treating electrical conductivity as a simple linear covariate in digital soil mapping exposes models to dual errors: inflating local predictive accuracy in freshwater regions, and masking salinity's true role as a categorical regime shift when it is binned or left out. To overcome this in coastal environments, models should instead employ interaction structures where EC modulates other geochemical relationships, rather than acting as an independent predictor.

This finding is further confirmed by the total failure of Leave-Zone-Out cross-validation, where models trained on three zones proved incapable of predicting the fourth. These zones operate as independent biogeochemical regimes, demonstrating that accurate prediction requires individual zone representative data. This mirrors the cross-delta result: what looks like a single global relationship at a macro scale breaks down into separate regime-specific dynamics at a micro scale.

5.3 Management outweighs the climate pathway

Process modeling across five deltas reveals that mangrove restoration yields a per-unit-area impact 4.6 to 7.2 times greater than the gap between RCP 4.5 and RCP 8.5 climate pathways. This does not mean climate change is unimportant, since business-as-usual scenarios still lead to a 5.9–11.6 % drop in baseline stocks. Rather, it shows that near-term land-use decisions dominate the outcome in the configuration examined here.

In the context of Bangladesh's Nationally Determined Contributions, which aim for a 15 % emissions reduction by 2030, restoring 50 % of hyper-saline areas to mangroves increases projected 2050 soil organic carbon by 26.6 %, providing a measurable and actionable boost to national carbon management. Because this identical priority ordering extends to the Mississippi, Mekong, Rhine-Meuse, and Yellow River deltas, it is unlikely to be a site-specific anomaly of the Ganges.

Two caveats qualify this result and bound its interpretation. First, the G-RothC model employs a single, homogeneous SOC pool; multi-pool frameworks like RothC or Millennial could compress or expand this ratio depending on how carbon is partitioned between fast and slow reservoirs. Second, these projections lack empirical validation against observed rates of SOC change because the field data rely on a single temporal survey. Consequently, the 4.6–7.2$\times$ ratio must be interpreted as a scenario contrast rather than a calibrated forecast, with the specific values representing management-to-climate leverage under the model's assumptions rather than absolute predictions of stock change.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

5.4 Comparison with the wider literature

The location-grouped R^2 of 0.64 sits within the range reported by SOC prediction studies operating at substantially larger sample sizes: Were et al. (2015) achieved $R^2 = 0.73$ with Random Forest on 321 samples, Zhou et al. (2020) reported 0.78 with XGBoost on 689 samples, Keskin et al. (2024) achieved 0.81 with optimized feature engineering on 284 samples, and the global SoilGrids product achieved 0.56 on over 150,000 samples (Hengl et al., 2017). The comparison is only indicative, since the cited values were not necessarily obtained under location-blocked validation and may themselves be optimistic, and the studies differ in sample size, spatial extent and target variable. Even so, reaching a comparable leak-free R^2 from only 100 samples underscores the value of domain-informed feature engineering over brute-force data scaling. As an independent check on the SOC magnitudes reported here, WoSIS 2023 profiles from the same latitude and longitude window give a mean organic-carbon concentration ($\sim$12 g kg^{-1}) and bulk density ($\sim$1.46 g cm^{-3}) consistent with the sample-level means of this study (16 g kg^{-1} and 1.52 g cm^{-3}).

Within deltaic systems specifically, the observed SOC range (1.1–275.5 t C ha^{-1}) and the 6.5-fold saline-to-freshwater contrast align with published Sundarbans SOC stocks (20–150 t C ha^{-1}; Sarker et al., 2011) and with broader mangrove soil-carbon estimates (100–200 t C ha^{-1} in productive mangrove soils; Alongi, 2012). The management dominance over climate found here extends earlier findings from Indo-Pacific mangrove systems (Lovelock et al., 2017) to a multi-delta context. None of these headline numbers are anomalous. What is new is the integration of the deltas into a unified pipeline where the geochemical, dynamical, predictive, and cross-delta layers are tested against a common dataset, allowing the failure modes of each to be traced to the specific layer at which the framework breaks down.

