



# Compositional spatial modelling of soil organic and inorganic carbon fractions with calibrated joint uncertainty propagation

Raphael A. Viscarra Rossel<sup>1</sup>, Lewis Walden<sup>1</sup>, and Farid Sepanta<sup>1</sup>

<sup>1</sup>Soil and Landscape Science, School of Molecular and Life Sciences, Curtin University, Bentley, WA 6102, Australia.

**Correspondence:** Raphael A. Viscarra Rossel (r.viscarra-rossel@curtin.edu.au)

**Abstract.** Farm-scale soil-carbon assessments require more than a map on total carbon. They need the organic fractions, the inorganic pool, and calibrated uncertainty around each estimate. We developed a probabilistic compositional framework that propagates uncertainty jointly from mid-infrared (mid-IR) spectroscopic predictions through probabilistic trend and Bayesian spatial modelling. The framework preserves closure among particulate organic carbon (POC), mineral-associated organic carbon (MAOC) and an instrument-defined residual organic carbon (ROC), and preserves mass balance among total organic carbon (TOC), total inorganic carbon (TIC) and total carbon (TC). We applied the framework at a Mediterranean-type semi-arid farm to map POC, MAOC, ROC, TOC, TIC and TC at 0–10, 10–30 and 0–30 cm. Spectroscopic uncertainty was represented by bootstrap prediction distributions, propagated through Natural Gradient Boosting (NGBoost) trend models and Bayesian spatial models based on stochastic partial differential equations (SPDE), estimated using the Integrated Nested Laplace approximation (INLA). Predictive calibration was strong: 95 % probability-integral-transform (PIT) coverage was 0.94–0.95 across all response-depth combinations. Posterior intervals also bracketed bulk laboratory measurements (Kling–Gupta efficiency, KGE 0.64–0.79) and independent measurements (KGE 0.12 for ROC to 0.66 for MAOC, and up to 0.84 in the managed-pasture cohort). The maps showed consistent land-use effects on organic carbon. Cropping, managed pasture and natural vegetation formed the ordering crop < managed < natural for every organic-C pool and depth, with the largest deficits at the surface. Cropping also shifted composition toward the protected pools, with a lower labile-to-protected ratio (POC/[MAOC+ROC]) than pasture. TIC and ROC showed little land-use contrast. Spatial controls differed among pools: gamma-radiometric ratios dominated MAOC, electromagnetic induction conductivity dominated POC at depth, and topographic redistribution organised pools integrating multiple mechanisms. The calibrated posterior, rather than the point estimate, is the appropriate basis for soil-C management, monitoring and accounting.



## 1 Introduction

Soil holds more carbon (C) than the world's vegetation and atmosphere combined, partitioned into a dynamic organic pool and a slow inorganic pool of pedogenic carbonate. Organic-C dynamics directly affect atmospheric CO<sub>2</sub> concentrations and the global climate; inorganic C is a geological-timescale store that nonetheless dominates the inventory of many arid and semi-arid soils (Huang et al., 2024). Organic C is also the substrate of soil function: it drives microbial activity, nutrient cycling, water-holding capacity, and aggregate stability (Hu et al., 2026). A soil losing organic C is both a failing climate sink and a degrading habitat, and the inorganic component sets the geological baseline of the inventory against which those losses are measured. Total organic carbon (TOC), however, hides much of what matters within the organic pool. Its kinetic fractions, the particulate (POC) and mineral-associated (MAOC) organic carbon, differ markedly in turnover times, stabilisation mechanisms, microbial accessibility, and vulnerability to disturbance (Lavalée et al., 2020; Sun et al., 2026, Viscarra Rossel et al. (2019)). A soil dominated by particulate C is not equivalent to one of the same total C content dominated by mineral-associated C.

Semi-arid agroecosystems cover roughly 40 % of Earth's land surface and host much of its rainfed grain production (Plaza et al., 2018). They are also informative natural laboratories for soil-C dynamics. Water-limited primary production and high evapotranspiration result in soils with low organic-C contents and large sensitivity to small changes in climate, vegetation and management (Viscarra Rossel et al., 2024). Pedogenic carbonate accumulates in the same systems, and the total inorganic carbon (TIC) it carries is often comparable to or greater than the organic-C content at depth, so a whole-soil C inventory in semi-arid agroecosystems has to reconcile both forms. Where C is scarce, the partitioning into fast-cycling particulate forms, mineral-stabilised forms, and the more recalcitrant residue determines whether a management gain endures or whether a perturbation releases C to the atmosphere. Yet most field-scale evidence on soil-C dynamics has come from cooler, wetter systems; semi-arid agroecosystems remain underexplored, and their controls and spatial expression less well understood (Zhang et al., 2025b).

Spatial modelling translates point measurements of soil properties into continuous fields, scaling insights from the sample to the field, catchment, and the region. The methods have evolved from geostatistical variants, machine-learning trend models with geostatistical residuals to contextual multiscale methods (Behrens et al., 2010, 2018, 2019) and Bayesian models, which now produce predictions accurate enough to inform management at sub-hectare resolution (Poggio et al., 2016; Huang et al., 2017; Arshad et al., 2020). Yet much of this work in digital soil mapping (McBratney et al., 2003) has tracked the mean prediction more carefully than the uncertainty around it. Measurement error is treated as a single point of truth that downstream models use to estimate the trend, whose predictive distribution is collapsed into a single residual variance. The maps produced conceal the uncertainty around each pixel's value and around the aggregated estimates derived from them. The unquantified uncertainty propagates through the analyses, and the resulting maps are considered more certain than they really are.



Visible–near infrared and mid infrared spectroscopy are now indispensable partners of spatial modelling in soil  
55 science (Viscarra Rossel et al., 2010, 2015). Spectroscopic measurements are rapid, repeatable and non-destructive, a  
single spectrum carries information on multiple soil properties at once, and the cost of a spectroscopic estimate is  
roughly one-tenth that of the corresponding wet-chemistry determination (Viscarra Rossel et al., 2022). Calibrating  
soil properties against laboratory chemistry on a representative subset, then predicting the rest from spectra, lifts  
measurement densities by orders of magnitude. This makes spatial modelling of properties that would otherwise be  
60 too costly to characterise across hundreds of samples, including the granulometric C fractions and TIC used here,  
routinely feasible. We refer to this combination, soil mapping in which the values entering the spatial stage are  
spectroscopically inferred rather than directly measured, as *spectropedometry*.

Despite these gains, every spectroscopic prediction carries an explicit uncertainty, so every value flowing into the  
spatial stage is an estimate rather than a measurement. Ignoring that uncertainty conceals it from the spatial model's  
65 predictive distribution, and it propagates through to the maps as artificial confidence. Several recent studies have  
addressed error propagation in this setting: Heuvelink (2002) formalised uncertainty propagation in soil mapping;  
Viscarra Rossel et al. (2016) propagated uncertainties in farm-scale soil carbon accounting and in continental-scale  
carbon stocks baseline mapping (Viscarra Rossel et al., 2014); Padarian et al. (2019) coupled deep-learning trend  
models to spatial residuals with confidence intervals, and Poggio et al. (2016) and Hu et al. (2025) embedded fully  
70 Bayesian frameworks in their mapping efforts. None of these treat all three uncertainty sources jointly within a single  
posterior while preserving compositional closure across C fractions and totals.

Here, we aim to demonstrate a probabilistic compositional framework with joint uncertainty propagation across  
the spectroscopic, trend and spatial parts, for farm-scale mapping of POC, MAOC, ROC, TOC, TIC and TC. Our  
objectives are to: (i) develop the spatial-modelling framework that propagates spectroscopic, trend-model and spatial  
75 uncertainty into a single posterior while preserving compositional closure between the organic fractions and TOC, and  
mass balance between TOC and TIC; (ii) implement the framework in a Mediterranean-type semi-arid agroecosystem  
to produce, jointly and with calibrated credible intervals at the per-pixel and farm-aggregated scales, the full soil-C  
inventory: the kinetic organic fractions POC and MAOC, total organic carbon (TOC), total inorganic carbon (TIC),  
total carbon (TC = TOC + TIC), and the residual organic-C component (ROC); and (iii) validate and interpret the  
80 fraction-level spatial patterns and the process signatures of the different forms of soil C that measurements of totals  
obscure. We hypothesise that the dominant spatial control on each C fraction reflects its stabilisation mechanism:  
MAOC by mineralogy, POC by soil moisture and conductivity, with topographic redistribution dominating the pools  
that integrate multiple mechanisms at the farm scale where parent-material variation is constrained. The spatial  
residual field is expected to capture fine-scale structure not explained by the covariates, with the trend–spatial  
85 coupling sharpest where the two parts compete. By propagating uncertainty across all three stages, our approach  
delivers per-pixel posteriors and farm-aggregated credible intervals whose calibration is testable against held-out data  
rather than asserted.



## 2 Materials and methods

### 2.1 Study site and sampling

90 The Muresk Institute Farm, in the Western Australian semi-arid region near Northam, is a representative farm of the region's mixed cropping–grazing systems. The property comprises approximately 800 ha. It sits across a landscape of Avon River flats, granite outcrops and dolerite dykes, dominated by duplex soils, which include Chromosols, Kandosols, Sodosols and some Vertosols (Isbell, 2016) near the flats, which broadly correspond in WRB terms to Luvisols/Lixisols, Ferralsols/Luvisols/Lixisols, Solonetz/Planosols and Vertisols, respectively (IUSS Working Group  
95 WRB, 2022). The region experiences a Mediterranean-type semi-arid climate, with mean annual rainfall of around 437–445 mm falling mostly during the May–October growing season, and an estimated mean annual temperature of approximately 17.4°C. Current management is no-till and no-burn, with dryland crop rotations of wheat, canola, lupins and grazing oats and barley, oaten hay and field peas integrated with Merino and Dohne sheep and cattle grazing enterprises.

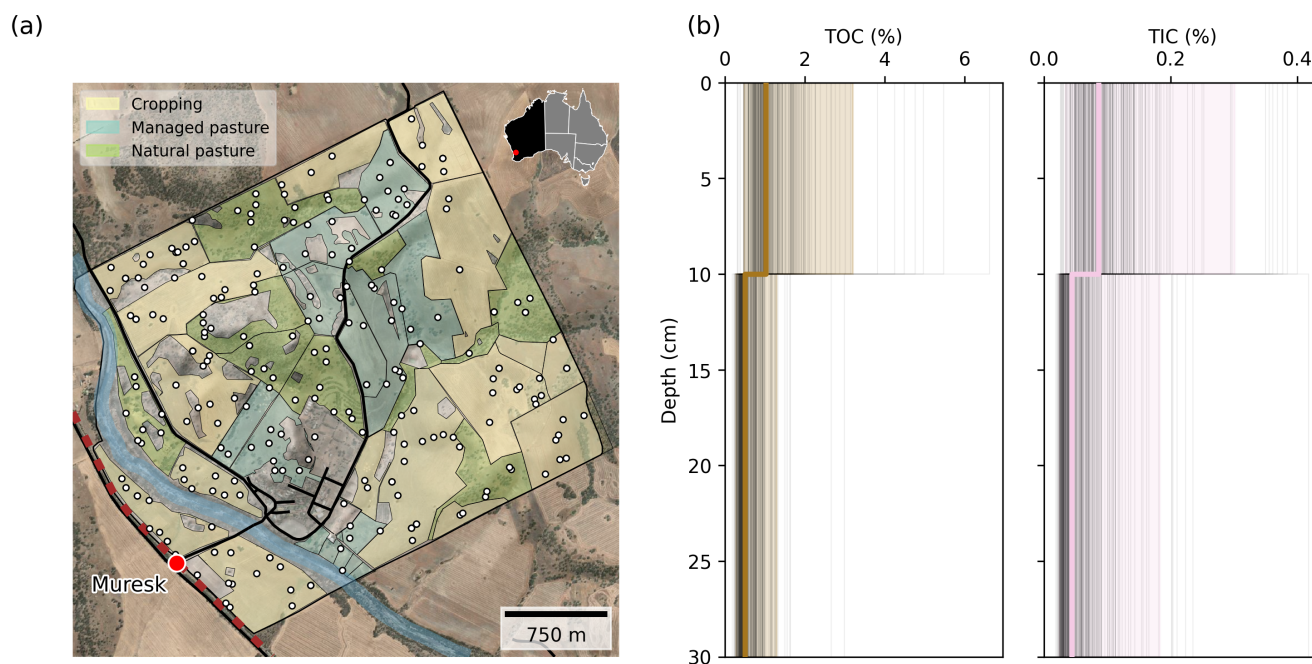
100 The sampling design was similar to that described by Viscarra Rossel et al. (2016), and involved selecting 260 sampling sites (Fig. 1) at random from a stratification of land use by principal component of terrain attributes and electromagnetic-induction and gamma-radiometric sensor maps. The terrain covariates were derived from a 5 m digital elevation model (DEM), and the conductivity and radiometrics from kriging on the 5 m grid (Viscarra Rossel et al., 2016). At each site we sampled two depth intervals, 0–10 cm and 10–30 cm, for 520 specimens in total. Three  
105 cores spaced 1 m apart in a triangular configuration were composited by depth to dampen small-scale variability. The specimens were air-dried, gently crushed, and sieved to < 2 mm before any subsequent analysis. After depth-quality checks and Cook's-distance leverage screening (Section 2.6.2), 250 sites (500 specimens) were retained for the spatial modelling.

### 2.2 Soil analyses

110 We measured the mid-infrared (mid-IR) spectrum of every specimen and a subset of those were selected for direct analytical measurements. The mid-IR measurements are inexpensive and non-destructive; the wet-chemistry measurements are neither. The subset for direct measurements was selected by conditioned Latin hypercube sampling (cLHS) (Ramírez-López et al., 2014) on the mid-IR spectra so that it spans the spectral diversity of the full population. This approach maximised information while minimising analytical costs, as physical fractionation and direct carbon  
115 measurements are substantially more time- and resource-intensive than spectroscopic analysis.

#### 2.2.1 Mid-infrared spectroscopy

A subsample ( $\approx 5$  g) of each specimen was finely ground to pass an 80  $\mu\text{m}$  sieve (Retsch RM 200; Retsch GmbH, Haan, Germany). mid-IR spectra were recorded on a Bruker Invenio HTS-XT diffuse-reflectance Fourier-transform



**Figure 1. Study site and measured depth profiles.** (a) Muresk Institute Farm showing land uses, sampling locations ( $n = 260$  white dots), and surrounding context features (roads, waterway, railway). (b) Measured concentrations of total organic carbon (left) and total inorganic carbon (right) at the two sampled depth intervals (0–10 cm and 10–30 cm). Thin lines represent individual profiles, while the bold steps represent the median. Shaded band represents 95 % coverage.

spectrometer (Bruker Optics, Billerica, MA, USA) over  $4000\text{--}450\text{ cm}^{-1}$  at  $4\text{ cm}^{-1}$  resolution, averaging 64 scans  
 120 per measurement. Soil samples were placed in a 24-well plate with a diffuse-gold reference on the first sample well, which was used to normalise the spectra recorded from the specimens on the plate. All specimens were measured in quadruplicate (technical replicates) and we required a coefficient of variation  $< 2\%$  between replicates and averaged accepted spectra before subsequent analyses.

### 2.2.2 Physical fractionation and operational carbon pools

125 We selected  $n = 50$  representative profiles by cLHS for granulometric fractionation following Walden et al. (2025), yielding  $n = 100$  specimens across the two depth intervals. Briefly, 30 g of  $< 2\text{ mm}$  soil was suspended in 135 mL of deionised water and dispersed with an ultrasonic probe (Sonics VCX 500; Newtown, CT, USA) at  $500\text{ J mL}^{-1}$ , an energy chosen to disperse aggregates without degrading organic matter. The dispersed suspension was wet-sieved through  $250\text{ }\mu\text{m}$  and  $50\text{ }\mu\text{m}$  sieves on an Analysette 3 Pro (Fritsch GmbH, Idar-Oberstein, Germany), giving three  
 130 particle-size fractions:  $250\text{--}2000\text{ }\mu\text{m}$ ,  $50\text{--}250\text{ }\mu\text{m}$ , and  $< 50\text{ }\mu\text{m}$ . The  $< 50\text{ }\mu\text{m}$  fraction was recovered from suspension by gypsum-induced flocculation ( $\text{CaSO}_4 \cdot 2\text{H}_2\text{O}$  at  $1\text{ g L}^{-1}$ ). All three fractions were oven-dried at  $60^\circ\text{C}$  and ground



to  $< 80 \mu\text{m}$  before soil analysis using a thermal ramping analyser (SoliTOC Cube, Elementar Analysensysteme GmbH, Hanau, Germany).

135 From the SoliTOC measurements on the three physical fractions we defined three operational organic-carbon pools at every fractionated specimen, each on a mass-weighted whole-soil basis:

**POC:** the sum of the TOC measured on the coarser soil fractions, weighted by their whole-soil mass proportions,

$$\text{POC} = \text{TOC}_{2000-250\mu\text{m}} m_{2000-250\mu\text{m}} + \text{TOC}_{250-50\mu\text{m}} m_{250-50\mu\text{m}}.$$

**MAOC:** the TOC measured on the silt+clay fraction, weighted by its whole-soil mass proportion,  $\text{MAOC} =$

$$\text{TOC}_{<50\mu\text{m}} m_{<50\mu\text{m}}.$$

140 **ROC:** an instrument-defined residual oxidation component reported by the SoliTOC analyser (Elementar Analysensysteme GmbH, 2023), summed across the physical fractions on a mass-weighted whole-soil basis,

$$\text{ROC} = \text{ROC}_{2000-250\mu\text{m}} m_{2000-250\mu\text{m}} + \text{ROC}_{250-50\mu\text{m}} m_{250-50\mu\text{m}} + \text{ROC}_{<50\mu\text{m}} m_{<50\mu\text{m}}.$$

Following the temperature-programmed protocol (DIN, 2016), the SoliTOC Cube heats the sample at  $70^\circ\text{C min}^{-1}$  and integrates the evolved  $\text{CO}_2$  over three temperature windows:  $\text{TOC}_{400}$  at  $400^\circ\text{C}$  in  $\text{O}_2$ , ROC between  $400$  and  $600^\circ\text{C}$ , and  $\text{TIC}_{900}$  between

145  $600$  and  $900^\circ\text{C}$ ; we used the variant with carrier-gas switching ( $\text{N}_2$  pyrolysis at  $900^\circ\text{C}$  followed by  $\text{O}_2$  re-oxidation) for sharper ROC–TIC separation. The manufacturer markets the ROC step as elemental carbon (Elementar Analysensysteme GmbH, 2023), but the combustion window is not particle-selective; we therefore treat ROC as an

analyser-defined residual and return to the question of its biogeochemical identity in the Discussion.

