

UC Irvine

UC Irvine Previously Published Works

Title

NONSTATIONARY SPATIAL PREDICTION OF SOIL ORGANIC CARBON: IMPLICATIONS FOR STOCK ASSESSMENT DECISION MAKING.

Permalink

<https://escholarship.org/uc/item/88654309>

Journal

The Annals of Applied Statistics, 13(1)

ISSN

1932-6157

Authors

Risser, Mark
Calder, Catherine
Berrocal, Veronica
[et al.](#)

Publication Date

2019-03-01

DOI

10.1214/18-AOAS1204

Peer reviewed



Published in final edited form as:

Ann Appl Stat. 2019 March ; 13(1): 165–188. doi:10.1214/18-AOAS1204.

NONSTATIONARY SPATIAL PREDICTION OF SOIL ORGANIC CARBON: IMPLICATIONS FOR STOCK ASSESSMENT DECISION MAKING¹

MARK D. RISSER^{*}, CATHERINE A. CALDER^{†,2}, VERONICA J. BERROCAL[‡], CANDACE BERRETT[§]

^{*} Lawrence Berkeley National Laboratory

[†] Ohio State University

[‡] University of Michigan

[§] Brigham Young University

Abstract

The Rapid Carbon Assessment (RaCA) project was conducted by the US Department of Agriculture's National Resources Conservation Service between 2010–2012 in order to provide contemporaneous measurements of soil organic carbon (SOC) across the US. Despite the broad extent of the RaCA data collection effort, direct observations of SOC are not available at the high spatial resolution needed for studying carbon storage in soil and its implications for important problems in climate science and agriculture. As a result, there is a need for predicting SOC at spatial locations not included as part of the RaCA project. In this paper, we compare spatial prediction of SOC using a subset of the RaCA data for a variety of statistical methods. We investigate the performance of methods with off-the-shelf software available (both stationary and nonstationary) as well as a novel nonstationary approach based on partitioning relevant spatially-varying covariate processes. Our new method addresses open questions regarding (1) how to partition the spatial domain for segmentation-based nonstationary methods, (2) incorporating partially observed covariates into a spatial model, and (3) accounting for uncertainty in the partitioning. In applying the various statistical methods we find that there are minimal differences in out-of-sample criteria for this particular data set, however, there are major differences in maps of uncertainty in SOC predictions. We argue that the spatially-varying measures of prediction uncertainty produced by our new approach are valuable to decision makers, as they can be used to better benchmark mechanistic models, identify target areas for soil restoration projects, and inform carbon sequestration projects.

¹Supported in part by the Statistical Methods for Atmospheric and Oceanic Sciences (STATMOS) research network (NSF DMS awards 1106862, 1106974, and 1107046).

²Supported in part by NSF DMS award 1209161.

SUPPLEMENTARY MATERIAL

Marginal likelihood estimation and supporting plots (DOI: [10.1214/18-AOAS1204SUPP](https://doi.org/10.1214/18-AOAS1204SUPP); .pdf). We provide additional supporting text regarding marginal likelihood estimation and supporting plots that show marginal likelihood estimates and holdout sets.

Keywords

Gaussian process; spatial clustering; model averaging; soil carbon; spatial regression

1. Introduction.

Oceans, terrestrial systems, and the atmosphere form the three primary carbon reservoirs on the earth (Batjes (1996)). The amount of carbon in terrestrial ecosystems is nearly three times that of the atmosphere, and while its size is dwarfed by the ocean's carbon storage, terrestrial carbon is much more dynamic (Batjes (1996)). Soil organic carbon (SOC), a generic term for the carbon found in soil's organic matter, constitutes approximately two-thirds of the carbon in terrestrial ecosystems and is an important component in the earth's carbon cycle (Post et al. (1982)). The carbon in soil is in continuous interaction with the atmosphere via processes such as plant growth and decomposition (Bliss et al. (2014)), and SOC helps mitigate the negative consequences of global changes in climate by sequestering carbon released into the atmosphere by fossil fuel combustion (Jobbágy and Jackson (2000); Post et al. (1982)). In addition to its relevance to the earth's climate, SOC is important in forestry and agriculture, as organic matter contributes to soil fertility by helping retain moisture and supply plant nutrients (Bliss et al. (2014); Post and Kwon (2000)). Furthermore, SOC is one of the soil properties used by hydrologists to better understand how precipitation is processed by different land surfaces and contributes to surface and ground water quality (Bliss et al. (2014)).