5.5 Limitations

The primary limitations of this study track its design. (1) The 100-sample field dataset limits statistical power, most acutely for the climate analysis, which operates at the location level ($n = 20$), meaning only strong correlations ($|r| \geq 0.44$ at 80 % power) can be detected. Random cross-validation is inappropriate for the depth-replicated design as it artificially inflates accuracy by $\Delta R^2 = 0.27$. Thus location grouped cross-validation should be regarded as the rigorous standard for future deltaic SOC mapping, and predictions should not be extrapolated to unsampled areas without local calibration (Mahmoudzadeh et al., 2024). (2) The cross-delta panel uses publicly available data with highly uneven sample sizes (Pearl 55 and Senegal 57 against Mississippi 2,043 and Rhine-Meuse 2,652) and inconsistent variable coverage. Particularly the scarcity of co-measured EC, CEC and bulk density across deltas restricted the common predictor set to pH and texture, so the transferability of the geochemistry-informed features could not be tested externally. The datasets are also dominated by temperate deltas. That skew works against rather than for the central finding, since transfer fails even between the data-rich deltas, and the BIS augmentation of Rhine-Meuse left the cross-delta outcome unchanged. Residual non-normality further indicates that the highest-SOC samples are systematically under-predicted. (3) Finally, because the training data derive from a single survey, the process-based projections received no temporal validation and are best read as scenario contrasts rather than calibrated prediction.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

6 Conclusions

Field measurement, process modeling, machine learning and a fifteen-delta transferability test combine here into three integrated findings. First, geochemical controls on Ganges Delta SOC operate hierarchically. EC, CEC and moisture dominate macro-scale variation between zones, with the most saline samples storing about 6.5 times more SOC than freshwater ones, while accurate micro-scale prediction within a zone requires the interaction terms (CEC × pH, clay × moisture) that pooled linear analyses miss. The dominance of EC is scale-dependent, and any digital soil mapping framework operating in coastal systems must explicitly accommodate that behavior.

Second, within the process model, land management outweighed climate forcing on SOC change by a factor of 4.6 to 7.2 across the five deltas examined. Restoring half of the most saline land to mangrove raised projected 2050 SOC by 26.6 %, against a business-as-usual baseline decline of 5.9–11.6 %. The same ordering, management first and climate pathway second, held in every delta tested, spanning temperate, tropical, arid, Mediterranean and subarctic regimes at the 2050 horizon.

Third, with profound implications for scaling deltaic carbon science, the predictive framework demonstrates layered transferability. While individual models succeeded at micro-scales within the same delta (median $R^2 = +0.25$), cross-delta transfer at the macro scale almost universally failed, with only 6 % of 210 directional pairings yielding positive skill (median $R^2 = -0.55$), a failure that persisted even when training on fourteen combined deltas. The higher-level structures remained portable, namely the relative importance of pH and the texture fractions and the functional form of the process model. Trained machine learning weights therefore cannot be swapped between deltas, and local data collection cannot be replaced by data from other deltas, even at scale.

Layered transferability offers a practical, integrated framing for SOC assessment in data-limited deltaic environments. It has direct consequences for blue-carbon accounting, for the use of global SOC products in coastal policy, and for the design of future sampling efforts: prioritize local calibration data, but leverage the parts of the framework that do travel, namely the process forms and the feature-importance structures. Four questions remain open. Whether geochemistry-informed engineered features such as EC × moisture or TN/EC remain effective in deltas with attenuated salinity gradients (Niger, Amazon) has yet to be tested, because those deltas lack the requisite predictors in current global archives. Whether multi-temporal Sentinel-2 and Landsat time series can track SOC change over time, reducing reliance on process-based projection, deserves systematic investigation. Whether deep-soil paleosols below 100 cm contribute meaningfully to deltaic SOC totals is likewise unresolved. And whether the management dominance over climate found in this study holds under multi-pool process models remains to be verified. Methodologically, spatially blocked cross-validation should become the default for deltaic SOC mapping, and prioritizing co-measurement of EC, CEC and bulk density in global soil archives would greatly extend the reach of any future transferability study.