### 2.2.3 Bulk carbon analysis

150 We selected  $n = 200$  representative samples by cLHS on the mid-IR spectra for direct total organic and total inorganic carbon measurements. Each specimen was analysed on a SoliTOC Cube temperature-ramp protocol described above,

returning  $\text{TOC}_{400}$ , ROC and  $\text{TIC}_{900}$  from the three combustion windows; the analyser reports total carbon as  $\text{TC} = \text{TOC}_{400} + \text{ROC} + \text{TIC}_{900}$  rather than as an independent integration. We used the bulk-measured TOC, TIC and TC directly. The bulk-soil ROC reported by the SoliTOC was measured but not used.

## 155 2.3 Spectroscopic prediction of soil-C variables

From the mid-IR spectra of all 500 specimens we estimated the soil-C variables for the remaining analyses. The organic-C composition was spectroscopically modelled in the ‘sums-to-total’ scheme of Zhang et al. (2025a). The

four sub-fractions ( $\text{POC}_{2000-250\mu\text{m}}$ ,  $\text{POC}_{250-50\mu\text{m}}$ , MAOC, ROC) were rescaled per specimen to sum to one and modelled in centred-log-ratio (CLR) coordinates against the spectra; bulk-TOC was modelled separately in  $\log_{10}$

160 units. At the back-transform, each fraction-bootstrap’s softmax-recovered proportions were scaled by the matching bulk-TOC bootstrap, so that the four fraction predictions sum to the bulk-TOC prediction at every bootstrap. Bulk

TIC and TC were modelled directly in  $\log_{10}$  units. The spectroscopic models were derived for the two sub-fractions,



POC<sub>2000–250 $\mu$ m</sub> and POC<sub>250–50 $\mu$ m</sub> defined above, but for the spatial analyses they were summed per bootstrap at the back-transform to produce a single POC fraction. Thus, the soil C variables for the spatial modelling were POC, MAOC, ROC, TOC, TIC and TC.

We preprocessed the spectra with a Savitzky–Golay first-derivative transform (window 17, quadratic polynomial; Savitzky and Golay (1964)) to suppress baseline drift and accentuate the molecular absorptions. We fitted both support-vector-machine (SVM) regressions with a radial-basis kernel (Cortes and Vapnik, 1995) and Cubist regression trees (Quinlan, 1993). For each response we retained the model (SVM or Cubist) that gave the better cross-validated  $R^2$  and Lin’s concordance correlation coefficient ( $\rho_c$ ) on the back-transformed (%C) scale, and used it for all subsequent bootstraps and propagation (Table 1). Hyperparameters (cost  $C$ , kernel bandwidth  $\gamma$  for SVM; number of committees and neighbours for Cubist) were tuned by random search inside a 10-fold cross-validation (CV) on the calibration set; folds were grouped by site so that paired 0–10 cm and 10–30 cm samples never split across folds, and we selected the combination minimising the root mean-squared error (RMSE).

Spectroscopic uncertainty was quantified by the non-parametric bootstrap ( $B = 30$ ; Hastie et al. (2009); Viscarra Rossel (2007)). For each bootstrap, we refitted the model retained for that response (SVM or Cubist; Table 1) and predicted in transformed units ( $\log_{10}$  for the totals; CLR for the fractions), then back-transformed per bootstrap to %C. The 30 back-transformed predictions per specimen and per variable were retained as a discrete predictive distribution and carried forward into the multiple-imputation (Section 2.6); from these we report mean, median and 95 % intervals (q025–q975) at each specimen for context. Cross-validation performance per response is reported in the Results (Table 1).

## 2.4 Proximal sensing and terrain covariates

We assembled a set of predictors that could capture the controls on soil-C variation at the farm scale: terrain, proximally-sensed soil conductivity, and gamma radiometrics. Climate variables were not used as the study extent is too small to contain meaningful within-farm variation in temperature or rainfall.

Terrain attributes were derived in GRASS GIS (GRASS Development Team, 2025) from the 5 m DEM: slope, aspect, profile and tangential curvature, flow accumulation, the topographic wetness index (TWI) and topographic convergence index (TCI), and the first- and second-order directional derivatives (dxEW, dyNS, dxx, dxy, dyy). To capture the multi-scale character of soil-C variation, each terrain attribute was decomposed with the discrete wavelet transform (DWT) using a Daubechies-2 wavelet at five scales (the original 5 m grid plus DWT levels 3–6; Daubechies (1988); Lark and Webster (1999)) and reconstructed back to the 5 m grid. Apparent soil electrical conductivity was acquired with a DUALEM electromagnetic-induction (EMI) instrument on a mobile platform; the two coil-pair geometries returned the shallower DUALEM 50 and the deeper DUALEM 150 conductivity readings. These, together with gamma-radiometric K, Th, U and total counts, were kriged to the 5 m grid. All rasters were projected to EPSG:32750, resampled by bilinear interpolation, and aligned to the 5 m grid.



To reduce dimensionality and avoid overfitting, we performed a principal component analysis (PCA) on the terrain attributes scale-by-scale. At each scale, the standardised attributes were projected onto the principal components (PCs), and we retained those that individually explained at least 10 % of the variance. The second PC carried rectilinear ridge-like artefacts that projected into the prediction maps as diagonal stripes; we dropped it at every  
200 scale. The candidate covariate set passed to NGBoost combined the retained terrain PCs with the DUALEM 50 and DUALEM 150 conductivities, the gamma-radiometric K, Th, U and total counts, sensor ratios (the DUALEM 50 / DUALEM 150 conductivity ratio; the gamma ratios K/Th, K/U, Th/U; and the K-normalised count K/total), and multiscale fine-minus-coarse terrain differences (TWI, TCI, profile and tangential curvature).

The candidate covariate sets were compared by 5-fold cross-validated performance of natural gradient boosting  
205 [NGBoost; Duan et al. (2020)], selecting the combination that gave the lowest mean negative log-likelihood (NLL). NGBoost was tuned across a 32-combination grid: learning rate  $\in \{0.01, 0.025\}$ , minibatch fraction  $\in \{0.5, 1.0\}$ , column subsample  $\in \{0.5, 1.0\}$ , decision-tree base learner with `max_depth`  $\in \{3, 4\}$  and Friedman-MSE splits, and scoring rule  $\in \{\text{LogScore}, \text{CRPScore}\}$  (the log score and the continuous ranked probability score);  $n_{\text{estimators}} = 1500$  acted as a cap with early stopping after 50 rounds on a 20 % within-fold validation set. LogScore was selected for all  
210 but three variable–depth combinations, where CRPScore had marginally lower NLL; selected ensemble sizes ranged from 23 to 139 trees.

## 2.5 Compositional structure of soil-C variables

The three organic-C fractions (POC, MAOC, ROC) sum to TOC, so the absolute values and their proportions of TOC carry a redundant degree of freedom: only two of the three can vary independently. Compositional data analysis  
215 (Aitchison, 2003) handles this by mapping the composition into unconstrained Euclidean coordinates that an ordinary Gaussian model can work on. The CLR transform,  $\text{CLR}(\mathbf{c}) = \log \mathbf{c} - \frac{1}{3} \sum_i \log c_i$ , returns a three-vector that sums to zero and is unchanged when the three values are scaled by a common factor,  $\text{CLR}(k\mathbf{c}) = \text{CLR}(\mathbf{c})$  for any  $k > 0$ . The redundancy remains, however, because the three CLR components always sum to zero, so they cannot be modelled as three independent Gaussians. The isometric-log-ratio (ILR) transform resolves this by rotating CLR onto two new  
220 coordinates that are free to vary independently,

$$\text{ILR}(\mathbf{c}) = \text{CLR}(\mathbf{c}) H,$$

where  $H \in \mathbb{R}^{3 \times 2}$  is an orthonormal basis (its columns sum to zero, are unit-length and orthogonal). Because the columns of  $H$  sum to zero, the geometric-mean centring in CLR cancels under the rotation, so equivalently  $\text{ILR}(\mathbf{c}) = (\log \mathbf{c}) H$ .

225 We chose a custom orthonormal balance basis  $H$  (Egozcue and Pawlowsky-Glahn (2005)) for biogeochemical interpretability: the first balance contrasts the labile pool against the protected pool (POC vs. MAOC+ROC); the



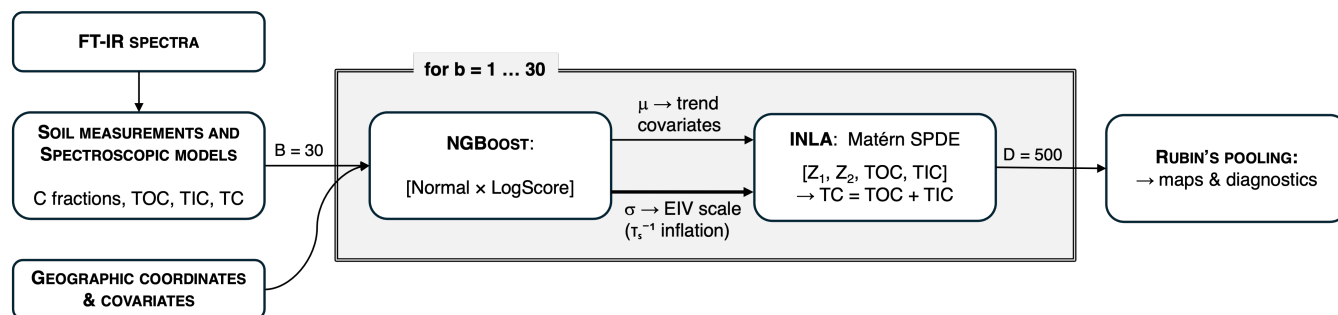
second contrasts mineral-associated against residual carbon (MAOC vs. ROC). The resulting ILR coordinates are

$$Z_1 = \frac{1}{\sqrt{6}} [2 \log \text{POC} - \log \text{MAOC} - \log \text{ROC}], \quad Z_2 = \frac{1}{\sqrt{2}} [\log \text{ROC} - \log \text{MAOC}]. \quad (1)$$

Total carbon also decomposes as  $\text{TC} = \text{TOC} + \text{TIC}$ , but the two summands are not exchangeable parts of one process: TOC reflects biological inputs and turnover, while TIC reflects parent-material weathering, pedogenic-carbonate accumulation and potentially also soil management. Imposing a single compositional ILR parameterisation on (TOC, TIC) would conflate these distinct processes. We therefore modelled  $\log(\text{TOC})$  and  $\log(\text{TIC})$  as two independent latent processes, each with its own covariates and SPDE, and derived  $\text{TC} = \text{TOC} + \text{TIC}$  at the prediction step. Mass conservation is preserved at every posterior draw, and the joint posterior dependence between the two pools was propagated into the predictive distribution of TC by summing draw-by-draw.

## 2.6 Hierarchical modelling and multiple-imputation propagation

The modelling has three parts (Fig. 2): (1) Spectroscopic modelling of the soil C variables with a bootstrapped prediction to quantify spectroscopic uncertainty; (2) NGBoost probabilistic gradient boosting of those predictions on environmental covariates, capturing the non-linear trend together with its per-point predictive distribution; and (3) Bayesian spatial modelling with the integrated nested Laplace approximation (INLA) on the residuals, using a Matérn stochastic partial differential equation (SPDE) random field. In part two, the predictive standard deviation enters the observation likelihood as an errors-in-variables (EIV) variance ( $\tau_s^{-1}$ ). Thus, each stage carries a distribution, and we propagated the spectroscopic, trend, and spatial-model uncertainties jointly by multiple imputation (Section 2.7).



**Figure 2. Hierarchical pipeline.** The spectroscopic model is fit once and returns  $B = 30$  bootstrap predictions per specimen per soil C variable. The multiple-imputation loop iterates over  $b = 1, \dots, 30$ : each iteration refits NGBoost, passes its predictive mean  $\mu$  into INLA as a trend covariate and its predictive variance  $\sigma^2$  as the site-specific observation-error variance in the errors-in-variables (EIV) likelihood (Eq. 3), fits four Matérn-SPDE models for  $Z_1$ ,  $Z_2$ , TOC and TIC, and draws  $D = 500$  posterior samples, with TC derived per draw as  $\text{TC} = \text{TOC} + \text{TIC}$ . Pooling across imputations produces the per-pixel predictive distributions, maps and diagnostics.



### 2.6.1 Spectroscopic predictions

245 The spectroscopic modelling and prediction is described in Section 2.3. For each specimen and each soil C variable, the input to part two is the  $B = 30$ -vector of bootstrap predictions on %C.

### 2.6.2 NGBoost probabilistic regression on environmental covariates

For the trend model we use NGBoost, which updates the parameters of a chosen distribution family with the natural gradient of a proper scoring rule (Duan et al., 2020). Each fit returns a full predictive distribution per point in a  
250 single pass. The per-point standard deviation enters the spatial model (part three) as the measurement-error variance in an EIV likelihood (Eq. 3), and calibration of the predictive spread enters the model selection through the NLL scoring rule.

We fit NGBoost separately for each soil C variable and depth interval on the spectroscopic predictions and the selected covariates, modelling the response as Gaussian and selecting the proper scoring rule (LogScore or CRPScore) by  
255 cross-validated NLL (Duan et al., 2020). Hyperparameters were tuned per variable and depth by 5-fold cross-validation as described above (Section 2.4), with NLL as the selection metric. NGBoost was fit per depth, so each fold contains at most one specimen per site (site-grouping reduces to plain  $k$ -fold). We then filtered features within the chosen set by permutation importance on a single hold-out fit, retaining those with a positive lower 95 % CI and importance exceeding 5 % of the baseline RMSE. The chosen hyperparameters and selected covariates were set at the baseline  
260 spectroscopic mean and held fixed across multiple-imputation iterations; only the model itself is refit each iteration on the drawn calibration set. Influential outliers were screened by Cook's distance (Cook, 1977) via a ridge regression surrogate and removed when they distorted the trend.

To prevent the spatial model from overfitting, we passed cross-validated NGBoost predictions  $\hat{y}_{cv}$  as the covariate to INLA, together with their per-point predictive standard deviations  $\hat{\sigma}_{cv}$  which carry the trend uncertainty into the  
265 spatial modelling (Section 2.7).

NGBoost was fitted in the same transformation the spatial model received: CLR for the C fractions,  $\log_{10}$  for TOC and TIC. No reprojection onto the simplex was needed between stages as CLR is invariant under closure,  $\text{CLR}(\mathbf{c}) = \text{CLR}(k\mathbf{c})$  for any positive composition, so applying CLR to the spectroscopic %C predictions reproduces the spectroscopic model's coordinates regardless of how the fractions sum. We checked that predictions were positive  
270 and within range at the start of each multiple-imputation iteration.

### 2.6.3 Bayesian spatial model via INLA

The spatial model splits the observed value at each location into three parts: a trend that the environmental covariates explain, a spatial random field that absorbs remaining structure varying smoothly between nearby points, and a residual observation error. We modelled the spatial field as a Matérn Gaussian random field, the standard choice



275 in geostatistics, with a range parameter that sets how far the spatial correlation extends and a marginal standard deviation that sets its amplitude. For each response, depth and multiple-imputation iteration, the spatial model was fitted in the transformed response space,

$$Y(\mathbf{s}) = \beta_0 + \beta_1 \hat{y}_{cv}(\mathbf{s}) + w(\mathbf{s}) + \varepsilon(\mathbf{s}), \quad \varepsilon(\mathbf{s}) \sim \mathcal{N}(0, \tau_{\mathbf{s}}^{-1}), \quad (2)$$

where  $Y(\mathbf{s})$  is the observed transformed response at location  $\mathbf{s}$ : ILR units for  $Z_1$  and  $Z_2$ , and  $\log_{10}$  units for TOC  
 280 and TIC. The term  $\beta_0 + \beta_1 \hat{y}_{cv}(\mathbf{s})$  is the calibrated trend, with  $\hat{y}_{cv}(\mathbf{s})$  the out-of-fold NGBoost prediction at sampled locations;  $w(\mathbf{s})$  is the zero-mean Matérn field; and  $\varepsilon(\mathbf{s})$  is residual observation error with location-specific variance  $\tau_{\mathbf{s}}^{-1}$ . This residual variance is inflated by NGBoost's per-point predictive variance so that observations whose trend is more uncertain are down-weighted in the spatial fit,

$$\tau_{\mathbf{s}}^{-1} = \tau^{-1} \left[ 1 + \hat{\sigma}_{cv}^2(\mathbf{s}) / \tilde{\sigma}^2 \right], \quad \tilde{\sigma}^2 = \text{median}_{\mathbf{s} \in \mathcal{S}_{\text{obs}}} \hat{\sigma}_{cv}^2(\mathbf{s}). \quad (3)$$

285 This EIV term propagates NGBoost trend uncertainty directly into the SPDE fit. The intercept  $\beta_0$  and slope  $\beta_1$  allow the spatial model to recalibrate the NGBoost trend while estimating residual spatial structure.