Because of the broad relevance of SOC to climate and agriculture, the National Resources Conservation Service (NRCS) initiated the Rapid Carbon Assessment (RaCA) project in 2010 (Wills et al. (2013)). The goal of the RaCA project was to collect measurements of the carbon content of soil across the conterminous United States at a single point in time. Specifically, the project emphasized collecting spatially-referenced SOC measurements or "stocks," that is, the amount of SOC in a volume (area and depth of soil) to produce "statistically reliable quantitative estimates of amounts and distribution of carbon stocks for U.S. soils under various land covers" (Wills et al. (2013)). Since the collection of SOC data is highly limited by time and cost constraints (Sleutel et al. (2003); Goidts and van Wesemael (2007)), SOC is measured only at limited locations. Thus, there is a need for statistical methods to predict SOC concentration at unobserved locations using the RaCA data.

In this paper, we consider the problem of spatial prediction of SOC concentration based on RaCA measurements. While geostatistical methods have previously been used for this purpose (albeit not for the RaCA data set; see Simbahan et al. (2006)), these analyses have relied on an assumption of second-order stationarity in SOC over space. As we demonstrate in Section 2, this assumption is inappropriate because SOC shows evidence of second-order nonstationary behavior on regional scales. However, existing methods (both stationary and nonstationary) with off-the-shelf software are inadequate for spatial prediction of SOC, as they yield either unrealistic prediction maps or a noninformative characterization of prediction error. Since it is well documented that SOC is influenced by

covariate information such as land use (Jobbágy and Jackson (2000)) and soil properties (Mishra et al. (2009)), we propose a novel approach for spatial prediction of SOC that uses these spatially-referenced covariates to describe the second-order nonstationarity in SOC. Existing approaches for covariate-driven nonstationary spatial modeling (Calder (2008); Schmidt, Guttorp and O’Hagan (2011); Reich et al. (2011); Vianna Neto, Schmidt and Guttorp (2014); Ingebrigtsen, Lindgren and Steinsland (2014); Risser and Calder (2015)) require fully-observed spatial covariates, however, in this case the relevant covariates are not observed everywhere a prediction is desired. To address this limitation, we propose a covariate-partitioning approach to nonstationary spatial modeling that uses spatially-referenced covariate information to divide the spatial domain into distinct “segments.” Once defined, these segments partition the domain such that every location is contained in exactly one of these segments, and the SOC process within each segment is assumed to be locally stationary conditional on a particular segmentation. Since it is the segment membership that will be used in our statistical model, not the value of the covariate, spatial prediction does not require covariates to be observed at the prediction location.

In spite of the fact that the SOC data display meaningful nonstationarities on regional scales, we acknowledge that quantitative evaluation criteria do not indicate a strong preference for our nonstationary statistical model, relative to approaches with off-the-shelf software. This is not surprising (see, e.g., Fuglstad et al. (2015)); in any case, another important feature of our approach is that we are able to capture spatial variation in the uncertainty associated with SOC concentration predictions. Unlike traditional geostatistical methods where, for a fixed set of monitoring sites, spatial variation in prediction uncertainty is driven primarily by variation in the geographical distribution of the data (e.g., uncertainty is greatest at locations far from the observations), our approach readily captures how covariates such as land use and soil properties impact the strength of spatial dependence in SOC, which directly informs the assessment of soil carbon stocks.

The paper proceeds as follows: in Section 2, we introduce the RaCA data in more detail and conduct exploratory analyses, and in Section 3 we introduce relevant explanatory variables and their subsequent partitioning. Section 4 outlines our covariate-partitioning approach for nonstationary spatial prediction of SOC, and in Section 5 we present the results of our fitted model, predictions, and implications for carbon stock assessment. Section 6 concludes the paper.

2. RaCA data and exploratory analyses.

The RaCA data and additional variables such as land use-land cover (LULC) classes, soil series, and soil moisture are available through the `soilDB` package in `R` (Beaudette and Skovlin (2015)), and a summary report of the sampling methods and data description is provided in Wills et al. (2013). While measurements of SOC stocks are available for each site to depths of five, 10, 20, 30, 50, and 100 centimeters (in Mg C ha^{-1}), we use the one hundred centimeter depth measurement since our focus is on estimating total SOC, not a soil depth profile (e.g., Minasny et al. (2006); Mishra et al. (2009)). Our study focuses on a subset of the RaCA SOC measurements collected in the Great Lakes region of the midwestern United States, shown in Figure 1, which contains 790 observations.

