Major points

1. Quantify the loss of predictive performance on the same target profiles

Under the alternative geomorphological delineation, median cross-delta R^2 changes from −0.55 to −0.66, while median within-delta R^2 falls from +0.25 to +0.05. These results show that models using pH and texture generally predict poorly across the deltas examined. However, the two overall medians do not establish how much predictive performance is lost when a model is transferred to each target delta.

A more informative comparison would evaluate the local model, the model transferred directly from another delta, and a simple baseline on the same held-out target profiles. For the few-shot analysis, the local-only and external-plus-local models should use the same number of local training profiles and the same test profiles. Reporting MAE and mean bias alongside R^2 would help distinguish poor predictive discrimination from systematic overprediction or underprediction and from low SOC variability in the target data.

The manuscript should also document how depth ranges, profile aggregation, and analytical methods were harmonised across databases, and identify which results hold in the WoSIS-only analysis. The transferability conclusion should be limited to the datasets, predictors, and models tested. The cross-delta analysis did not test EC, cation exchange capacity (CEC), or the full predictor set used in the local Ganges Delta model.

2. Reconcile SOC concentration, individual-layer stocks, and whole-profile stocks

Table S2 summarises 25 layer samples per zone. The mean across the four zones is approximately 45.1 t C ha^{-1}, which is also the initial value used in the process model. Figure 2d, however, presents stocks for five profiles per zone. If 45.1 t C ha^{-1} is the mean stock of an individual layer, summing the five equal-thickness layers would give a mean 0–100 cm profile stock of approximately 225.5 t C ha^{-1}. This calculation should be checked against the layer-level data.

The manuscript should provide a traceable calculation from measured SOC concentration, bulk density, layer thickness, and any coarse-fragment correction to individual-layer and whole-profile stocks. It should identify the quantity and depth interval underlying both the reported 6.5-fold difference and the 2050 scenario results. Any extrapolation from the sampled profiles to regional stocks or regional restoration benefits also requires a clear explanation of sampling representativeness and area weighting.

3. Reanalyse the salinity relationships at the profile level and verify SOC measurements in saline samples

The 100 layer samples come from 20 profiles, with five profiles per geomorphological zone. The correlation tests, interaction tests, and confidence intervals should account for dependence among layers from the same profile. The abstract describes the within-zone EC–SOC relationship as effectively disappearing, yet Table S10 reports correlations of r = +0.65 and +0.55 in the Mature and Moribund zones, respectively. The uncertainty around these estimates should be reassessed at the profile level. A profile-level resampling procedure or a hierarchical model that accounts for depth would help distinguish within-zone and between-zone patterns. It would also be useful to assess how strongly the pooled association depends on the Tidal Active zone.

High chloride concentrations can cause positive interference in the Walkley–Black determination of SOC. Because this interference could bias the comparison in the same direction as the reported SOC enrichment at high EC, the authors should document chloride concentrations, any correction applied, or independent measurements for representative saline samples. The available material identifies a potential measurement bias that needs to be assessed. Even if the association remains after reanalysis, hydrology, vegetation, land use, and sedimentary setting remain plausible explanations for the observed pattern. A salinity threshold and its proposed effect on decomposition should therefore be treated as hypotheses.

4. Align machine-learning validation with the intended prediction task

Random fivefold cross-validation gives $R^2 = 0.91$, location-grouped fivefold cross-validation gives $R^2 = 0.64$, leave-one-location-out validation gives $R^2 = 0.87$, and leave-one-zone-out validation gives $R^2 = -0.37$. These schemes differ in the spatial structure of the holdout and the amount of training data available. To attribute the difference between 0.91 and 0.64 to validation leakage, the manuscript should confirm that both analyses used the same response variable, predictors, algorithm, and scoring procedure. Table S5 also reports $R^2 = 0.74$ with engineered predictors. The relationship between this result and Table 2 should be made explicit, and any data-driven feature selection should take place entirely within the training folds.

The intended application also needs clarification. Predicting SOC at a site where soil properties have already been measured differs from mapping SOC at an unsampled site. Bulk density enters directly into the stock calculation, so its use as a predictor should be accompanied by an explanation of how it would be obtained at prediction locations. The most informative performance summaries would include profile-level uncertainty, errors relevant to the intended application, and the degree of underprediction for high-SOC samples. The deep-learning comparison could be shortened because the original training code was not retained and the public implementation is a reconstruction.

5. Treat the estimated restoration benefits and process-model transfer as scenario results

The process model uses SOC stocks from a single survey in 2025 to infer steady-state carbon inputs, then uses those stocks to initialise simulations spanning 2000–2050. The restoration scenario prescribes an increase in carbon inputs and a reduction in decomposition, while the reported parameter ensemble mainly varies the base decomposition rate and temperature sensitivity. The projected 26.6% increase in SOC and the 4.6–7.2-fold ratio of the management effect to the difference between the climate pathways are therefore conditional on the scenario design and parameter values. The same direction of response across five deltas does not establish predictive skill in those deltas.

If this analysis is retained, the manuscript should clarify the initialisation year, the stock represented in the model, the area assumed to undergo restoration, and the basis for the management parameters. Sensitivity analyses should examine the assumed strength of restoration, the management parameters, and departures from initial steady state. With the available observations, the results can be presented as exploratory scenarios. Retaining quantitative restoration benefits or cross-delta process-model validity as central findings would require independent temporal observations, restoration controls, or other external validation.