For each ILR coordinate, the covariate  $\hat{y}_{cv}$  in Eq. 2 is the rotation of the NGBoost CLR prediction onto column  $k$  of  $H$ :

$$\hat{y}_{cv}^{(Z_k)} = \hat{\mathbf{c}}_{cv} H_{:,k}, \quad k = 1, 2.$$

290 Because the coordinates of  $H$  sum to zero, any common shift across the three CLR components cancels, so the rotation is well-defined whether or not the three independently-fitted NGBoost models predict CLR vectors that strictly close. No back-transform through the simplex is taken between parts two and three.

We approximated  $w(\mathbf{s})$  by the SPDE representation of Lindgren et al. (2011), which casts the Matérn field as a numerical solution on a triangular mesh of points spanning the study area. This makes the field computationally  
 295 tractable on a fine grid while preserving the Matérn covariance. The SPDE,

$$(\kappa^2 - \Delta)^{\alpha/2} (\tau_w w(\mathbf{s})) = \mathcal{W}(\mathbf{s}),$$

where  $\Delta$  is the Laplacian and  $\mathcal{W}$  is spatial Gaussian white noise, has solutions whose covariance is Matérn with smoothness  $\nu = \alpha - 1$  in two dimensions. We fixed  $\alpha = 2$ , giving the once-differentiable surface ( $\nu = 1$ ) typical for environmental fields. The remaining parameters ( $\kappa, \tau_w$ ) were reparameterised into the practical range  $\rho = \sqrt{8\nu}/\kappa$  (the  
 300 distance over which the spatial correlation drops to about 0.13) and the marginal standard deviation  $\sigma = (\tau_w \kappa \sqrt{4\pi})^{-1}$ , on which we placed priors directly. Inference used INLA (Rue et al., 2009) via R-INLA. The SPDE mesh was built with `inla.mesh.2d` with inner and outer maximum edge lengths of 70 m and 350 m, a cutoff of 70 m, and inner and outer offsets of 175 m and 1050 m beyond the convex hull of the sample locations.



We placed a joint penalised-complexity prior on  $(\rho, \sigma)$  (Fuglstad et al., 2019), which shrinks the spatial field towards  
305 a simpler baseline of no spatial structure unless the data clearly support it,

$$\Pr(\rho < 50\text{m}) = 0.05, \quad \Pr(\sigma > 1) = 0.05.$$

For the residual precision  $\tau$  we used the default INLA log-Gamma prior. For TOC and TIC we worked in  $\log_{10}$  units; both are strictly positive and the transform stabilises the variance. The transform is fixed across multiple-imputation iterations and inverted draw-by-draw on back-transformation.

310 We fitted four INLA models per depth: (i) a non-spatial regression on  $Z_1$ . In preliminary fits the SPDE for  $Z_1$  could not be identified. The marginal variance posterior collapsed onto its prior, so we dropped the spatial term for this coordinate; (ii) an SPDE on  $Z_2$ ; (iii) a scalar SPDE on TOC (transformed units); and (iv) a scalar SPDE on TIC (transformed units).

From the joint posterior of each model we drew 500 samples per multiple-imputation iteration and back-transformed  
315 every draw to units of %C. The back-transforms were applied draw-by-draw to preserve the compositional constraints, the mass-conservation identity  $\text{TC} = \text{TOC} + \text{TIC}$ , and the joint posterior dependence between pools.

**Organic-C simplex.** For each draw  $i$  we have  $(Z_1^{(i)}, Z_2^{(i)}, \log_{10} \text{TOC}^{(i)})$  in transformed units ( $Z_1, Z_2$  in ILR coordinates; TOC in  $\log_{10}$ ). Inverting the  $\log_{10}$  on the total recovers  $\text{TOC}^{(i)}$  in %C. The fraction proportions follow from the inverse of the orthonormal balance basis: let  $\mathbf{I}^{(i)} = H(Z_1^{(i)}, Z_2^{(i)})^\top$  denote the CLR of the C composition,  
320 where the columns of  $H \in \mathbb{R}^{3 \times 2}$  are the orthonormal balance basis contrasts implicit in Eq. 1. Closure of  $\exp(\mathbf{I}^{(i)})$ —equivalently, softmax, yields the fraction proportions  $\boldsymbol{\pi}^{(i)}$ , and absolute concentrations are

$$(\text{POC}, \text{MAOC}, \text{ROC})^{(i)} = \text{TOC}^{(i)} \boldsymbol{\pi}^{(i)}.$$

**Total carbon by mass conservation.** Total carbon was derived per draw as the sum of the two independently-modelled pools, after inverting each pool's saved transform:

325  $\text{TC}^{(i)} = \text{TOC}^{(i)} + \text{TIC}^{(i)}.$

Mass conservation holds at every posterior draw. The joint posterior correlation between TOC and TIC, induced jointly by the shared environmental covariates and by the (independent) spatial fields, propagates into the predictive distribution of TC when the two pools are sampled together and summed draw-by-draw.

## 2.7 Multiple-imputation propagation of uncertainty

330 Spectroscopic and trend uncertainty are propagated forward jointly by multiple imputation (Rubin, 1987). At iteration  $m = 1, \dots, M$  (with  $M = 30$ ), we draw one bootstrap realisation per specimen and per response from the spectroscopic predictive distribution. We then refit NGBoost once per variable on the drawn calibration set, evaluate its predictive



distribution at training and prediction points to obtain a per-point  $(\hat{y}_{cv}(\mathbf{s}), \hat{\sigma}_{cv}(\mathbf{s}))$ , refit INLA with  $\hat{y}_{cv}$  as the trend  
covariate and  $\hat{\sigma}_{cv}$  propagated into the observation as the EIV likelihood of Eq. 3, and sample 500 joint posterior  
335 draws. Each draw was back-transformed to %C before pooling. A single NGBoost fit per (variable, iteration) returns  
the predictive mean and standard deviation directly; the five per-variable fits within each iteration (POC, MAOC,  
ROC, TOC, TIC) are independent and run concurrently.

We pool across iterations using Rubin’s rules (Rubin, 1996). The pooled posterior mean is the average of the  
within-iteration means, and the total predictive variance combines the within-imputation variance  $\bar{U}$  and the between-  
340 imputation variance  $\mathcal{B}$  as  $V = \bar{U} + (1 + 1/M)\mathcal{B}$ , where the  $(1 + 1/M)$  inflation accounts for the finite number of  
imputations.

We do not summarise per-pixel intervals as  $\hat{\mu} \pm 1.96\sqrt{V}$  because the back-transformed posterior is non-Gaussian—  
both the  $\log_{10}$  inversion on the totals (TOC and TIC) and the softmax closure of the balance coordinates ( $Z_1$   
and  $Z_2$ ) onto the  $C$  simplex are nonlinear, so the symmetric Gaussian summary misplaces the tails (Zhou and  
345 Reiter, 2010). We therefore report 95 % credible intervals as the empirical 2.5 / 97.5 % quantiles of the union of all  
 $M \times 500 = 15,000$  back-transformed draws per pixel, and the median raster as the empirical median of that pool,  
which is the mixed-draws approximation of Gelman et al. (2004).

## 2.8 Diagnostics

We assessed our modelling with six diagnostics, each exercising a different aspect of the joint model. Reported  
350 numbers are pooled across the  $M = 30$  multiple imputations unless stated otherwise.

**Variance decomposition.** To trace where uncertainty originates, we allocated the per-pixel predictive variance  
in transformed units to five components: spectroscopic (between-iteration variance of the linear-predictor mean),  
the trend variance  $\text{Var}(T)$  from NGBoost (where  $T$  denotes the trend stage), the SPDE variance  $\text{Var}(S)$  from the  
Matérn random field (where  $S$  denotes the spatial stage), the trend–SPDE coupling  $c = 2\text{Cov}(T, S)$  (signed), and the  
355 residual observation noise  $1/\tau$ . The decomposition tells us which stage of the analyses drives the credible interval at  
each response and depth combination, and where the trend and SPDE compete to explain the same spatial signal.  
Components are reported as percentages of total predictive variance and sum to 100 % per pixel. The coupling term  
is signed; when the trend and SPDE are anti-correlated (negative coupling), the marginal trend and SPDE shares  
can each individually exceed 100 % of the total predictive variance, and the negative coupling absorbs the cross-stage  
360 covariance so the row still sums to 100 %. We keep the trend–SPDE covariance as a separate signed coupling term  
rather than folding it into the marginal stages, so the cross-stage interaction remains visible. Components are pooled  
across multiple-imputation iterations as the median and inter-quartile range to dampen any occasional far-tail  
 $\sigma$ -hyperparameter draws.



**Observed-vs-fitted performance.** Performance was assessed on held-out CV residuals from the per-iteration model fit. We report the Kling–Gupta efficiency (KGE) as the skill metric alongside RMSE and mean error (ME). KGE =  $1 - \sqrt{(r - 1)^2 + (\alpha - 1)^2 + (\beta - 1)^2}$  decomposes skill into the Pearson correlation  $r$ , the variance ratio  $\alpha = \sigma_{\hat{y}}/\sigma_y$ , and the bias ratio  $\beta = \mu_{\hat{y}}/\mu_y$  (Gupta et al., 2009), rather than the Nash–Sutcliffe efficiency, because KGE penalises the variance shrinkage that Bayesian spatial smoothers introduce. We plot posterior-mean predictions against observations in transformed units, the scale on which the spatial model fits, where any systematic departure from the 1:1 line is most easily seen.

**Predictive 95 % coverage.** For each held-out CV observation  $y_i$  we evaluated the posterior predictive distribution at the observed value,  $u_i = F_i(y_i)$ —the probability-integral transform (PIT) (Gneiting et al., 2007). Under a calibrated posterior the  $u_i$  are uniform between 0 and 1; we report 95 % coverage as the fraction of  $u_i$  falling within [0.025, 0.975]. Departures from nominal indicate under- or over-dispersion of the predictive intervals.

**Nested model comparison.** Deviance Information Criterion (DIC), Widely Applicable Information Criterion (WAIC), Log-Predicted Marginal Likelihood (LPML) for null, trend-only (NGBoost), SPDE-only and full models, per response. We report each metric relative to the null model, signed so that positive values indicate improvement of the variant over the null baseline. Because DIC and WAIC are lower-is-better while LPML is higher-is-better, the deltas are:

$$\Delta\text{DIC} = \text{DIC}_{\text{null}} - \text{DIC}_{\text{variant}}, \quad \Delta\text{WAIC} = \text{WAIC}_{\text{null}} - \text{WAIC}_{\text{variant}}, \quad \Delta\text{LPML} = \text{LPML}_{\text{variant}} - \text{LPML}_{\text{null}}.$$

Each iteration refits the four variants on the same INLA stack and the deltas are pooled across the  $M = 30$  multiple-imputation iterations.

**Coherence check against laboratory measurements.** We made two complementary checks against the bulk analytical (SoliTOC) observations at sample sites. (i) Compositional closure: the back-transformed sum  $\text{POC}^{(i)} + \text{MAOC}^{(i)} + \text{ROC}^{(i)}$  from the ILR fraction model compared against the bulk-measured TOC; this exercises the spectroscopic, NGBoost and SPDE stages jointly. (ii) TC mass-balance: the per-draw derived total  $\text{TC}^{(i)} = \text{TOC}^{(i)} + \text{TIC}^{(i)}$ , summed from the two independently-modelled SPDEs, compared against the bulk-measured combustion TC; this tests the internal consistency of the two C pools against the analytical measurement of TC.

**Validation at monitoring sites.** A subset of  $n = 59$  sites was resampled in a subsequent year (118 specimens at 0–10 cm and 10–30 cm) and analysed by SoliTOC on the same granulometric protocol. The monitoring measurements were not used to calibrate any stage of the pipeline. We compared the joint posterior from the first time sampling at each monitoring location (54 of 59 sites within 1 m of the first-time coring location, the remaining five within 168 m) to the monitoring bulk and granulometric values, stratified by land-use class at the first time sampling. We report KGE, RMSE, ME and the proportion of monitoring observations bracketed by the first-time 95 % credible interval, both pooled and per land-use class.



## 2.9 Compositional metrics and land-use contrasts

To interpret the fraction-level spatial patterns and quantify how they vary with land use and with environmental drivers, we compute three sets of derived quantities from the joint posterior.

**Organic-C composition and labile-to-protected ratio.** We display the compositional balance between the three C fractions in a ternary RGB map: at every pixel and every posterior draw, the proportions of POC, MAOC and ROC in TOC are encoded as the red, green and blue channels of an RGB triplet, with the per-pixel posterior median shown in the figures. For each pixel and posterior draw we compute the labile-to-protected ratio (L:P),

$$L:P = \frac{\text{POC}}{\text{MAOC} + \text{ROC}},$$

the labile particulate pool relative to the combined protected pools. Other authors have used L:P as a vulnerability index, a measure of the share of organic C at risk of disturbance-driven loss (Viscarra Rossel et al., 2019); we adopt the ratio without committing to that single interpretation and discuss the agroecological interpretation in the Discussion.

**Inorganic share of total carbon.** For each pixel and posterior draw we compute

$$\text{TIC}/\text{TC} = \frac{\text{TIC}}{\text{TOC} + \text{TIC}},$$

the share of total carbon held in the inorganic pool, with  $\text{TC}^{(i)} = \text{TOC}^{(i)} + \text{TIC}^{(i)}$  taken per draw so that the ratio inherits propagated uncertainty from both pools and lies in  $[0, 1]$ . The ratio reveals where the inorganic pool holds an appreciable share of TC even when both absolute concentrations are low, a contrast the absolute-TIC map (Fig. 4 d-f) does not display directly.

**Land-use contrasts.** The contrast target is the per-class areal-mean mapped concentration across the farm under the joint posterior, integrated over all prediction-grid pixels in each LU class, so the spectroscopic, trend and SPDE uncertainty all propagate into the contrast, and the inference is over the mapped population rather than the calibration sample. For each posterior draw and each land-use (LU) class (cropping, managed pasture, natural vegetation) we computed the spatial-mean concentration over all prediction-grid pixels in that class, then took pairwise differences between LU classes paired within the same draw,

$$\bar{Y}_{L,m,d} = \frac{1}{n_L} \sum_{p \in L} Y_{p,m,d}, \quad \Delta_{m,d}^{L_1,L_2} = \bar{Y}_{L_1,m,d} - \bar{Y}_{L_2,m,d},$$

where  $m$  indexes the  $M = 30$  multiple-imputation iterations and  $d$  the 500 INLA posterior draws per iteration. The within-draw pairing cancels farm-wide fluctuations from spectroscopic-bootstrap and INLA-hyperparameter variability. It isolates the land-use-class component of the difference. Each contrast posterior is summarised by  $\bar{\Delta}$ , its 95 % credible interval, and  $\Pr(\Delta > 0 \mid \mathbf{y})$  across the  $30 \times 500 = 15,000$  paired draws.