Code and data availability. The analysis code and data underlying the geochemical analyses, the process-based modeling, the leak-free cross-validation and feature-engineering ablation, and the cross-delta transferability results are openly available. Analysis scripts (geochemical hypothesis tests, process model, leak-free cross-validation, engineered-feature ablation, per-zone correlations and descriptive statistics, cross-delta transfer analysis, the polygon-delineation robustness check, a reference implementation of the deep-learning architectures, and figure generation) are archived on GitHub (https://github.com/rakibhhridoy/ganges-delta-soc). The field dataset, derived feature matrix, geomorphological-zone layers, digital elevation model and harmonized cross-delta sample tables are deposited on Zenodo (https://doi.org/10.5281/zenodo.21564960). Scripts are self-contained and reproduce the corresponding reported numbers from fixed random seeds. Cross-delta transferability data were extracted from the WoSIS 2023 December snapshot (Batjes et al., 2024) (ISRIC; CC BY 4.0; https://files.isric.org/public/wosis_snapshot/) and, for Rhine-Meuse, augmented with the BIS Nederland national soil survey (Wageningen Environmental Research, 2020) (4TU.ResearchData; CC BY 4.0; https://doi.org/10.4121/c90215b3-bdc6-4633-b721-4c4a0259d6dc); both extracted subsets, the bounding-box panel (6,652 profiles across 14 deltas) and the polygon panel (1,449 profiles), are included in the archive, together with the active-delta polygons of Nienhuis et al. (2020). The Bangladesh national boundary used for the study-area inset is from geoBoundaries (gbOpen ADM0; CC BY 4.0). Satellite imagery (Landsat-8, Sentinel-2, MODIS, CHIRPS) is publicly available via Google Earth Engine; SoilGrids data (Hengl et al., 2017) are at https://soilgrids.org.

Author contributions. Md Rakib Hasan: conceptualization, methodology, software, formal analysis, investigation, data curation, writing (original draft), writing (review and editing), visualization. Fazla Zawadul Arabi and Mst Anika Khatun Rupa: investigation, data curation, writing (review and editing). Shoumik Zubyer: data curation, writing (review and editing). Mohammad Fazle Alam Rabbi: supervision, writing (review and editing). Md Jashim Uddin: supervision, project administration, writing (review and editing).

Competing interests. The authors declare that they have no competing interests.

Statement on inclusion in global research

This study was conceived, designed and carried out by researchers based in Bangladesh, working in the delta in which they live. Fieldwork, laboratory analysis, modeling and writing were led from the University of Dhaka, and all authors are based in Bangladesh. The choice of study system reflects local research priorities rather than an external agenda: the western Ganges Delta supports more than 160 million people and is among the most under-represented major deltas in global soil archives, a gap this work documents directly. The central finding, that models trained elsewhere do not transfer and that local calibration data remain indispensable, speaks to the same asymmetry from the side of the data-poor delta.

Acknowledgements. The authors acknowledge the use of Google Earth Engine for satellite data acquisition and thank the field assistants who collected soil samples across the Ganges Delta study sites.

Financial support. This research has not received any external funding.

References

Allison, M. A. and Meselhe, E. A.: The sensitivity of the Mississippi delta to river sediment and wave processes, Journal of Coastal Research, 30, 586–608, https://doi.org/10.2112/JCOASTRES-D-13-00060.1, 2014.

Alongi, D. M.: Carbon sequestration in mangrove forests, Carbon Management, 3, 313–322, https://doi.org/10.4155/cmt.12.20, 2012.

Arrouays, D., Grundy, M. G., Hartemink, A. E., Hempel, J. W., Heuvelink, G. B. M., Hong, S. Y., Lagacherie, P., Lelyk, G., McBratney, A. B., McKenzie, N. J., et al.: GlobalSoilMap: Toward a fine-resolution global grid of soil properties, Advances in Agronomy, 125, 93–134, 2014.

Baldock, J. A. and Skjemstad, J. O.: Role of the soil matrix and minerals in protecting natural organic materials against biological attack, Organic Geochemistry, 31, 697–710, https://doi.org/10.1016/S0146-6380(00)00049-8, 2000.