2.1. Variogram analysis.

As discussed in Section 1, accurate estimation of the spatial distribution of SOC is hindered by costly and time-consuming data collection, and we are thus motivated to consider spatial statistical methods to predict SOC concentration at unobserved locations. Focusing on the Great Lakes region, we begin with a simple variogram analysis of the SOC data, first calculating the empirical semivariogram and corresponding fitted exponential semivariogram for the entire region. Here and throughout the remainder of the paper, the SOC data is transformed to the log scale; for the variogram analyses, we use the ordinary least squares (OLS) residuals from a regression of log SOC on latitude, longitude, and the longitude/latitude interaction. Exploratory analysis indicates that the exponential correlation model fits the data well; the fitted exponential semivariogram is shown in Figure 2(b) with 95% confidence band (using the parametric bootstrap).

A simple way to assess the presence of second-order nonstationary behavior in a spatial data set is to split up the spatial domain into subregions and conduct a variogram analysis separately for each subregion. Arbitrarily dividing the Great Lakes region into four parts, we fit an exponential semivariogram to the OLS residuals in each subregion. The subregion-specific semivariograms with 95% confidence bands (again using the parametric bootstrap) are shown in Figure 2(d): the fitted semivariograms indicate that the subregions display quite different spatial dependence patterns. These differences are significant, as indicated by the nonoverlapping uncertainty bands, which motivates a nonstationary spatial model for the SOC data where the spatial dependence properties vary over the spatial domain.

2.2. Off-the-shelf spatial prediction.

Various packages in \mathbb{R} are readily available for the analysis of spatial data, some developed for stationary processes and some accommodating nonstationary ones. Here, we summarize and report on the prediction results obtained by applying different models with off-the-shelf software to our SOC data. All models were fitted using 20,000 MCMC iterations with the first 10,000 discarded as burn-in and with no thinning of the chain. All the spatial models use an exponential correlation structure, as this was deemed to provide a good representation of the dependence structure in the data, and all use uninformative, most times uniform, priors on model parameters.

Bayesian additive regression trees.—Bayesian additive regression trees (BART; Chipman, George and McCulloch (2010)) is a Bayesian sum-of-trees model that has been shown to perform well in a variety of settings, particularly for prediction (see, e.g., Bonato et al. (2011), Ding et al. (2012), Green and Kern (2012)). Because of its predictive skill, we spatially predict (log) SOC using BART as implemented in the `BayesTree` package for \mathbb{R} . In applying BART to the SOC data, we use latitude and longitude as covariates for the mean function.

Treed Gaussian process.—Related to BART is the nonstationary treed Gaussian process (TGP) model of Gramacy and Lee (2008) in which tree partitions are defined by segmenting the coordinate axes, and partitioning uncertainty is accounted for through a model averaging approach (Hoeting et al. (1999)). The TGP model is implemented in the `tgp` package for \mathbb{R} .

(Gramacy (2007)), which we use with all the defaults on (log) SOC. As in BART, latitude and longitude are used as main effects in the mean function.

Gaussian predictive process.—Given the moderately large dimension of the RaCA dataset, we also consider an off-the-shelf implementation of the Gaussian predictive process (PP) model of Banerjee et al. (2008), which is specifically designed to account for large spatial data sets. The PP works by constructing an approximation to a Gaussian process by projecting realizations of the process of interest onto a lower dimensional space spanned by the predictive process knots, thereby reducing the computational burden. By construction, the PP is technically nonstationary, and the PP model provides a nonstationary approximation to any stationary covariance function. We fit a PP model to our SOC data using the `spBayes` package for R (Finley, Banerjee and Carlin (2007); Finley, Banerjee and Gelfand (2013)) specifying 103 knots and a constant mean.

Bayesian stationary Gaussian process.—The PP model reduces to a Bayesian stationary Gaussian process model if the predictive process knots are taken to be exactly the observation locations. Using the `spBayes` package and the `spLM` function with default prior settings, we generate spatial predictions of log SOC using a traditional Bayesian stationary (isotropic) spatial Gaussian process model with mean (log) SOC specified as a linear function of longitude and latitude.

Rows 2–5 in Table 1 summarize the different models applied to the (log) SOC data, with details on the mean function specification, the R package used to implement the model, and the computational time that it takes to fit each model. Surfaces of (log) SOC obtained by generating predictions on a fine grid covering the entire spatial domain using each of the aforementioned methods are presented in Figure 3. Specifically, from left to right, the panels show the posterior mean (top) and standard deviation (bottom) of the predictions using BART, TGP, SGP, and PP. Clearly, none of these models appear to be appropriate for the data: BART does not yield scientifically meaningful predictions with the artificial horizontal and vertical lines, and while the three other spatial models produce prediction maps that are generally smooth they are characterized by uncertainty prediction maps that are either unrealistic for an environmental process like log SOC (see TGP), or uninformative since they are almost constant in space.