### 3 Results

#### 425 3.1 Spectroscopic predictions of soil-C variables

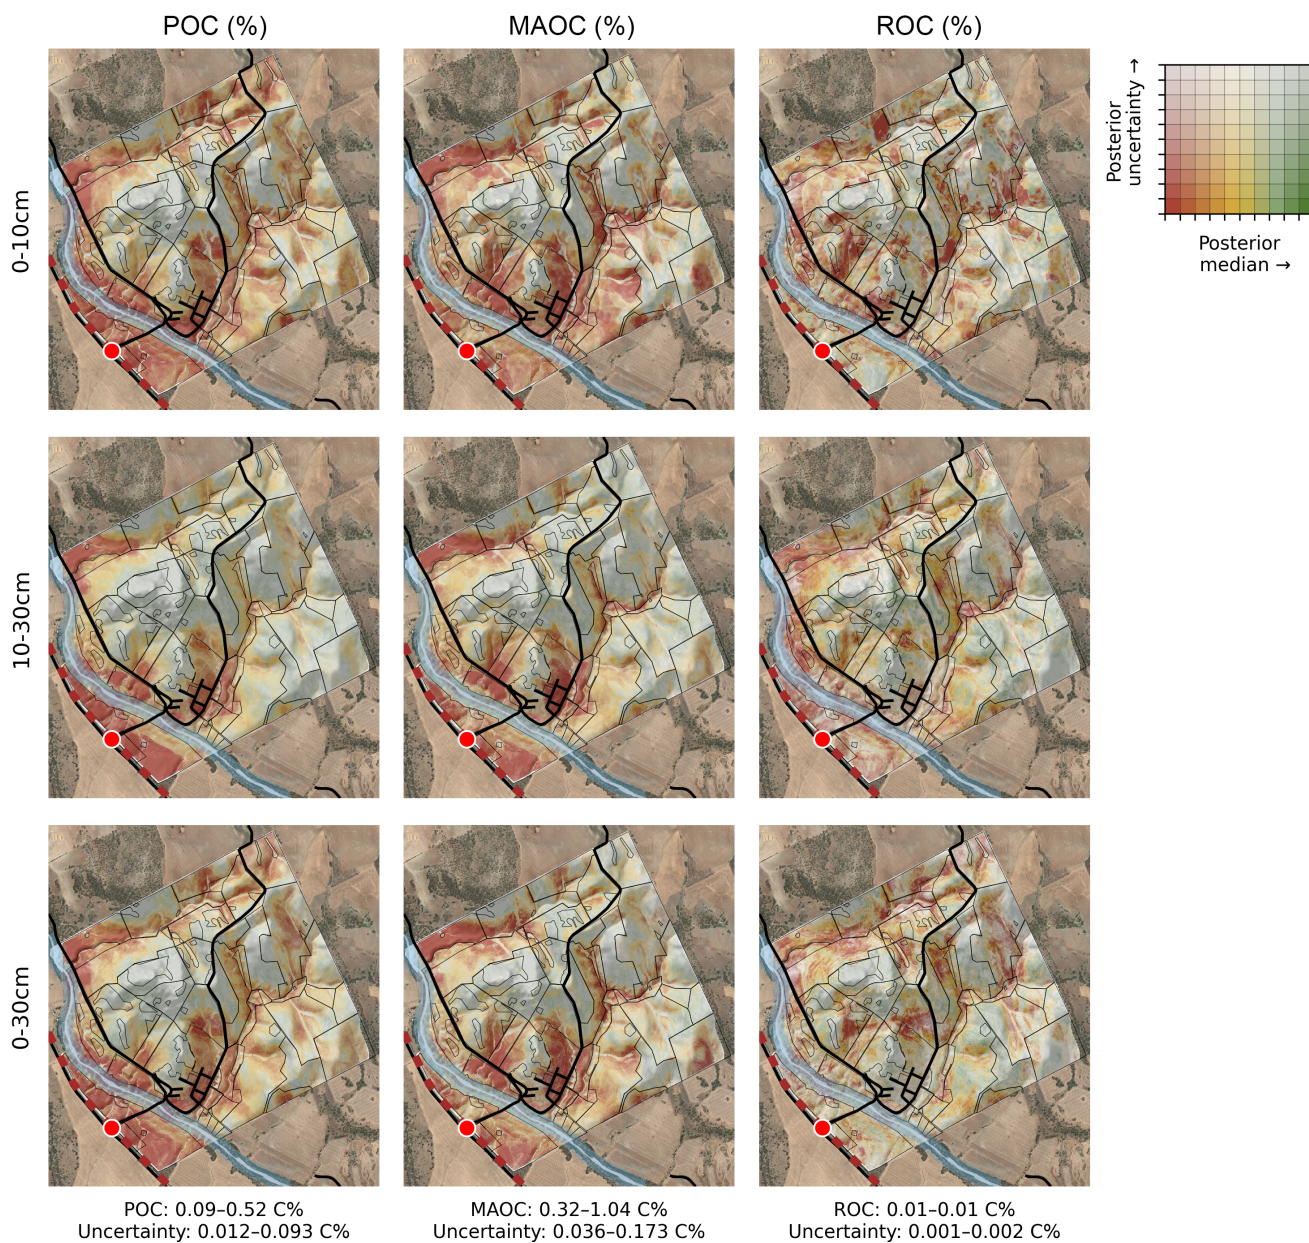
The spectroscopic models predicted the bulk total pools strongly and the C fractions less so but still usefully (Table 1). The  $B = 30$  bootstrap predictions per specimen form the discrete spectroscopic predictive distribution that enters the multiple-imputation loop as the uncertainty to propagate. The between-bootstrap standard deviation, averaged across specimens, was largest for MAOC (0.049 %C,  $\sim 7\%$  of its mean concentration), and tighter for the remaining  
430 responses (Table 1).

Table 1: **Spectroscopic cross-validation performance and bootstrap uncertainty.** Metrics are reported on the back-transformed (%C) scale from site-grouped 10-fold cross-validation ( $n = 100$  for organic-C sub-fractions;  $n = 200$  for TOC, TIC and TC). Mean SD is the average between-bootstrap standard deviation across the  $B = 30$  bootstrap predictions. The two POC sub-fractions were summed per bootstrap before the simplex-stage analysis.

Response	Model	$R^2$	RMSE ( %C)	$\rho_c$	KGE	Mean SD ( %C)
TOC	SVM	0.95	0.16	0.97	0.97	0.074
TC	Cubist	0.94	0.19	0.97	0.95	0.063
TIC	Cubist	0.84	0.03	0.92	0.90	0.010
ROC	SVM	0.88	0.16	0.94	0.93	0.001
POC <sub>250–50<math>\mu</math>m}</sub>	SVM	0.79	0.15	0.88	0.83	0.021
MAOC	SVM	0.64	0.13	0.80	0.79	0.049
POC <sub>2000–250<math>\mu</math>m}</sub>	SVM	0.60	0.17	0.75	0.68	0.010

#### 3.2 Posterior maps and calibrated uncertainty

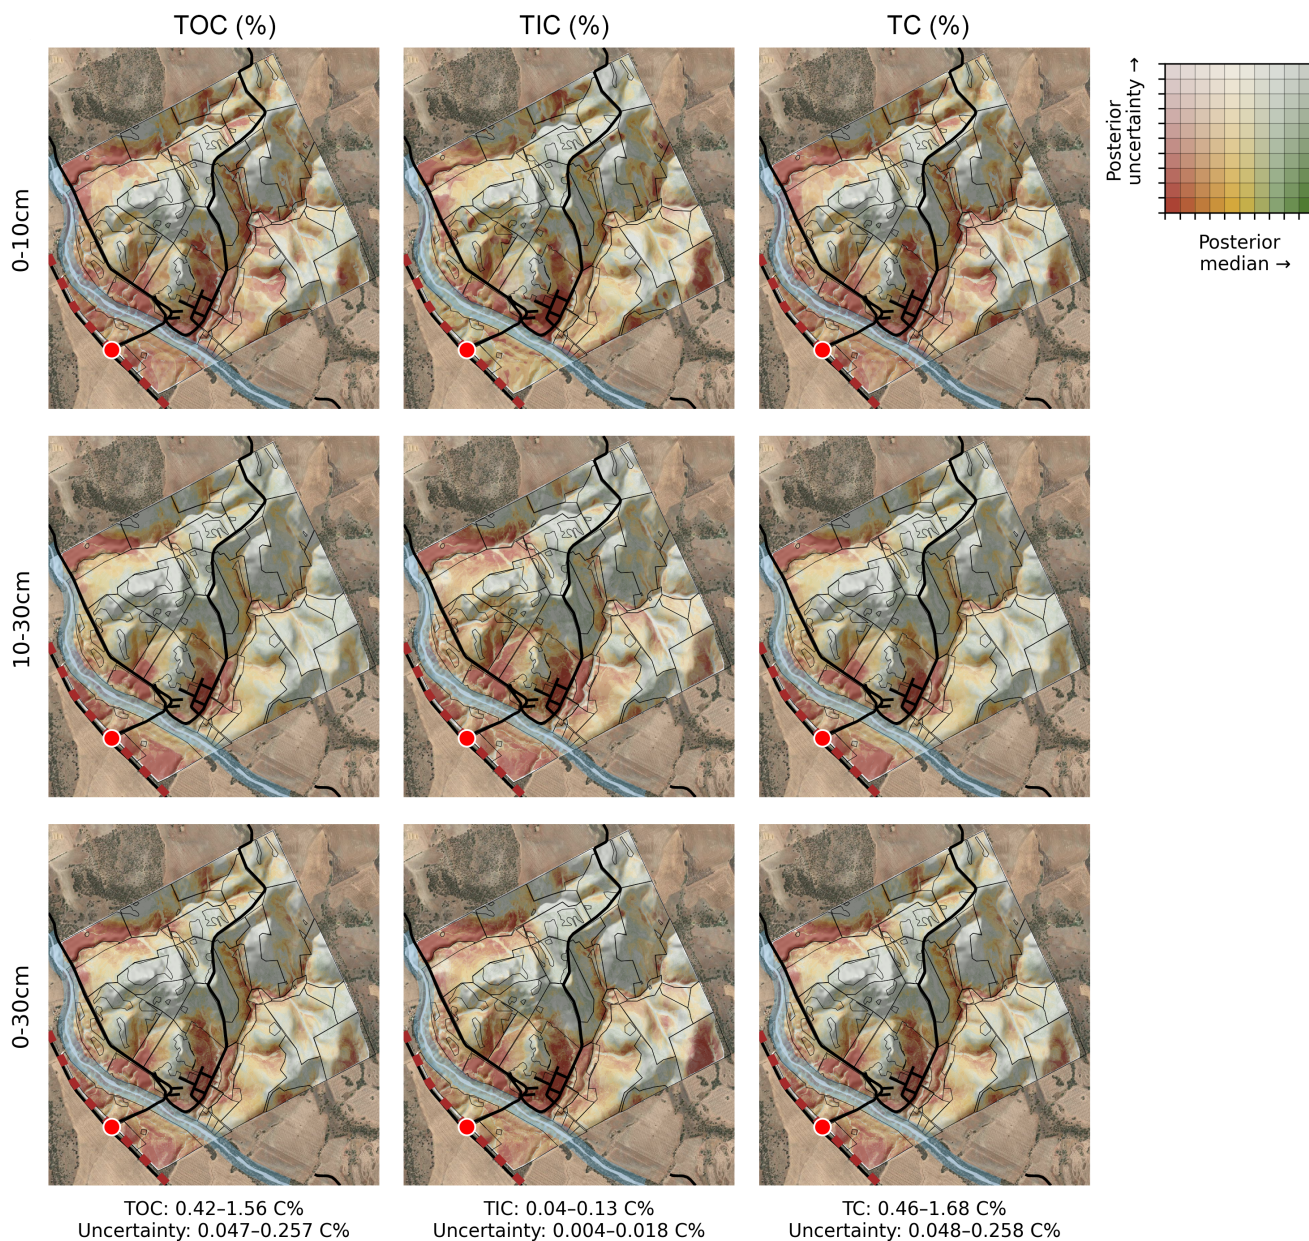
We show the posterior median and per-pixel uncertainty jointly in bivariate maps (Fig. 3, Fig. 4): each pixel's colour encodes both the predicted concentration and its posterior uncertainty. Posterior median concentrations range from 0.052–1.34 % C for POC, 0.20–1.88 % C for MAOC, and 0.005–0.018 % C for ROC at the three depths. MAOC is  
435 the largest of the three C fractions at every site and depth; ROC concentrations are an order of magnitude smaller than the other two fractions. The ROC magnitudes are in the range reported for the SoliTOC 400–600°C step on agricultural soils (Elementar Analysensysteme GmbH, 2023); despite the small absolute values, the ROC field shows coherent spatial structure across the farm (Fig. 3) that does not co-locate with the POC or MAOC patterns. All three fractions display visible spatial heterogeneity across the farm but ROC varies over a much narrower absolute  
440 range (Fig. 3).



**Figure 3. Bivariate posterior maps of organic carbon fractions.** Per-pixel posterior median and uncertainty for POC, MAOC, and ROC at 0–10, 10–30 and 0–30 cm. Colours jointly encode the median and posterior uncertainty, defined as the 95 % CI half-width / (2 × 1.96). Bins are response-specific quantiles over the 5–95 percentile range; outside values saturate at the corner bins.



For the totals, TOC ranges 0.26–3.11 % C, TIC 0.027–0.30 % C, and the derived TC = TOC + TIC ranges 0.29–3.31 % C. TC patterns track TOC closely at this site because TOC dominates TC by an order of magnitude (Fig. 4).



**Figure 4. Bivariate posterior maps of total carbon pools.** Per-pixel posterior median and uncertainty for TOC, TIC and TC at 0–10, 10–30 and 0–30 cm. TC is derived per posterior draw as  $TC^{(i)} = TOC^{(i)} + TIC^{(i)}$ . Colours and bins follow Fig. 3.



Both TOC and TC show their strongest spatial contrasts at 0–10 cm (Fig. 4). Lower-concentration areas lie on the southern paddocks below the river and on the lower-lying eastern part of the farm; higher-concentration patches occur on the steeper north-eastern parts of the farm and have larger uncertainty. The contrasts attenuate at 10–30 cm, and the 0–30 cm aggregate inherits the surface pattern (Fig. 4). TIC remains relatively low across the farm at every depth and is somewhat higher at 0–10 cm than at 10–30 cm, particularly in parts of the central and eastern paddocks (Fig. 4); we return to its possible origins in the Discussion.

Confident high-carbon predictions, the green corner of the bivariate, are essentially absent across all responses and depths (less than 0.1 % of pixels in the high-median, low-uncertainty quadrant). The model’s uncertainty grows systematically with predicted concentration, evident in the maps as a fading from confident-red at low values to uncertain-grey at high values without ever reaching confident-green (Fig. 3, Fig. 4). This coupling between concentration and uncertainty indicates sparser observational support at the upper tail of the distribution, where high-carbon areas are spatially localised and rare.

### 3.2.1 Uncertainty and its decomposition

The variance decomposition partitions the per-pixel predictive variance into five components: spectroscopic, trend, SPDE, trend–SPDE coupling, and residual (Table 2, Fig. 5).

Table 2: **Variance decomposition by response and depth.** Posterior variance of the transformed linear predictor is partitioned into spectroscopic, trend, SPDE, trend–SPDE coupling and residual components. Values are percentages of total predictive variance; rows sum to 100 %. Aggregation is the median across multiple-imputation iterations, then the mean across 368,666 grid pixels. For  $Z_1$ , the model is non-spatial, so SPDE and coupling shares are zero.

Depth (cm)	Response	Spectroscopic	Trend (T)	SPDE (S)	T–S coupling	Residual
0–10	$Z_1$	7.0	0.6	0.0	0.0	92.4
0–10	$Z_2$	28.1	1.9	10.3	-3.0	62.7
0–10	TOC	13.4	1.6	29.8	-0.8	56.0
0–10	TIC	24.9	121.2	128.5	-240.6	65.9
10–30	$Z_1$	10.1	0.6	0.0	0.0	89.3
10–30	$Z_2$	16.3	35.6	41.5	-70.2	76.8
10–30	TOC	17.1	4.5	26.0	-7.2	59.5
10–30	TIC	24.6	224.0	231.2	-445.5	65.8
0–30	$Z_1$	10.1	0.6	0.0	0.0	89.3
0–30	$Z_2$	19.5	7.8	13.3	-14.6	73.9
0–30	TOC	19.9	9.6	28.8	-17.2	58.9
0–30	TIC	22.1	143.0	151.8	-284.0	67.0

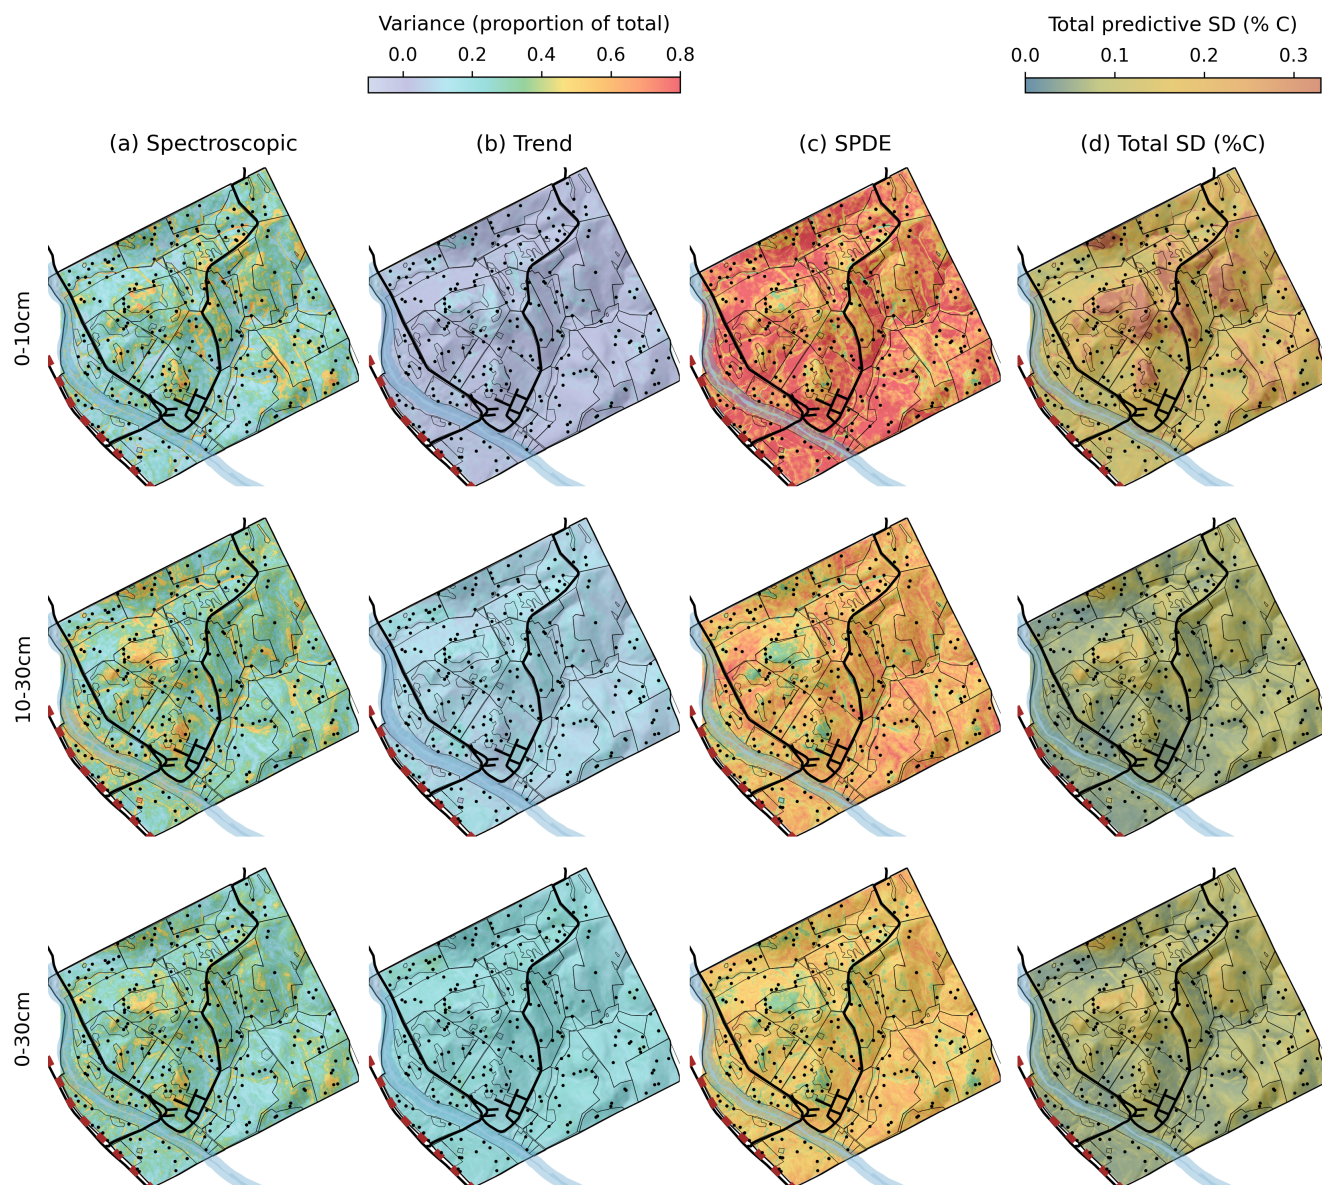


Depth (cm)	Response	Spectroscopic	Trend (T)	SPDE (S)	T-S coupling	Residual
------------	----------	---------------	-----------	----------	--------------	----------

The residual component carries the largest single positive share at every (response, depth) combination, dominating for  $Z_1$  (89–92 % across depths) and remaining substantial for the spatial responses (56–77 %). The spectroscopic part contributes a small but consistent share (7–28 %, mean  $\approx$  18 %). For  $Z_1$  the model is non-spatial, the SPDE and coupling parts are zero, the trend is also essentially zero ( $\approx$  0.6 % across depths), and the remaining variance partitions between spectroscopic and residual variances (Table 2).