Batjes, N. H., Ribeiro, E., and van Oostrum, A.: Standardised soil profile data to support global mapping and modelling (WoSIS snapshot – December 2023), Earth System Science Data, 16, 4735–4765, https://doi.org/10.5194/essd-16-4735-2024, 2024.

Bond-Lamberty, B. and Thomson, A.: Temperature-associated increases in the global soil respiration record, Nature, 464, 579–582, https://doi.org/10.1038/nature08930, 2010.

Bouyoucos, G. J.: Hydrometer method improved for making particle size analyses of soils, Agronomy Journal, 54, 464–465, 1962.

Bremner, J. M.: Nitrogen—total, in: Methods of Soil Analysis, Part 3: Chemical Methods, edited by Sparks, D. L., pp. 1085–1121, SSSA, 1996.

Chapman, H. D.: Cation-exchange capacity, in: Methods of Soil Analysis, Part 2, edited by Black, C. A., pp. 891–901, American Society of Agronomy, 1965.

Chen, S., Arrouays, D., Leatitia Mulder, V., Poggio, L., Minasny, B., Roudier, P., et al.: A global synthesis of deep learning applications for soil organic carbon estimation from remote sensing data, Earth-Science Reviews, 248, 104 616, https://doi.org/10.1016/j.earscirev.2023.104616, 2024.

Coleman, K., Jenkinson, D., Crocker, G. J., Grace, P. R., Klir, J., Körschens, M., and Richter, D. D.: Simulating trends in soil organic carbon in long-term experiments using RothC-26.3, Geoderma, 81, 29–44, https://doi.org/10.1016/S0016-7061(97)00079-7, 1997.

Conant, R. T., Ryan, M. G., Ågren, G. I., Birge, H. E., Davidson, E. A., Eliasson, P. E., Evans, S. E., Frey, S. D., Giardina, C. P., Hopkins, F. M., et al.: Temperature and soil organic matter decomposition rates—synthesis of current knowledge and a way forward, Global Change Biology, 17, 3392–3404, https://doi.org/10.1111/j.1365-2486.2011.02496.x, 2011.

Davidson, E. A. and Janssens, I. A.: On the variability of respiration in terrestrial ecosystems: Moving beyond Q_{10}, Global Change Biology, 12, 154–164, https://doi.org/10.1111/j.1365-2486.2005.01065.x, 2006.

Donato, D. C., Kauffman, J. B., Murdiyarso, D., Kurnianto, S., Stidham, M., and Kanninen, M.: Mangroves among the most carbon-rich forests in the tropics, Nature Geoscience, 4, 293–297, https://doi.org/10.1038/ngeo1123, 2011.

Ericson, J. P., Vörösmarty, C. J., Dingman, S. L., Ward, L. G., and Meybeck, M.: Effective sea-level rise and deltas: Causes of change and human dimension implications, Global and Planetary Change, 50, 63–82, https://doi.org/10.1016/j.gloplacha.2005.07.004, 2006.

Gorishniy, Y., Rubachev, I., Khrulkov, V., and Babenko, A.: Revisiting deep learning models for tabular data, in: Advances in Neural Information Processing Systems, vol. 34, pp. 18 932–18 943, 2021.

Grinsztajn, L., Oyallon, E., and Varoquaux, G.: Why do tree-based models still outperform deep learning on typical tabular data?, Advances in Neural Information Processing Systems, 35, 507–520, 2022.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

Hengl, T., Mendes de Jesus, J., Heuvelink, G. B. M., Ruiperez Gonzalez, M., Kilibarda, M., Blagotić, A., Shangguan, W., Wright, M. N., Geng, X., Bauer-Marschallinger, B., et al.: SoilGrids250m: Global gridded soil information based on machine learning, PLoS ONE, 12, e0169 748, 2017.

Hossain, M. S., Dearing, J. A., Rahman, M. M., and Salehin, M.: Sediment dynamics and geomorphic changes in the Ganges-Brahmaputra-Meghna delta, Estuarine, Coastal and Shelf Science, 163, 49–62, https://doi.org/10.1016/j.ecss.2015.05.032, 2015.