To address the limitations of these off-the-shelf methods, and to model the globally nonstationary but locally stationary nature of the SOC data highlighted in Section 2, in Sections 3 and 4 we move on to introduce a novel segmentation-based model for nonstationary spatial processes with partitions informed by covariates.

3. Covariate-driven partitioning.

Returning to the exploratory analysis in Section 2, it is clear that partitioning the domain captures the nonstationary behavior in SOC for the Great Lakes region. However, there are three problems associated with the subregion-specific variogram analyses. First, the arbitrary partitioning of our domain as in Figure 2 is not scientifically meaningful: in other words, it provides no way of understanding why the different subregions exhibit nonstationarities.

Second, generating predictions of SOC based only on separately or independently fitted variograms for each subregion does not comprise an appropriate spatial model, because no information is shared across the subregions. Third, the fitted semivariogram estimates are likely sensitive to the specific partition used in Figure 2; in other words, we might want to account for uncertainty in the partition itself.

A solution to the first problem is to partition the spatial domain based on covariate information. As discussed in Section 1, there are well documented relationships between SOC and a number of covariate variables, for example, land use-land cover (LULC) class and drainage class. Both of these categorical variables are available in the `soilDB` package in \mathbb{R} (Beaudette and Skovlin (2015)). According to Wills et al. (2013), the LULC classes were developed specifically for the RaCA project to correspond to the classes and definitions of the Natural Resources Inventory. RaCA designated five specific LULC classes (four of which are represented in the Great Lakes region subset), namely cropland, farmland, pastureland, rangeland, and wetland, with one additional category for any cropland site that was also known to correspond to a Conservation Reserve Program (CRP). The drainage class variable refers to “the frequency and duration of wet periods under conditions similar to those under which the soil developed. Alteration of the water regime by man, either through drainage or irrigation, is not a consideration unless the alterations have significantly changed the morphology of the soil” (NRCS Soil Survey Manual, Chapter 3). Plots of the subsetted LULC and drainage class variables are shown in Figure 4.

Given the known relationships between SOC and these variables as well as the rich literature on covariate-driven nonstationary modeling, we are motivated to use these covariates to describe the second-order nonstationary behavior exhibited by SOC over the Great Lakes region. In other words, our hypothesis is that both the first- and second-order properties of SOC might be similar in areas where the covariates are homogeneous. However, measurements of LULC and drainage class are not available for every prediction location of interest, which render the nonstationary models of Reich et al. (2011), Vianna Neto, Schmidt and Guttorp (2014), Ingebrigtsen, Lindgren and Steinsland (2014), and Risser and Calder (2015) unusable in this situation.

Hence, here we propose to partition the spatial domain based on the multivariate spatial distribution of LULC and drainage class. This provides scientifically meaningful partitioning, and also accounts for the fact that we are dealing with incompletely observed covariates. Having defined a partitioning of the spatial domain, we can use an approach similar to one outlined in Fuentes (2001) (see Section 4) to define a globally nonstationary, locally-stationary spatial process. Finally, to both account for uncertainty in the partitioning *and* introduce spatial dependence in SOC across subregions, we propose to use a Bayesian model averaging framework (Hoeting et al. (1999)).

There are a variety of methods in the statistics literature for obtaining partitions of a multivariate space (here, geographic space, i.e., latitude and longitude) based on multivariate inputs (here, LULC and drainage class). For this paper, we use multivariate cluster-wise regression, also called multivariate latent class regression (see, e.g., Leisch (2004); Müller, Quintana and Rosner (2011)). Before describing this approach, we establish some

terminology: we define a *partition* of the spatial domain \mathcal{D} as the assignment of each location in the geographical space to a particular subregion. In a partition, each location is assigned to exactly one subregion and collectively the subregions comprise the entire spatial domain. Also, we define a *segment* as an individual subregion: therefore, a partition is made up of a set of segments. We use \mathcal{P} to denote a generic partition, made up of segments $\{\mathcal{S}_1, \dots, \mathcal{S}_K\}$. By definition, $\mathcal{P} = \cup_k \mathcal{S}_k$ and $\mathcal{S}_k \cap \mathcal{S}_{k'} = \emptyset$ for $k \neq k'$.