The trend–SPDE coupling is negative for  $Z_2$ , TOC and TIC, with magnitudes largest for TIC at every depth (Table 2). The strongest case is TIC at 10–30 cm, where the trend and SPDE shares are both large and approximately equal (224.0 % and 231.2 % of the total predictive variance), and a negative coupling of  $-445.5$  % offsets them (so the row sums to 100 %). TIC at 0–30 cm ( $-284.0$  %) and TIC at 0–10 cm ( $-240.6$  %) follow, with  $Z_2$  at 10–30 cm ( $-70.2$  %) the next-most-extreme.

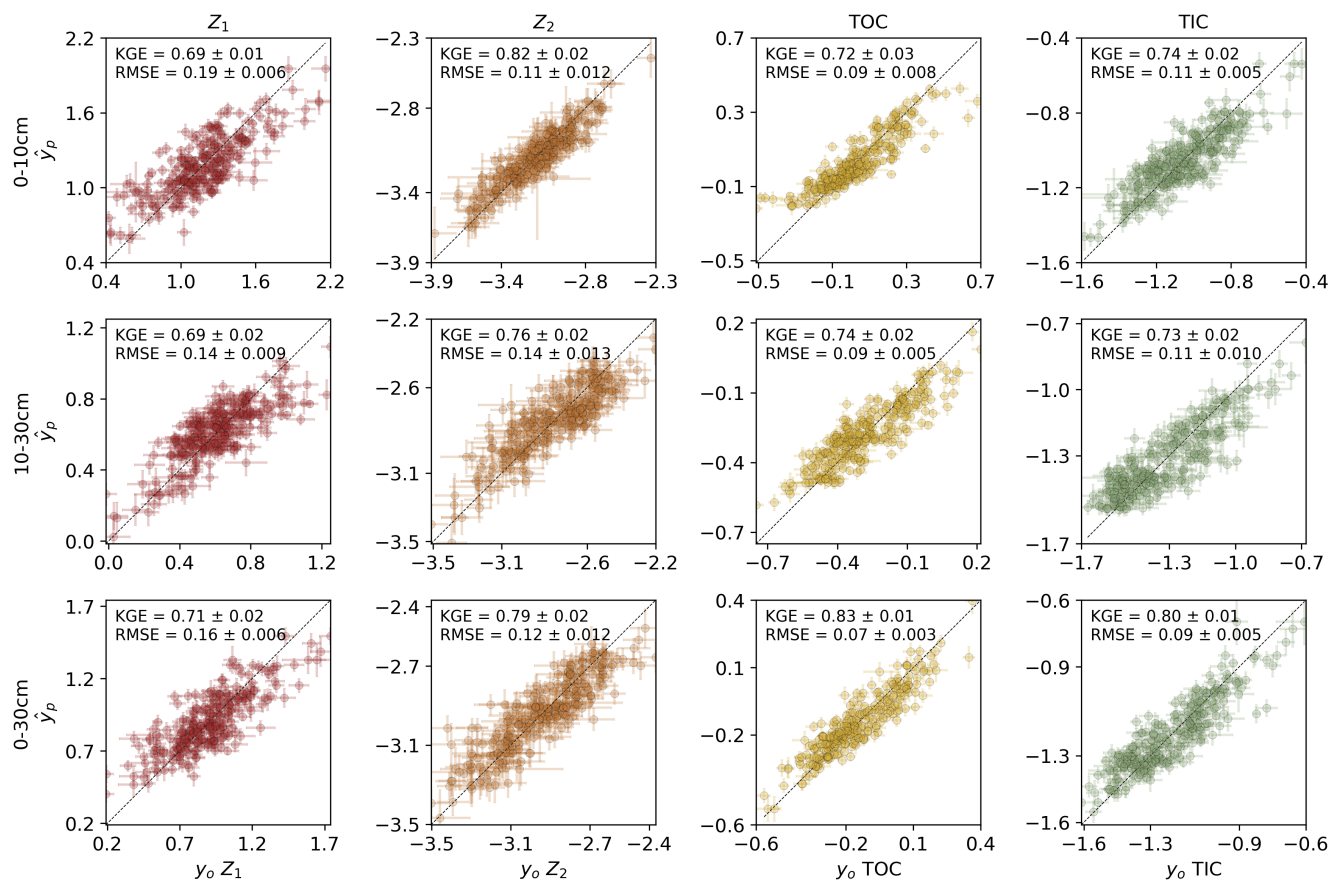
TOC at every depth and  $Z_2$  at 0–10 and 0–30 cm show modest couplings ( $|c| \leq 18$  %), an order of magnitude smaller than for TIC at any depth. The spatial decomposition of TOC predictive uncertainty (Fig. 5) illustrates how the per-pixel variance shares partition for one response: the SPDE component dominates at every depth, with the spectroscopic component intermediate and the trend component the smallest of the three, near-zero at 0–10 cm and rising modestly with depth.



**Figure 5. Spatial decomposition of TOC predictive uncertainty.** Per-pixel shares of latent-stage variance attributable to (a) spectroscopy, (b) NGBoost trend and (c) SPDE spatial field, shown for the three depth intervals. Panel (d) gives the pooled total predictive standard deviation (%C). Shares sum to one per pixel and use  $\text{Spec} + \text{Var}(T) + \text{Var}(S)$  as the denominator; coupling and residual variance are reported separately in Table 2.

### 3.2.2 Observed-vs-fitted performance and predictive coverage

The four modelled responses are recovered in transformed units across all three depths (Fig. 6).



**Figure 6. Observed-vs-fitted performance in transformed units.** Per-response posterior-mean prediction  $\hat{y}$  versus observed  $y$  in transformed units ( $Z_1, Z_2$  in ILR units; TOC and TIC in  $\log_{10}$  units; Eq. 1) at the three depths. Error bars indicate the per-point 95 % predictive interval. KGE and RMSE annotated as mean  $\pm$  SD across the  $M = 30$  multiple-imputation iterations. Dashed line: 1:1 reference.

475 KGE ranges 0.69–0.83 and RMSE 0.07–0.19 across the twelve response  $\times$  depth combinations, and is broadly consistent depth-to-depth for each response. TOC is the best-recovered response (KGE 0.72–0.83, RMSE 0.07–0.09), followed by  $Z_2$  (KGE 0.76–0.82, RMSE 0.11–0.14), TIC (KGE 0.73–0.80, RMSE 0.09–0.11) and  $Z_1$  (KGE 0.69–0.71, RMSE 0.14–0.19);  $Z_1$  shows a larger spread in values (Table 3).



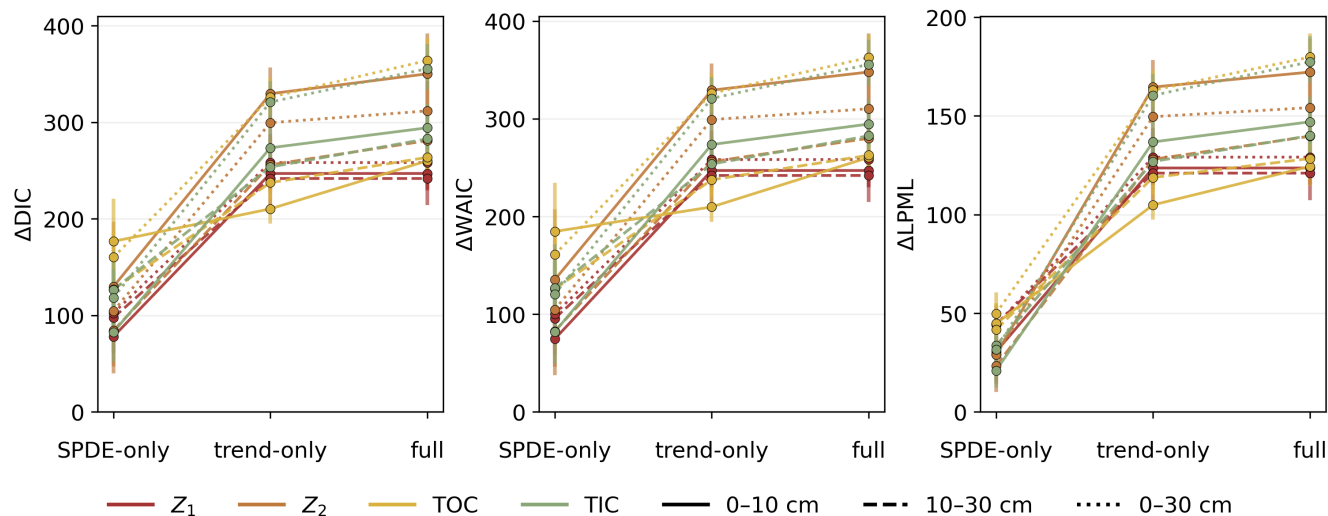
Table 3: **Pooled performance and predictive calibration per response and depth.** KGE and RMSE in transformed units ( $Z_1, Z_2$  in ILR units; TOC and TIC in  $\log_{10}$  units; Eq. 1); 95 % predictive coverage from the per-iteration PIT at the held-out CV residuals. All values are mean  $\pm$  SD across the  $M = 30$  multiple-imputation iterations.

Depth	Response	KGE	RMSE	95 % coverage
0–10 cm	$Z_1$	$0.69 \pm 0.01$	$0.19 \pm 0.01$	$0.94 \pm 0.01$
0–10 cm	$Z_2$	$0.82 \pm 0.02$	$0.11 \pm 0.01$	$0.94 \pm 0.01$
0–10 cm	TOC	$0.72 \pm 0.03$	$0.09 \pm 0.01$	$0.95 \pm 0.01$
0–10 cm	TIC	$0.74 \pm 0.02$	$0.11 \pm 0.01$	$0.95 \pm 0.01$
10–30 cm	$Z_1$	$0.69 \pm 0.02$	$0.14 \pm 0.01$	$0.95 \pm 0.01$
10–30 cm	$Z_2$	$0.76 \pm 0.02$	$0.14 \pm 0.01$	$0.95 \pm 0.01$
10–30 cm	TOC	$0.74 \pm 0.02$	$0.09 \pm 0.01$	$0.95 \pm 0.01$
10–30 cm	TIC	$0.73 \pm 0.02$	$0.11 \pm 0.01$	$0.95 \pm 0.01$
0–30 cm	$Z_1$	$0.71 \pm 0.02$	$0.16 \pm 0.01$	$0.94 \pm 0.01$
0–30 cm	$Z_2$	$0.79 \pm 0.02$	$0.12 \pm 0.01$	$0.95 \pm 0.01$
0–30 cm	TOC	$0.83 \pm 0.01$	$0.07 \pm 0.00$	$0.95 \pm 0.01$
0–30 cm	TIC	$0.80 \pm 0.01$	$0.09 \pm 0.01$	$0.95 \pm 0.01$

The 95 % predictive coverage from the per-iteration PIT distribution lies in 0.94–0.95 across all twelve response by depth combinations (Table 3), bracketing the nominal 0.95 within  $\pm 1$  percent in every case. Three dip marginally to 0.94 ( $Z_1$  at 0–10 and 0–30 cm,  $Z_2$  at 0–10 cm); the remaining nine sit on the nominal 0.95. The observed deviation of 0.01 is within one binomial SE of nominal,  $\sqrt{0.95 \times 0.05/250} \approx 0.014$  for  $n = 250$ .

### 3.2.3 Nested model comparison

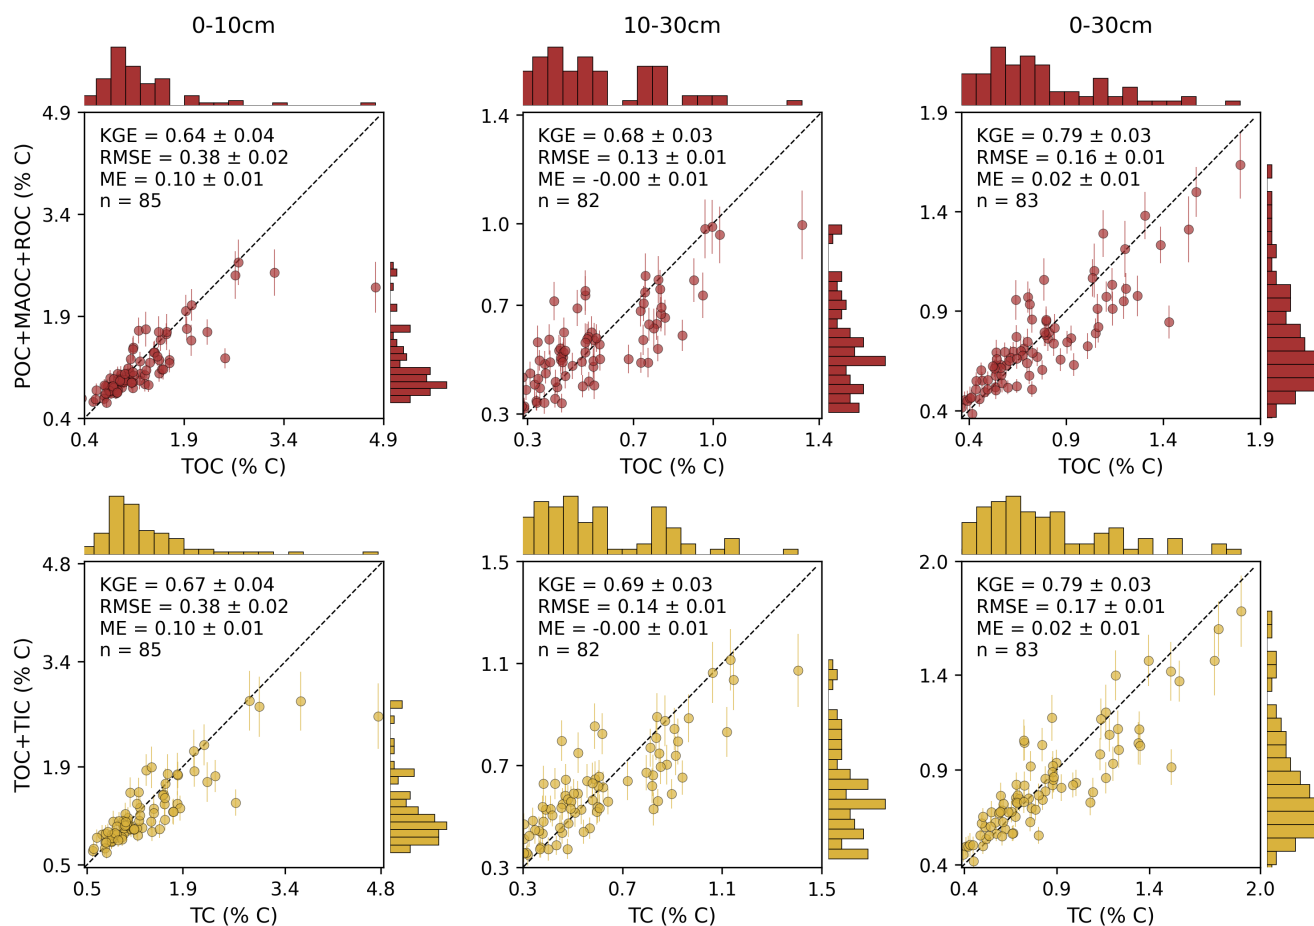
Across all twelve response and depth combinations, the full (trend + SPDE) variant achieves the largest  $\Delta$ DIC relative to the null model, with  $\Delta$ WAIC and  $\Delta$ LPML reproducing the same ordering (Fig. 7). For  $Z_1, Z_2$  and TIC the trend-only  $\Delta$ DIC exceeds the SPDE-only  $\Delta$ DIC by 150–200 units, indicating that the covariates carry the bulk of the predictive information at this site, with the SPDE adding a smaller further improvement on top of the trend (none for the non-spatial  $Z_1$ , where full and trend-only coincide). TOC is the exception as trend-only and SPDE-only are close ( $\Delta$ DIC gap of  $\sim 50$ ), and the full model adds  $\sim 100$  further units on top of trend-only, so for TOC the spatial field contributes comparably to, or more than, the covariates.