Keskin, H., Grunwald, S., and Harris, W. G.: Digital mapping of soil organic carbon stocks in a data-scarce environment using machine learning with optimized feature engineering, Geoderma, 442, 116 795, https://doi.org/10.1016/j.geoderma.2024.116795, 2024.

Kleber, M., Eusterhues, K., Keiluweit, M., Georgiou, C., and Liang, C.: Mineral–organic associations: Formation, properties, and relevance in soil environments, Advances in Agronomy, 130, 1–140, https://doi.org/10.1016/bs.agron.2014.10.005, 2015.

Lehmann, J. and Kleber, M.: The contentious nature of soil organic matter, Nature, 528, 60–68, https://doi.org/10.1038/nature16069, 2015.

Lloyd, J. and Taylor, J. A.: On the temperature dependence of soil respiration, Functional Ecology, 8, 315–323, https://doi.org/10.2307/2389824, 1994.

Lovelock, C. E., Feller, I. C., Reef, R., Hickey, S., and Ball, M. C.: Mangrove dieback during fluctuating sea levels, Scientific Reports, 7, 1680, https://doi.org/10.1038/s41598-017-01927-6, 2017.

Mahmoudzadeh, H., Matinfar, H. R., Taghizadeh-Mehrjardi, R., and Kerry, R.: Spatial prediction of soil properties using random forest, k-nearest neighbors and cubist approaches in western Iran, Geoderma Regional, 36, e00 762, https://doi.org/10.1016/j.geodrs.2024.e00762, 2024.

McBratney, A. B., Mendonça Santos, M. L., and Minasny, B.: On digital soil mapping, Geoderma, 117, 3–52, https://doi.org/10.1016/S0016-7061(03)00223-4, 2003.

Murdiyarso, D., Purbopuspito, J., Kauffman, J. B., Warren, M. W., Sasmito, S. D., Donato, D. C., Manuri, S., Krisnawati, H., Taberima, S., and Kurnianto, S.: The potential of Indonesian mangrove forests for global climate change mitigation, Nature Climate Change, 5, 1089–1092, https://doi.org/10.1038/nclimate2734, 2015.

Nienhuis, J. H., Ashton, A. D., Edmonds, D. A., Hoitink, A. J. F., Kettner, A. J., Rowland, J. C., and Törnqvist, T. E.: Global-scale human impact on delta morphology has led to net land area gain, Nature, 577, 514–518, https://doi.org/10.1038/s41586-019-1905-9, 2020.

Padarian, J., Minasny, B., and McBratney, A. B.: Using deep learning for digital soil mapping, Soil, 5, 79–89, https://doi.org/10.5194/soil-5-79-2019, 2019.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in Python, Journal of Machine Learning Research, 12, 2825–2830, 2011.

Rasmussen, C., Heckman, K., Wieder, W. R., Keiluweit, M., Lawrence, C. R., Berhe, A. A., Blankinship, J. C., Crow, S. E., Druhan, J. L., Hicks Pries, C. E., et al.: Beyond clay: Towards an improved set of variables for predicting soil organic matter content, Biogeochemistry, 137, 297–306, https://doi.org/10.1007/s10533-018-0424-3, 2018.

Rogers, K., Kelleway, J. J., Saintilan, N., Megonigal, J. P., Adams, J. B., Holmquist, J. R., and Woodroffe, C. D.: Wetland carbon storage controlled by millennial-scale variation in relative sea-level rise, Nature, 567, 91–95, https://doi.org/10.1038/s41586-019-0951-7, 2019.

Sarker, M. H., Huque, I., Alam, M., and Koudstaal, R.: Soil salinity and carbon sequestration in the coastal zone of Bangladesh, Journal of Soil Science and Environmental Management, 2, 293–303, 2011.