**Figure 7. Nested model comparison.**  $\Delta$ DIC,  $\Delta$ WAIC and  $\Delta$ LPML for SPDE-only, trend-only and full trend + SPDE models relative to the null model, by response and depth. Markers show posterior means; whiskers show 95 % multiple-imputation intervals across  $M = 30$  iterations. Positive values indicate improvement over the null. For  $Z_1$ , trend-only and full models coincide because the production model is non-spatial.

### 3.2.4 Coherence check against bulk laboratory measurements

The spatial-model predictions are checked for coherence across all parts of the modelling against the bulk laboratory measurements at sample sites (Fig. 8). First, compositional closure: the back-transformed sum POC + MAOC + ROC from the ILR spatial model agrees with the bulk-measured TOC with KGE =  $0.64 \pm 0.04$  and RMSE =  $0.38 \pm 0.02$  % C at 0–10 cm ( $n = 85$ ), KGE =  $0.68 \pm 0.03$  and RMSE =  $0.13 \pm 0.01$  % C at 10–30 cm ( $n = 82$ ), and KGE =  $0.79 \pm 0.03$  and RMSE =  $0.16 \pm 0.01$  % C at 0–30 cm ( $n = 83$ ); the mean error remains within  $|ME| \leq 0.10$  % C across depths. Second, mass-balance: the posterior-derived total  $TC^{(i)} = TOC^{(i)} + TIC^{(i)}$ , summed per draw from two independently-modelled SPDEs, agrees with the bulk-measured combustion TC with KGE =  $0.67 \pm 0.04$  and RMSE =  $0.38 \pm 0.02$  % C at 0–10 cm, KGE =  $0.69 \pm 0.03$  and RMSE =  $0.14 \pm 0.01$  % C at 10–30 cm, and KGE =  $0.79 \pm 0.03$  and RMSE =  $0.17 \pm 0.01$  % C at 0–30 cm, with  $|ME| \leq 0.10$  % C. All metrics are reported as the mean  $\pm$  SD across the  $M = 30$  multiple-imputation iterations.



**Figure 8. Coherence check against bulk analytical (SoliTOC) measurements.** (a) Back-transformed POC + MAOC + ROC from the ILR fraction model versus bulk-measured TOC. (b) Posterior-derived TC = TOC + TIC versus bulk-measured combustion TC. Columns show 0–10 cm, 10–30 cm and 0–30 cm ( $n = 83$ ). Panel metrics are mean  $\pm$  SD across  $M = 30$  multiple-imputation iterations.

### 3.2.5 Validation at monitoring sites

The propagated 95 % credible intervals are tested against an independent monitoring resampling at sample sites (Table 4, Fig. 9). The dominant managed-pasture cohort recovered the monitoring totals with KGE up to 0.84, mean errors within  $|ME| \leq 0.08$  % C, and the 95 % credible intervals bracketed 0.75–0.94 of the held-out measurements across the six responses and both depths. Natural vegetation sites, at the high-C end of the farm distribution, showed a negative ME of  $-0.27$  to  $-0.40$  % C on the organic-C totals at 0–10 cm; the credible intervals nonetheless captured 0.82 of MAOC observations in this class. Cropping sites showed near-zero or slightly positive ME on TOC and MAOC

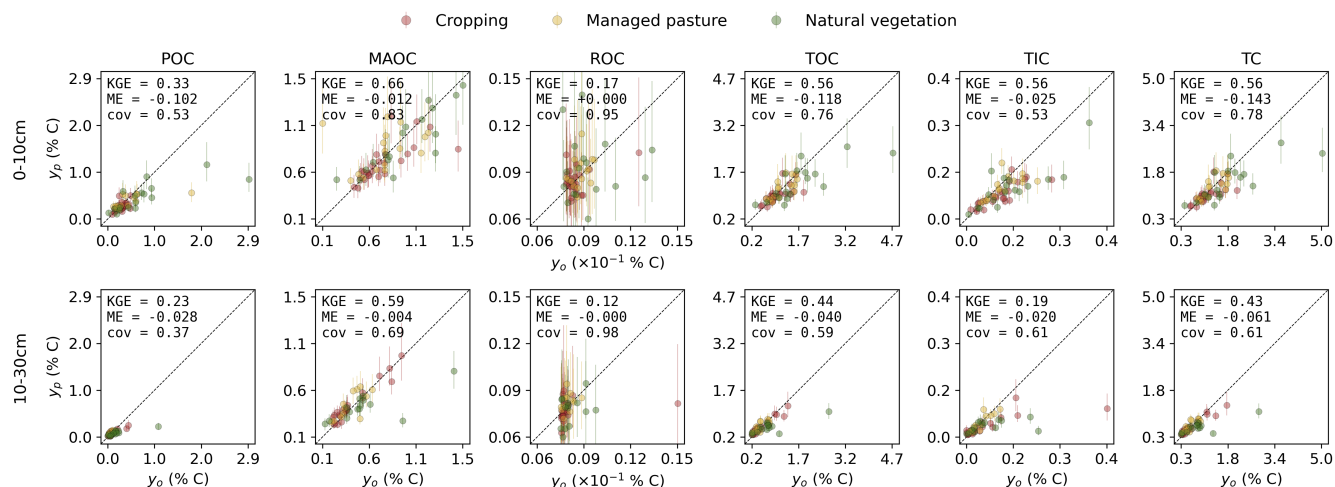


510 at 10–30 cm. TIC carried a small uniform offset of approximately  $-0.02$  to  $-0.03$  % C across every land-use class and depth.

Table 4: **Validation at monitoring sites.** Posterior-mean KGE, RMSE, ME and 95 % credible-interval coverage ( $\text{cov}_{95}$ ) against measurements from  $n = 59$  resampled sites. Land-use-specific MEs are reported for natural vegetation, managed pasture and cropping at the first sampling. ROC values are of order  $10^{-3}$  % C, so per-class ROC MEs appear as  $\approx 0$ .

Depth (cm)	Response	$n$	KGE	RMSE (%C)	ME (%C)	$\text{cov}_{95}$	$\text{ME}_{\text{nat}}$ (%C)	$\text{ME}_{\text{man}}$ (%C)	$\text{ME}_{\text{crop}}$ (%C)
0–10	POC	59	0.33	0.36	-0.102	0.53	-0.271	-0.070	-0.017
0–10	MAOC	59	0.66	0.22	-0.012	0.83	-0.039	0.066	-0.052
0–10	ROC	59	0.17	0.002	0.000	0.95	-0.000	0.000	-0.000
0–10	TOC	59	0.56	0.45	-0.118	0.76	-0.361	0.008	-0.049
0–10	TIC	59	0.56	0.05	-0.025	0.53	-0.036	-0.015	-0.026
0–10	TC	59	0.56	0.49	-0.143	0.78	-0.398	-0.007	-0.075
10–30	POC	59	0.23	0.13	-0.028	0.37	-0.102	0.025	-0.012
10–30	MAOC	59	0.59	0.16	-0.004	0.69	-0.109	0.043	0.028
10–30	ROC	59	0.12	0.001	-0.000	0.98	-0.000	0.000	-0.000
10–30	TOC	59	0.44	0.29	-0.040	0.59	-0.242	0.077	0.008
10–30	TIC	59	0.19	0.06	-0.020	0.61	-0.032	-0.000	-0.026
10–30	TC	59	0.43	0.33	-0.061	0.61	-0.274	0.077	-0.017

ROC was poorly predicted because its concentrations are near the analytical detection limit and vary over a much narrower range than the other carbon pools, so small absolute discrepancies translate into small KGE values despite being unbiased and with high credible-interval coverage (Table 4, Fig. 9).




**Figure 9. Validation at monitoring sites.** First-time posterior-mean prediction  $y_p$  against monitoring analytical (SolitoC) observation  $y_o$  at the 59 resampled sites, by response (columns) and depth (rows). Error bars are the per-site 95 % credible interval from the joint posterior; colour codes land-use class at the first time sampling. Dashed line: 1:1 reference.

### 3.3 Compositional patterns and ecological interpretation

#### 515 3.3.1 Organic-C composition and vulnerability

The compositional balance between the three C fractions varies spatially across the farm. They are not independently distributed. MAOC dominates the C composition at every site and depth, averaging 68 % of TOC at 0–10 cm and 76 % at 10–30 cm; POC contributes 31 % and 23 % respectively, and ROC remains  $\sim 1$  %. To bring out the spatial gradients the ternary RGB (Fig. 10) stretches each channel independently to its 1–99 percentile, emphasising where

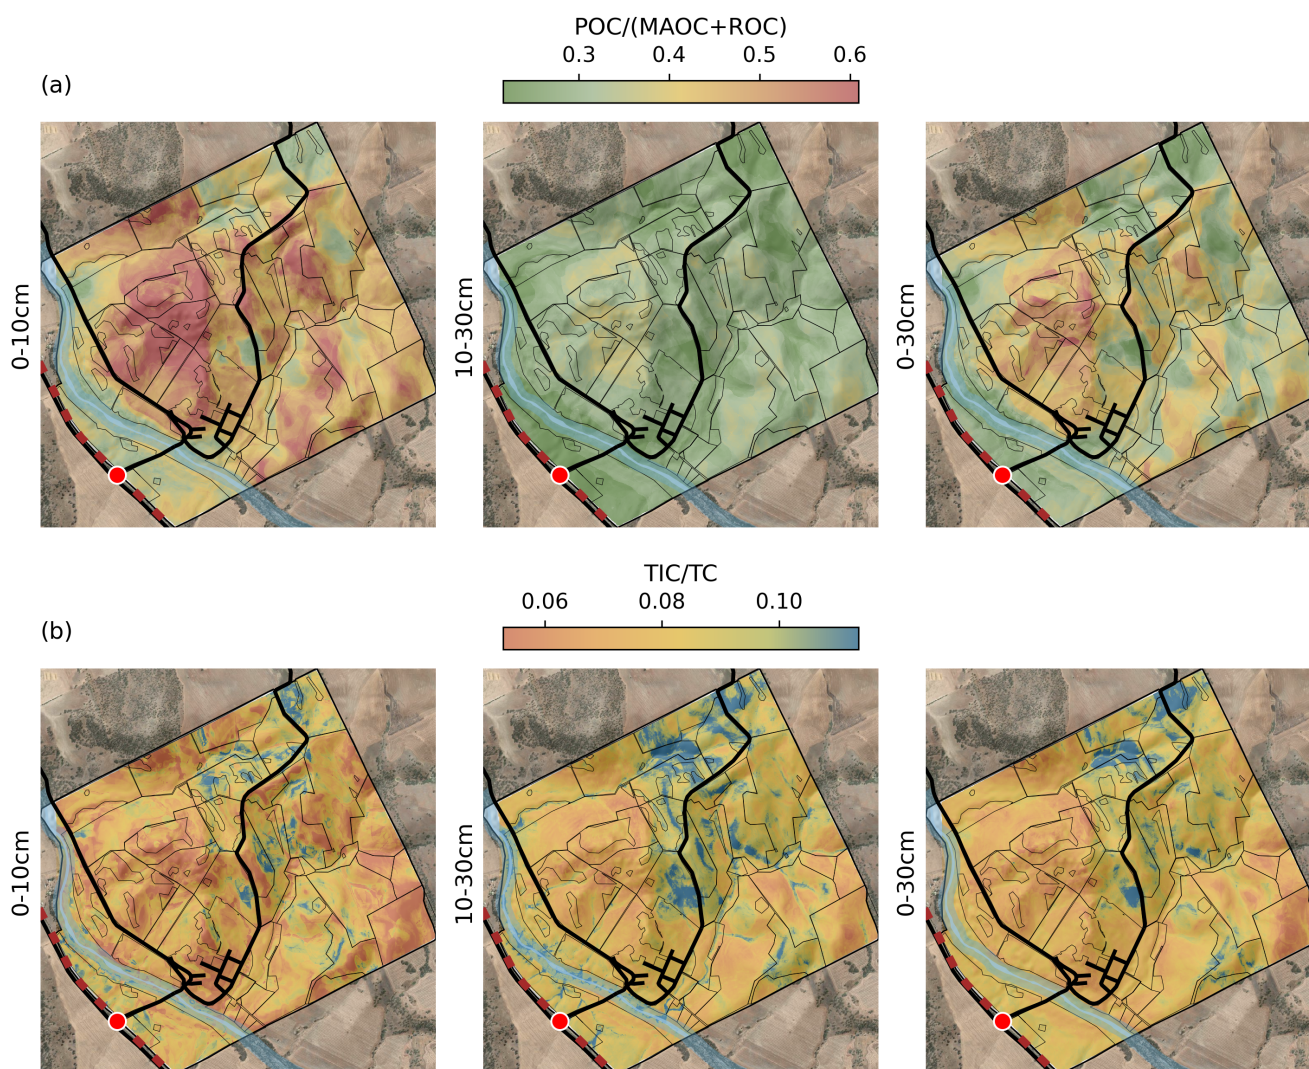
520 each fraction’s contribution rises or falls across the field rather than its absolute share. The gradients are sharpest at 0–10 cm, where field-shape boundaries and management blocks emerge as colour transitions; at 10–30 cm they flatten and the field reads as a more uniform composition (Fig. 10). The 0–30 cm depth integrates the two. The ternary map shows that the POC share is highest on the steeper parts of the farm, and the MAOC share is highest on the lower parts and below the river. ROC’s share is more spatially localised.



**Figure 10. Ternary RGB composition map of organic carbon fractions.** Per-pixel POC:MAOC:ROC proportions at 0–10, 10–30 and 0–30 cm, encoded as red, green and blue channels. Each channel gives the fraction’s share of TOC; the ternary inset maps colours to the three-part simplex.

525 The L:P ratio (Fig. 11 a) ranges from 0.23–0.84 in the surface soil layer and the map shows clear spatial heterogeneity, generally with larger values on steeper slopes and lower values in flatter areas. At 10–30 cm the ratio compresses to 0.13–0.40 and the map appears more uniform, as the ratio attenuates with depth. The 0–30 cm depth (0.18–0.70) integrates the two intervals and shows a similar spatial expression to the 0–10 cm depth (Fig. 11 a).

The inorganic share of total carbon, TIC/TC, is small but spatially structured across the farm (typically TIC/TC < 530 0.10; Fig. 11 b). At 0–10 cm discrete high-TIC/TC areas occur in the central and eastern, northeastern fields against a predominantly low-TIC/TC background where TOC dominates the C inventory; the same fields have the highest absolute TIC at this depth (Fig. 4 d). At 10–30 cm the high-TIC/TC areas are more distinct and the high–low contrast is stronger. The 0–30 cm depth integrates the two. The high-TIC/TC areas are not the lowest-TOC areas, so TIC/TC tracks local carbonate accumulation rather than organic-pool depletion.




**Figure 11. Maps of two compositional ratios.** Posterior-median maps of (a) labile-to-protected ratio  $POC/(MAOC+ROC)$  and (b) inorganic carbon share  $TIC/TC$  at 0–10, 10–30 and 0–30 cm. Ratios were computed draw-by-draw from the joint posterior before summarising per pixel.

### 535 3.3.2 Environmental controls on spatial patterns

Three groups of covariates contribute to the model (Fig. 12). The terrain ratio  $TWI/TCI$  is the leading predictor for TOC, TC and ROC at every depth, with within-response shares peaking at 26.2 % for TC at 0–30 cm and 23.3 % for TOC at 0–10 cm. Gamma-radiometric ratios dominate MAOC, K/U ratio at 0–10 cm (17.7 %), K/Th ratio at 0–30 cm (10.6 %) and the K-normalised (K/total counts, a weathering indicator) at 10–30 cm (10.2 %), and the



540 K/Th ratio is also the top predictor for POC at 0–10 and 0–30 cm (14.0 % in both). EMI conductivity contributes consistently across the organic-C pools and is the leading predictor for POC at 10–30 cm (DUALEM 50, 20.1 %). At 0–10 cm TWI/TCI is the leading predictor of TIC (11.5 %), but at 10–30 and 0–30 cm they are tangential-curvature difference and flow-accumulation/slope.

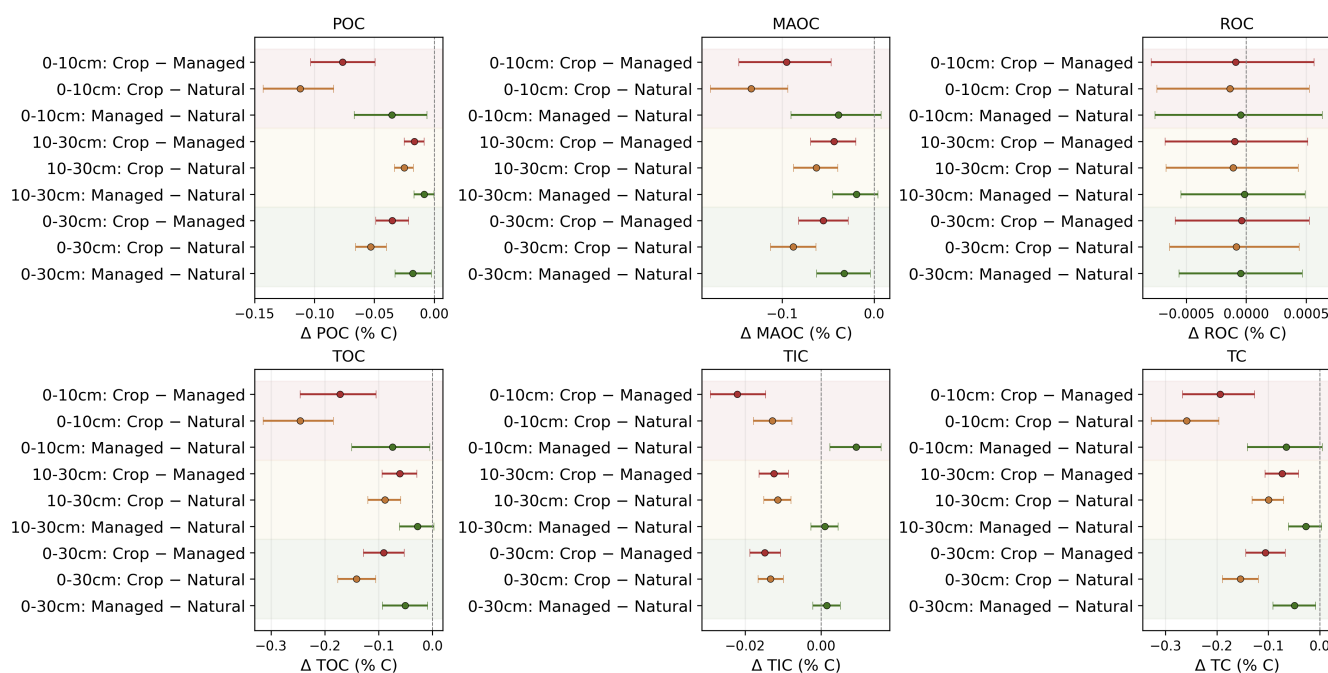


**Figure 12. NGBoost permutation feature importance by C fraction and depth.** Colour encodes the per-response normalised feature importance (% of within-response total). Circle size is inversely proportional to the coefficient of variation of importance across the  $M = 30$  multiple-imputation iterations, so larger discs indicate more stable rankings (low uncertainty) and smaller discs indicate less stable ones (high uncertainty).