Schmidt, M. W. I., Torn, M. S., Abiven, S., Dittmar, T., Guggenberger, G., Janssens, I. A., Kleber, M., Kögel-Knabner, I., Lehmann, J., Manning, D. A. C., Nannipieri, P., Rasse, D. P., Weiner, S., and Trumbore, S. E.: Persistence of soil organic matter as an ecosystem property, Nature, 478, 49–56, https://doi.org/10.1038/nature10386, 2011.

https://doi.org/10.5194/egusphere-2026-4799
Preprint. Discussion started: 11 September 2026

Setia, R., Gottschalk, P., Smith, P., Marschner, P., Baldock, J., Setia, D., and Smith, J.: Soil salinity decreases global soil organic carbon stocks, Science of The Total Environment, 465, 267–272, https://doi.org/10.1016/j.scitotenv.2012.08.028, 2013.

Shazeer, N., Azizzadenesheli, K., Maas, A., Baraniuk, R., Jaitly, N., and Le, Q. V.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer, in: International Conference on Learning Representations, 2017.

Shwartz-Ziv, R. and Armon, A.: Tabular data: Deep learning is not all you need, Information Fusion, 81, 84–90, https://doi.org/10.1016/j.inffus.2021.11.011, 2022.

Sierra, C. A., Müller, M., and Trumbore, S. E.: Soil organic matter persistence as a stochastic process: Age and transit time distributions of carbon in soils, Global Biogeochemical Cycles, 29, 1361–1372, https://doi.org/10.1002/2014GB005034, 2015.

Six, J., Conant, R. T., Paul, E. A., and Paustian, K.: Stabilization mechanisms of soil organic matter: Implications for C-saturation of soils, Plant and Soil, 241, 155–176, https://doi.org/10.1023/A:1016125726789, 2002.

Sollins, P., Homann, P., and Caldwell, B. A.: Stabilization and destabilization of soil organic matter: Mechanisms and controls, Geoderma, 74, 65–105, https://doi.org/10.1016/S0016-7061(96)00036-5, 1996.

Todd-Brown, K. E. O., Randerson, J. T., Post, W. M., Hoffman, F. M., Tarnocai, C., Schuur, E. A. G., and Allison, S. D.: Causes of variation in soil carbon simulations from CMIP5 Earth system models and comparison with observations, Biogeosciences, 10, 1717–1736, https://doi.org/10.5194/bg-10-1717-2013, 2013.

Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., and Bengio, Y.: Graph attention networks, in: International Conference on Learning Representations, 2018.

Wadoux, A. M. J.-C., Minasny, B., and McBratney, A. B.: Ten challenges for the future of pedometrics, Geoderma, 401, 115 155, https://doi.org/10.1016/j.geoderma.2021.115155, 2021.

Wagai, R. and Mayer, L. M.: Sorptive stabilization of organic matter in soils by hydrous iron oxides, Geochimica et Cosmochimica Acta, 73, 25–35, https://doi.org/10.1016/j.gca.2008.09.034, 2009.

Wageningen Environmental Research: BIS Nederland – Bodemkundig Informatiesysteem reference soil profile database, https://doi.org/10.4121/c90215b3-bdc6-4633-b721-4c4a0259d6dc, cC BY 4.0, 2020.

Walkley, A. and Black, I. A.: An examination of the Degtjareff method for determining soil organic matter, and a proposed modification of the chromic acid titration method, Soil Science, 37, 29–38, 1934.

Were, K., Bui, D. T., Dick, Ø. B., and Singh, B. R.: A comparative assessment of support vector regression, artificial neural networks, and random forests for predicting and mapping soil organic carbon stocks across an Afromontane landscape, Ecological Indicators, 52, 394–403, https://doi.org/10.1016/j.ecolind.2014.12.028, 2015.

Wong, V. N. L., Greene, R. S. B., Dalal, R. C., and Murphy, B. W.: Soil carbon dynamics in saline and sodic soils: A review, Soil Use and Management, 26, 2–11, https://doi.org/10.1111/j.1475-2743.2009.00251.x, 2010.

Zhou, T., Geng, Y., Chen, J., Pan, J., Haase, D., and Lausch, A.: High-resolution digital mapping of soil organic carbon and soil total nitrogen using DEM derivatives, Sentinel-1 and Sentinel-2 data based on machine learning algorithms, Science of the Total Environment, 729, 138 244, https://doi.org/10.1016/j.scitotenv.2020.138244, 2020.