### 3.3.3 Land-use differences in soil-C fractions and totals

545 The three land-use classes on the farm, cropping, managed pasture and natural vegetation, show systematic differences in the soil C fractions and totals (Fig. 13). Cropping carries lower organic-C concentrations than either pasture class for every C pool and total at every depth, with the largest deficits at 0–10 cm. At 0–10 cm, TOC under cropping is  $\bar{\Delta} = -0.25$  % C (95 % CI: -0.31, -0.18) lower than under natural vegetation and  $-0.17$  % C (95 % CI: -0.25, -0.10) lower than under managed pasture; managed pasture is itself lower than natural vegetation by  $-0.07$  % C (95 % CI: -0.15, -0.005). The ordering cropping < managed pasture < natural vegetation holds at 10–30 cm and 0–30 cm with magnitudes about half those at 0–10 cm. POC and MAOC follow the same pattern at 0–10 cm: cropping is depleted by  $\bar{\Delta} = -0.11$  % C and  $-0.13$  % C respectively relative to natural vegetation. For every cropping–pasture contrast in POC, MAOC, TOC and TC at 0–10 cm, the 95 % credible interval excludes zero and  $\text{Pr}(\Delta < 0) \approx 1$ .



**Figure 13. Pairwise posterior contrasts among land-use classes.** Differences in spatial-mean carbon concentration between cropping, managed pasture and natural vegetation for each response and depth interval. Each contrast is paired within the same posterior draw; markers show posterior means across 15,000 draws and whiskers show empirical 95 % credible intervals. The dashed line marks zero difference. ROC is shown on a narrower x-axis because its concentrations are much smaller.

555 ROC is unresponsive to land use: all nine ROC contrasts have credible intervals that bracket zero, with magnitudes of order  $10^{-4}$  %C. TIC contrasts are also small in magnitude ( $|\bar{\Delta}| \leq 0.03$  % C). At 0–10 cm cropping carries marginally less TIC than managed pasture and natural vegetation, and managed pasture carries marginally more



TIC than natural vegetation ( $\bar{\Delta} = 0.009$  % C,  $\Pr(\Delta > 0) = 0.995$ ); these directional differences are small relative to the organic-C-pool contrasts and consistent with a parent-material rather than a management origin.

Land use is also reflected in the compositional balance (Fig. 11 a). At 0–10 cm the L:P ratio is largest under managed  
560 (median 0.48) and natural vegetation (0.48) and smallest under cropping (0.45); 41 % of pasture pixels have a ratio  $>$   
0.5 versus 28 % of cropping pixels. The pattern attenuates with depth. At 10–30 cm all classes converge to a median  
L:P ratio  $\approx$  0.30 with no areas above 0.5. The 0–30 cm depth integrates the two: median 0.35 under cropping versus  
0.38–0.39 under managed pasture and natural vegetation, with 3.5–12 % of pixels above 0.5 (Fig. 11 a).

#### 4 Discussion

565 We demonstrate a probabilistic compositional framework with joint uncertainty propagation across the spectroscopic,  
trend and spatial parts, applied at a Mediterranean-type semi-arid agroecosystem to deliver the joint soil-carbon  
inventory, POC, MAOC, ROC, TOC, TIC and TC, at three depths under a single calibrated posterior. Mapping  
every form of soil C apart from gaseous CO<sub>2</sub> in a single posterior, rather than mapping TOC, the fractions or  
TIC in isolation, makes the inventory testable: the mass balance  $\text{TOC} + \text{TIC} = \text{TC}$  closes within the propagated  
570 intervals, and the farm-aggregated TC posterior is the level at which whole-farm carbon decisions can be made. The  
propagated 95 % credible intervals are calibrated against held-out residuals, coherent with the analytical laboratory  
measurements at sample sites, and validated at monitoring sites. The calibrated posterior shows a graded land-use  
signal in the management-responsive C pools, while also highlighting distinct process signatures across the C fractions  
that are obscured when considering TOC alone. The variance decomposition further shows a genuine but limited  
575 trade-off between the trend and SPDE parts.

The propagated posterior is calibrated against held-out data. The 95 % PIT coverage falls within  $\pm 1$  percentage point of  
nominal across all twelve response  $\times$  depth combinations (Table 3), and the predictive intervals bracket bulk analytical  
(SoliTOC) measurements with KGE 0.64–0.79 for both compositional closure ( $\text{POC} + \text{MAOC} + \text{ROC} = \text{TOC}$ ) and  
TC mass balance ( $\text{TOC} + \text{TIC} = \text{TC}$ , summed per draw from two independently-modelled SPDEs; Fig. 8). Where  
580 prior work has propagated spectroscopic uncertainty into spatial models stage-by-stage (Ramírez-López et al., 2019,  
Viscarra Rossel et al. (2016), Viscarra Rossel et al. (2014)) or coupled trend models to spatial residuals with predictive  
intervals (Padarian et al., 2019; Poggio et al., 2016; Hu et al., 2025), here each multiple-imputation iteration draws  
from the spectroscopic bootstrap, refits NGBoost on that draw, threads the per-point predictive standard deviation  
into INLA as an EIV scale on the observation precision (Eq. 3), and Rubin’s-rules pooling closes the loop into a  
585 single posterior of 15,000 back-transformed draws per pixel.

The coherence KGEs strengthen with depth: 0.64 at 0–10 cm, 0.68 at 10–30 cm, 0.79 at 0–30 cm, and the same  
depth-graded pattern emerges in the TC mass balance. This could be interpreted as worse performance at the surface,  
but if read as a property of the soil, it is informative. The 0–10 cm layer carries the most transient variability from



biological turnover and recent residue inputs. The labile particulate fraction with turnover times of months to years  
590 (Lavallee et al., 2020; Cotrufo et al., 2013), and the propagated posterior translates that variability into wider 0–10  
cm credible intervals. The 0–30 cm aggregate integrates the surface fluctuations against the more stable subsoil signal,  
producing both better point agreement and narrower intervals.

The propagation is also tested against the monitoring measurements at sample sites, approximately 18 months later  
(Section 3.2.5, Fig. 9, Table 4). The managed-pasture cohort was recovered with KGE up to 0.84,  $|ME| \leq 0.08$  % C  
595 and 0.75–0.94 credible-interval coverage across the six responses. Natural vegetation sites carried a larger negative  
ME on the organic-C totals at the surface (–0.27 to –0.40 % C), consistent with the shrinkage of posterior means at  
the upper tail of the predictive distribution that the bivariate maps display (Fig. 3, Fig. 4); the calibrated intervals  
nonetheless bracketed 0.82 of MAOC observations in that class. A small TIC offset of about –0.02 to –0.03 %  
C across every land-use class might reflect a systematic shift on the SoliTOC 900 °C window in the monitoring  
600 measurements.

The joint posterior also reveals a trade-off between the trend ( $T$ ) and SPDE ( $S$ ) parts. At sample locations the  
data constrain only the sum  $T + S$ , not the two components individually. The posterior therefore assigns strongly  
anti-correlated draws to  $T$  and  $S$ , so  $\text{Var}(T)$  and  $\text{Var}(S)$  can each be large even where their sum varies little across  
the posterior. The constraint is sharpest for TIC at 10–30 cm (Table 2), where  $\text{Var}(T)$  and  $\text{Var}(S)$  each reach  $\sim 230$   
605 % of the total predictive variance and the trend–SPDE coupling  $2\text{Cov}(T, S)$  is –445.5 %, which cancels the overlap.  
The negative coupling is not an error; it identifies where the trend–SPDE split is poorly resolved and where any  
inference should rest on the joint trend + SPDE ( $T + S$ ) posterior rather than on either part alone. For TOC the  
constraint is mild at every depth (couplings within  $|c| \leq 18$  %); the SPDE share dominates, neither marginal share  
exceeds 100 % of the total predictive variance, and the trend–SPDE split is identifiable (Fig. 5).

610 The data support the hypothesis. The dominant spatial control on each C fraction relates to the mechanism by  
which that fraction is stabilised. MAOC, the only pool literally bound to the mineral matrix, is dominated by  
gamma-radiometric K-based ratios (K/Th, K/U, K-normalised; Fig. 12), proxies for mineralogy. POC at 10–30 cm  
tracks EMI conductivity, which represents texture and moisture that mediate physical protection and microbial  
turnover of the labile pool. The pools that integrate multiple stabilisation mechanisms (TOC, TC, TIC) are dominated  
615 by topographic redistribution, principally the TWI/TCI ratio, at this farm scale where parent-material variation is  
constrained. The expected gamma-radiometric signal on TIC, anticipated from continental-scale carbonate mapping  
by Wilford et al. (2015), did not emerge. TIC is instead affected by the same terrain features that drive the organic-C  
pools.

The compositional balance between POC and the stable pools (MAOC+ROC), captured by the  $Z_1$  ILR coordinate,  
620 is environmentally invariant: the trend's contribution to  $\text{Var}(Z_1)$  is essentially nil at every depth (Table 2), and  
the covariates do not move the POC-vs-(MAOC+ROC) balance across the farm. What does move it is land use.



Cropping has both reduced absolute carbon and shifted the labile-to-protected balance toward the protected end (median L:P 0.45 versus 0.48 under pasture; Fig. 11 a), a categorical management effect that the continuous-covariate trend cannot capture. Continuous gradients leave the composition alone; disturbance, a categorical, biological process, moves it.

These distinctions matter for whole-farm carbon accounting and management. A TOC-only map captures only the organic pool, leaving TIC unmapped; yet decisions about a farm's carbon inventory depend on the total C across organic and inorganic forms and on how the organic pool is distributed among labile and stable fractions. Fraction-level mapping with calibrated propagated uncertainty quantifies that inventory, TOC with its compositional split (POC, MAOC, ROC) and TIC, each with its own covariates, and the propagated credible intervals make the contrasts (between fractions, depths, and land uses) testable. Holistic farm management at the field or farm scale needs both the absolute content and its compositional distribution; fraction-level mapping provides both.

Cropping soil has less organic C than managed pasture, which in turn carries less than soil under natural vegetation. The largest gap is at the surface and attenuates with depth (Fig. 13). This points to a graded response to management intensity rather than a binary cropping-versus-pasture split. The managed pasture is itself resolvable below natural vegetation at the surface; the gradient is not a switch from disturbed to undisturbed but reflects how intensively the system has been worked. The depth attenuation is what we would expect when the difference is driven by surface biology and recent residue inputs rather than by mineralogy or hydrological position.

The depletion is asymmetric across the C pools. The POC and MAOC pools reflect the land-use contrast at the surface (Fig. 13), while ROC does not move detectably across LU at any depth and TIC contrasts are small in magnitude ( $|\bar{\Delta}| \leq 0.03 \% C$ ). The asymmetry is what the operational-fraction framework anticipates for the management-responsive pools. POC and MAOC carry the surface land-use signal because their stabilisation depends on biology and mineral binding, both affected by management; TIC is a parent-material signal of the carbonate mineralogy. The inorganic share of TC supports the parent-material attribution: the high-TIC/TC patches (Fig. 11 b) sit on the same central and eastern paddocks where TIC content peaks (Fig. 4 d), so the spatial structure of TIC/TC ties to where carbonate has accumulated, not to where TOC has been depleted. At 0–10 cm TIC nonetheless carries a small but high-probability managed > natural difference ( $\bar{\Delta} = 0.009 \% C$ ,  $\Pr(\Delta > 0) = 0.995$ ), plausibly reflecting historical liming or a parent-material confound across the landscape positions. ROC's unresponsiveness presents two non-exclusive interpretations: a pyrogenic one, in which the 400–600°C combustion window captures charcoal-derived carbon (Gustafsson et al., 1997; Schmidt and Noack, 2000) that turns over on timescales much longer than typical farm management; or an analyser-defined mixture in which the same window also releases thermally-stable non-pyrogenic organic carbon whose responses do not co-vary with land use. The combustion data alone cannot distinguish them. The land-use effect concentrates where the biology is, in the management-responsive labile and mineral-associated pools, and remains negligible in the pools where biology has less control.



655 Cropping has changed both the amount and the composition of organic C. L:P is lowest under cropping at the surface (median 0.45 versus 0.48 under managed pasture and natural vegetation; Fig. 11 a), and the fraction of high-L:P pixels drops from 41 % under pasture to 28 % under cropping. Cropping has selectively depleted its labile particulate pool, leaving the protected end relatively intact. The ratio has two interpretations. From a stability perspective, a larger labile share is more vulnerable to disturbance-driven loss (Lavalley et al., 2020; Viscarra Rossel et al., 2019); from an agroecological perspective the same fraction is the substrate for microbial activity, nutrient cycling and aggregate stabilisation (Janzen, 2006; Cotrufo et al., 2013), so high-L:P regions also mark areas of greater biological activity. The two perspectives converge at this site: cropping has lost labile carbon (the stability concern) and biological substrate (the agroecological concern), so its composition is more degraded under either framing.

These contrasts are spatial, not temporal; they are differences in the C state across LU classes today, not a causal estimate of management effects on rates of change, and the depletion under cropping could in principle be partly 665 confounded with baseline soil differences across the landscape positions occupied by each LU class. Repeat sampling at the same sites, or paired-paddock and chronosequence designs, would convert the spatial differences into temporal rates. The depth-attenuating, pool-selective pattern is nonetheless internally consistent with biology being the driver, and the calibrated propagated uncertainty places the LU effect outside sampling noise.

#### 670 4.1 Considerations and future research

Several considerations bound the inferences here. First, the farm is small ( $\sim 4 \times 4$  km) relative to the posterior SPDE range for TIC ( $\approx 5$ – $10$  km) and  $Z_2$  at depth ( $\approx 2.4$ – $2.8$  km), so spatially-buffered  $k$ -fold cross-validation (Roberts et al., 2017) could not be applied uniformly without leaving too few training points per fold. Calibration was instead assessed through PIT coverage on held-out CV residuals (Table 3), bulk SoliTOC measurements at sample sites 675 (Fig. 8), and validation at monitoring sites (Fig. 9, Table 4).

Second, the trend and SPDE components are weakly separable for TIC at 10–30 cm, where their coupling reaches  $-445.5$  % of total predictive variance. The individual trend ( $T$ ) and SPDE ( $S$ ) contributions should therefore not be interpreted separately for this case; inference should be based on the calibrated joint trend + SPDE posterior ( $T + S$ ).

Third, the interpretation of ROC remains open. The SoliTOC 400–600°C combustion window plausibly captures 680 pyrogenic carbon in this fire-affected landscape, but it is not particle-selective. Testing whether ROC spatial structure tracks fire history would require independent markers, such as  $^{13}\text{C}$ -NMR aromaticity, benzene polycarboxylic-acid molecular markers, or petrographic identification of char particles.

The framework can be applied to regional-, farm- and field-scale digital soil mapping where dense covariate stacks and spectropedometric data are available. Its main outputs are per-pixel and aggregated credible intervals for soil-C 685 fractions and totals, supporting comparisons among fractions, depths and carbon inventories.



## 5 Conclusions

We developed a probabilistic compositional framework for farm-scale mapping of soil-C fractions and totals with joint uncertainty propagation. The framework combines spectroscopic prediction, probabilistic trend modelling and Bayesian spatial modelling, while preserving both the closure of the organic fractions (POC, MAOC and ROC) and the mass balance between TOC, TIC and TC. The propagated posterior was well calibrated. Predictive 95 % coverage was 0.94–0.95 across all response–depth combinations, and the predictions were supported by three checks: held-out CV residuals, bulk SoliTOC measurements at sample sites (KGE 0.64–0.79), and independent monitoring measurements (KGE 0.12 for ROC to 0.66 for MAOC, and up to 0.84 in the managed-pasture cohort). The maps show that land use affects both the amount and composition of organic carbon. Organic-C pools followed the ordering cropping < managed pasture < natural vegetation at all depths, with the largest deficits in the surface layer. Cropping also shifted the organic pool toward a more protected composition, with a lower labile-to-protected ratio than pasture. In contrast, ROC showed little land-use response and TIC contrasts were small ( $|\bar{\Delta}| \leq 0.03 \text{ \% C}$ ). The spatial controls differed among C pools. MAOC was most strongly associated with gamma-radiometric K-based ratios, consistent with mineralogical control. POC at 10–30 cm was most strongly associated with EMI conductivity, consistent with texture and moisture controls on physical protection and turnover. TOC, TC, ROC and TIC were mainly organised by topographic covariates, indicating that redistribution processes dominate the spatial pattern of pools not directly tied to the mineral matrix at this field scale. The variance decomposition also shows where interpretation should be cautious. For TIC at 10–30 cm, the trend and SPDE components are weakly separable, so their individual variance shares should not be interpreted. The joint trend + SPDE posterior remains calibrated and is the appropriate basis for inference. Overall, the calibrated posterior, rather than a point prediction, provides the useful unit of evidence for soil-C monitoring, accounting and management decisions.

## 6 Code and data availability

The code and data supporting the findings of this study can be made available from the corresponding author upon reasonable request.

## 7 Author contributions

R.A.V.R. secured funding, conceived the research, performed the analysis and wrote the manuscript. L.W. and F.S. contributed to field work, laboratory analyses and manuscript review and editing.

## 8 Competing interests

R.A.V.R. is an Executive Editor of *SOIL*. The authors declare no competing interests.



## 715 9 Acknowledgements

We thank the Australian Government's Australia-China Science and Research Fund-Joint Research Centres (ACSRF-JRCs; grant ACSRIV000077) for funding. We also thank members of the Soil and Landscape Science Laboratory, Dr Andrea Bravo-Escobar, Dr Mingxi Zhang, Dr Yang (Teresa) Hu, Mr Nathan Wells, Ms Sadichhya Adhikari, Mr Ashton Reaid, Mr Tshering Tashi and Mr Jigme Tenzin, for their contributions to the field work. We thank  
720 the Muresk Institute and its farm staff for their support and access to the site and we acknowledge the traditional custodians of the land on which this research was conducted, the Noongar people, and pay our respects to their elders past and present.



## References

- 725 Aitchison, J.: The Statistical Analysis of Compositional Data, Blackburn Press, Caldwell, NJ, reprint of the 1986 Chapman & Hall edition, 2003.
- Arshad, M., Li, N., Della Bella, L., and Triantafyllis, J.: Field-scale digital soil mapping of clay: Combining different proximal sensed data and comparing various statistical models, *Soil Science Society of America Journal*, 84, 314–330, <https://doi.org/10.1002/saj2.20008>, 2020.
- 730 Behrens, T., Schmidt, K., Zhu, A. X., and Scholten, T.: The ConMap approach for terrain-based digital soil mapping, *European Journal of Soil Science*, 61, 133–143, 2010.
- Behrens, T., Schmidt, K., MacMillan, R. A., and Viscarra Rossel, R. A.: Multiscale contextual spatial modelling with the Gaussian scale space, *Geoderma*, 310, 128–137, 2018.
- Behrens, T., Viscarra Rossel, R. A., Kerry, R., MacMillan, R., Schmidt, K., Lee, J., Scholten, T., and Zhu, A.-X.: The relevant range of scales for multi-scale contextual spatial modelling, *Scientific Reports*, 9, 14 800, 2019.
- 735 Cook, R. D.: Detection of Influential Observation in Linear Regression, *Technometrics*, 19, 15–18, <https://doi.org/10.1080/00401706.1977.10489493>, 1977.
- Cortes, C. and Vapnik, V.: Support-vector networks, *Machine Learning*, 20, 273–297, <https://doi.org/10.1007/BF00994018>, 1995.
- 740 Cotrufo, M. F., Wallenstein, M. D., Boot, C. M., Deneff, K., and Paul, E.: The Microbial Efficiency-Matrix Stabilization (MEMS) framework integrates plant litter decomposition with soil organic matter stabilization: do labile plant inputs form stable soil organic matter?, *Global Change Biology*, 19, 988–995, <https://doi.org/10.1111/gcb.12113>, 2013.
- Daubechies, I.: Orthonormal bases of compactly supported wavelets, *Communications on Pure and Applied Mathematics*, 41, 909–996, <https://doi.org/10.1002/cpa.3160410705>, 1988.
- 745 DIN: DIN 19539: Investigation of solids — Temperature-dependent differentiation of total carbon (TOC<sub>400</sub>, ROC, TIC<sub>900</sub>), 2016.
- Duan, T., Avati, A., Ding, D. Y., Thai, K. K., Basu, S., Ng, A. Y., and Schuler, A.: NGBoost: Natural Gradient Boosting for Probabilistic Prediction, in: *Proceedings of the 37th International Conference on Machine Learning (ICML)*, vol. 119 of *Proceedings of Machine Learning Research*, pp. 2690–2700, <https://proceedings.mlr.press/v119/duan20a.html>, 2020.
- 750 Egozcue, J. J. and Pawlowsky-Glahn, V.: Groups of Parts and Their Balances in Compositional Data Analysis, *Mathematical Geology*, 37, 795–828, <https://doi.org/10.1007/s11004-005-7381-9>, 2005.
- Elementar Analysensysteme GmbH: soli TOC cube: The versatile instrument for temperature-dependent differentiation of carbon in solids, Product Flyer Art.-No. 200014990, Elementar Analysensysteme GmbH, Langenselbold, Germany, 2023.
- Fuglstad, G.-A., Simpson, D., Lindgren, F., and Rue, H.: Constructing Priors That Penalize the Complexity of Gaussian Random Fields, *Journal of the American Statistical Association*, 114, 445–452, <https://doi.org/10.1080/01621459.2017.1415907>, 2019.
- 755 Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B.: *Bayesian Data Analysis*, Chapman & Hall/CRC, Boca Raton, FL, 2 edn., 2004.
- Gneiting, T., Balabdaoui, F., and Raftery, A. E.: Probabilistic forecasts, calibration and sharpness, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 69, 243–268, <https://doi.org/10.1111/j.1467-9868.2007.00587.x>, 2007.
- 760 GRASS Development Team: Geographic Resources Analysis Support System (GRASS GIS) Software, Version 8.4, Open Source Geospatial Foundation, USA, <https://doi.org/10.5281/zenodo.5176030>, 2025.



- Gupta, H. V., Kling, H., Yilmaz, K. K., and Martinez, G. F.: Decomposition of the mean squared error and NSE performance criteria: Implications for improving hydrological modelling, *Journal of Hydrology*, 377, 80–91, <https://doi.org/10.1016/j.jhydrol.2009.08.003>, 2009.
- 765 Gustafsson, Ö., Haghseta, F., Chan, C., MacFarlane, J., and Gschwend, P. M.: Quantification of the dilute sedimentary soot phase: Implications for PAH speciation and bioavailability, *Environmental Science & Technology*, 31, 203–209, <https://doi.org/10.1021/es960317s>, 1997.
- Hastie, T., Tibshirani, R., and Friedman, J.: *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, Springer, New York, NY, 2 edn., <https://doi.org/10.1007/978-0-387-84858-7>, 2009.
- 770 Heuvelink, G. B. M.: Analysing uncertainty propagation in GIS: why is it not that simple?, in: *Uncertainty in Remote Sensing and GIS*, edited by Foody, G. M. and Atkinson, P. M., pp. 155–165, Wiley, Chichester, <https://doi.org/10.1002/0470035269.ch10>, 2002.
- Hu, B., Geng, Y., Ni, H., Shi, Z., Wang, Z., Wang, N., Luo, J., Xie, M., Zou, Q., Opitz, T., and Li, H.: Mapping and understanding the regional farmland SOC distribution in southern China using a Bayesian spatial model, *Geoderma*, 460, 117446, <https://doi.org/10.1016/j.geoderma.2025.117446>, 2025.
- 775 Hu, Y., Cross, A., Shen, Z., Bouma, J., and Viscarra Rossel, R. A.: On soil health and the pivotal role of sensing, *SOIL*, 12, 227–252, <https://doi.org/10.5194/soil-12-227-2026>, 2026.
- Huang, J., Malone, B. P., Minasny, B., McBratney, A. B., and Triantafyllis, J.: Evaluating a Bayesian modelling approach (INLA-SPDE) for environmental mapping, *Science of the Total Environment*, 609, 621–632, <https://doi.org/10.1016/j.scitotenv.2017.07.201>, 2017.
- 780 Huang, Y., Song, X., Wang, Y.-P., Canadell, J. G., Luo, Y., Ciais, P., Chen, A., Hong, S., Wang, Y., Tao, F., Li, W., Xu, Y., Mirzaeitalarposhti, R., Elbasiouny, H., Savin, I., Shchepashchenko, D., Viscarra Rossel, R. A., Goll, D. S., Chang, J., Houlton, B. Z., Wu, H., Yang, F., Feng, X., Chen, Y., Liu, Y., Niu, S., et al.: Size, distribution, and vulnerability of the global soil inorganic carbon, *Science*, 384, 233–239, <https://doi.org/10.1126/science.adi7918>, 2024.
- 785 Isbell, R.: *The Australian Soil Classification*, CSIRO Publishing, Melbourne, Australia., 2016.
- IUSS Working Group WRB: *World Reference Base for Soil Resources: International Soil Classification System for Naming Soils and Creating Legends for Soil Maps*, International Union of Soil Sciences (IUSS), Vienna, Austria, 4 edn., [https://wrb.isric.org/files/WRB\\_fourth\\_edition\\_2022-12-18.pdf](https://wrb.isric.org/files/WRB_fourth_edition_2022-12-18.pdf), 2022.
- Janzen, H. H.: The soil carbon dilemma: Shall we hoard it or use it?, *Soil Biology and Biochemistry*, 38, 419–424, <https://doi.org/10.1016/j.soilbio.2005.10.008>, 2006.
- 790 Lark, R. M. and Webster, R.: Analysis and elucidation of soil variation using wavelets, *European Journal of Soil Science*, 50, 185–206, <https://doi.org/10.1046/j.1365-2389.1999.t01-1-00234.x>, 1999.
- Lavallee, J. M., Soong, J. L., and Cotrufo, M. F.: Conceptualizing soil organic matter into particulate and mineral-associated forms to address global change in the 21st century, *Global Change Biology*, 26, 261–273, <https://doi.org/10.1111/gcb.14859>, 2020.
- 795 Lindgren, F., Rue, H., and Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73, 423–498, <https://doi.org/10.1111/j.1467-9868.2011.00777.x>, 2011.



- McBratney, A. B., Mendonça Santos, M. L., and Minasny, B.: On digital soil mapping, *Geoderma*, 117, 3–52,  
800 [https://doi.org/10.1016/S0016-7061\(03\)00223-4](https://doi.org/10.1016/S0016-7061(03)00223-4), 2003.
- 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.
- Plaza, C., Zaccone, C., Sawicka, K., Mendez, A. M., Tarquis, A., Gascó, G., Heuvelink, G. B. M., Schuur, E. A. G.,  
and Maestre, F. T.: Soil resources and element stocks in drylands to face global issues, *Scientific Reports*, 8, 13 788,  
805 <https://doi.org/10.1038/s41598-018-32229-0>, 2018.
- Poggio, L., Gimona, A., Spezia, L., and Brewer, M. J.: Bayesian spatial modelling of soil properties and  
their uncertainty: The example of soil organic matter in Scotland using R-INLA, *Geoderma*, 277, 69–82,  
<https://doi.org/10.1016/j.geoderma.2016.04.026>, 2016.
- Quinlan, J. R.: *C4.5: Programs for Machine Learning*, Morgan Kaufmann, San Mateo, CA, 1993.
- 810 Ramírez-López, L., Schäpe, A. M., Demattê, J. A. M., Behrens, T., Stevens, A., and Stenberg, B.: Sampling optimal calibration  
sets in soil infrared spectroscopy, *Geoderma*, 226–227, 140–150, <https://doi.org/10.1016/j.geoderma.2014.02.002>, 2014.
- Ramírez-López, L., Wadoux, A. M. J.-C., Franceschini, M. H. D., Terra, F. S., Marques, K. P. P., Sayão, V. M., and Demattê,  
J. A. M.: Robust soil mapping at the farm scale with vis–NIR spectroscopy, *European Journal of Soil Science*, 70, 378–393,  
<https://doi.org/https://doi.org/10.1111/ejss.12752>, 2019.
- 815 Roberts, D. R., Bahn, V., Ciuti, S., Boyce, M. S., Elith, J., Guillera-Arroita, G., Hauenstein, S., Lahoz-Monfort, J. J., Schröder,  
B., Thuiller, W., Warton, D. I., Wintle, B. A., Hartig, F., and Dormann, C. F.: Cross-validation strategies for data with  
temporal, spatial, hierarchical, or phylogenetic structure, *Ecography*, 40, 913–929, <https://doi.org/10.1111/ecog.02881>, 2017.
- Rubin, D. B.: *Multiple Imputation for Nonresponse in Surveys*, John Wiley & Sons, New York,  
<https://doi.org/10.1002/9780470316696>, 1987.
- 820 Rubin, D. B.: Multiple Imputation after 18+ Years, *Journal of the American Statistical Association*, 91, 473–489,  
<https://doi.org/10.1080/01621459.1996.10476908>, 1996.
- Rue, H., Martino, S., and Chopin, N.: Approximate Bayesian inference for latent Gaussian models by using integrated  
nested Laplace approximations, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71, 319–392,  
<https://doi.org/10.1111/j.1467-9868.2008.00700.x>, 2009.
- 825 Savitzky, A. and Golay, M. J. E.: Smoothing and Differentiation of Data by Simplified Least Squares Procedures, *Analytical  
Chemistry*, 36, 1627–1639, <https://doi.org/10.1021/ac60214a047>, 1964.
- Schmidt, M. W. I. and Noack, A. G.: Black carbon in soils and sediments: Analysis, distribution, implications, and current  
challenges, *Global Biogeochemical Cycles*, 14, 777–793, <https://doi.org/10.1029/1999GB001208>, 2000.
- Sun, S., Cotrufo, M. F., Viscarra Rossel, R. A., Mueller, C. W., Kida, M., Hardie, A. G., Mackay, A., et al.: Global hotspots of  
830 particulate organic carbon losses under climate change, *Nature Communications*, 2026.
- Viscarra Rossel, R. A.: Robust Modelling of Soil Diffuse Reflectance Spectra by ‘Bagging-Partial Least Squares Regression’,  
*Journal of Near Infrared Spectroscopy*, 15, 39–47, <https://doi.org/10.1255/jnirs.694>, 2007.
- Viscarra Rossel, R. A., Bui, E. N., de Caritat, P., and McKenzie, N. J.: Mapping iron oxides and the colour of Australian soil  
using visible–near-infrared reflectance spectra, *Journal of Geophysical Research: Earth Surface*, 115, 2010.



- 835 Viscarra Rossel, R. A., Webster, R., Bui, E. N., and Baldock, J. A.: Baseline map of organic carbon in Australian soil to support national carbon accounting and monitoring under climate change, *Global Change Biology*, 20, 2953–2970, <https://doi.org/10.1111/gcb.12569>, 2014.
- Viscarra Rossel, R. A., Chen, C., Grundy, M., Searle, R., Clifford, D., and Campbell, P.: The Australian three-dimensional soil grid: Australia’s contribution to the GlobalSoilMap project, *Soil Research*, 53, 845–864, 2015.
- 840 Viscarra Rossel, R. A., Brus, D. J., Lobsey, C., Shi, Z., and McLachlan, G.: Baseline estimates of soil organic carbon by proximal sensing: Comparing design-based, model-assisted and model-based inference, *Geoderma*, 265, 152–163, 2016.
- Viscarra Rossel, R. A., Lee, J., Behrens, T., Luo, Z., Baldock, J., and Richards, A.: Continental-scale soil carbon composition and vulnerability modulated by regional environmental controls, *Nature Geoscience*, 12, 547–552, <https://doi.org/10.1038/s41561-019-0373-z>, 2019.
- 845 Viscarra Rossel, R. A., Behrens, T., Ben-Dor, E., Chabrillat, S., Demattê, J. A. M., Ge, Y., Gomez, C., Guerrero, C., Peng, Y., Ramirez-Lopez, L., Shi, Z., Stenberg, B., Webster, R., Winowiecki, L., and Shen, Z.: Diffuse reflectance spectroscopy for estimating soil properties: A technology for the 21st century, *European Journal of Soil Science*, 73, e13271, <https://doi.org/10.1111/ejss.13271>, 2022.
- Viscarra Rossel, R. A., Webster, R., Zhang, M., Shen, Z., Dixon, K., Wang, Y.-P., and Walden, L.: How much organic
- 850 carbon could the soil store? The carbon sequestration potential of Australian soil, *Global Change Biology*, 30, e17053, <https://doi.org/10.1111/gcb.17053>, 2024.
- Walden, L., Sepanta, F., and Viscarra Rossel, R. A.: FT-MIR Spectroscopic Analysis of the Organic Carbon Fractions in Australian Mineral Soils, *European Journal of Soil Science*, <https://doi.org/10.1111/ejss.70084>, article ejss.70084, 2025.
- Wilford, J., de Caritat, P., and Bui, E.: Modelling the abundance of soil calcium carbonate across Australia using geochemical
- 855 survey data and environmental predictors, *Geoderma*, 259–260, 81–92, <https://doi.org/10.1016/j.geoderma.2015.05.003>, 2015.
- Zhang, M., Shen, Z., Walden, L., Sepanta, F., Luo, Z., Gao, L., Serrano, O., and Viscarra Rossel, R. A.: Deep learning of the particulate and mineral-associated organic carbon fractions using a compositional transform and mid-infrared spectroscopy, *Geoderma*, 455, 117207, <https://doi.org/10.1016/j.geoderma.2025.117207>, 2025a.
- 860 Zhang, M., Walden, L., Shen, Z., and Viscarra Rossel, R. A.: Environmental drivers of soil organic carbon in the Australian rangelands, *Catena*, 255, 108952, <https://doi.org/10.1016/j.catena.2025.108952>, 2025b.
- Zhou, X. and Reiter, J. P.: A Note on Bayesian Inference After Multiple Imputation, *The American Statistician*, 64, 159–163, <https://doi.org/10.1198/tast.2010.09109>, 2010.