It is trivial to note that determining the segments of a partition \mathcal{P} of the spatial domain \mathcal{D} is equivalent to clustering the geographical coordinates of points within \mathcal{D} . In particular, if the segments of the partition are to be characterized by the fact that the covariate processes are similar within each segment, then the clustering mechanism ought to be informed by the covariates. Leisch (2004) and Müller, Quintana and Rosner (2011) propose two different approaches to achieve covariate-informed clustering: the first approach falls in the model-based clustering category, and the second falls in the class of product partition models (thus, specified in a Bayesian nonparametric framework). For computational simplicity, we follow Leisch (2004) and use a model-based clustering approach. Working with a two-dimensional spatial domain (i.e., latitude and longitude of locations) we model the bivariate vector of coordinates (s_1, s_2) for each point $\mathbf{s} \in \mathcal{D}$ as arising from a mixture of bivariate normal distributions where the mean and the covariance matrix are mixture-component specific. In other words, given a pre-specified and finite number K of clusters, we assume

$$\mathbf{s} = (s_1, s_2) \sim \sum_{k=1}^K \pi_k(\mathbf{s}) \cdot N_2(\mathbf{s}; \mathbf{m}_k, \mathbf{D}_k), \tag{3.1}$$

where $N_d(x; m, V)$ is the d -variate Gaussian density with mean m and covariance V evaluated at x . The `flexmix` package in \mathbb{R} allows us to fit this class of models for fixed K using the EM algorithm. Taking the observation locations of the two spatial covariates (LULC and drainage class) as data for model (3.1) and using the `flexmix` package with K ranging from 2 to 6, we obtain multiple potential partitions of the observation locations of the covariate processes, shown in Figure 5. The mixture for a specified number of segments is nonunique (in other words, the EM algorithm converges to several different local modes), and after re-starting the algorithm several times for each K we selected the best mixture(s) based on EM convergence and the log-likelihood.

However, note that Figure 5 only provides partitions for locations where we have a measurement of the spatial covariates. As will be seen in Section 4, we require partitions of the locations where we have SOC measurements (which differ from the covariate measurement locations) as well as all locations on a fine grid for generating predictions. When fitting model (3.1) to the latitude and longitude coordinates of the covariate locations using the package `flexmix`, we obtain as a byproduct the bivariate Gaussian densities from (3.1), which can be defined for any location $\mathbf{s} \in \mathcal{D}$. Based on the K segment probabilities, we assign each location to a segment by taking the maximum of the bivariate Gaussian densities, that is, the segment for location \mathbf{s} is defined as $\max_k \{N_2(\mathbf{s}; \mathbf{m}_k, \mathbf{D}_k)\}$. Thus the

eight mixtures in Figure 5 yield eight partitions of the spatial domain (not shown), denoted $\{\mathcal{P}_j; j = 1, \dots, 8\}$, with segments denoted $\{\mathcal{S}_{jk}; k = 1, \dots, K_j\}$.

4. A partition-based nonstationary spatial Gaussian process model.

The covariate partitions $\{\mathcal{P}_j; j = 1, \dots, 8\}$ defined in Section 3 can be used to model both first- and second-order nonstationarities in SOC as follows. Let $Z(\cdot)$ represent observed log SOC, where we model $Z(\cdot)$ as a spatial stochastic process defined for all $\mathbf{s} \in \mathcal{D}$ (here, \mathcal{D} denotes the Great Lakes region of the United States). For all $\mathbf{s} \in \mathcal{D}$, let

$$Z(\mathbf{s}) = \mu(\mathbf{s}) + Y(\mathbf{s}) + \varepsilon(\mathbf{s}), \tag{4.1}$$

where $E[Z(\mathbf{s})] = \mu(\mathbf{s})$ is a deterministic mean function, $Y(\cdot)$ is a mean-zero latent spatial Gaussian process, and $\varepsilon(\cdot)$ is an error process that is assumed to be independent of $Y(\cdot)$. We observe the value of $Z(\cdot)$ at a fixed, finite set of locations $\{\mathbf{s}_1, \mathbf{s}_2, \dots, \mathbf{s}_n\} \in \mathcal{D}$ (see Figure 1) and wish to use these observations to learn about the underlying processes and generate predictions at unobserved locations.

In Section 4.1, we outline a statistical model for log SOC conditional on a single partition \mathcal{P}_j , and in Section 4.2 we outline the model fitting for each conditional model. Then, in Section 4.3 we describe a model-averaging approach to posterior prediction that incorporates all candidate partitions.

4.1. Conditional model specification.

Conditional on the K_j segments $\{\mathcal{S}_{jk}; k = 1, \dots, K_j\}$ of partition \mathcal{P}_j , we first model $Y(\cdot)$ as a mixture of stationary processes (Fuentes (2001)), that is,

$$Y(\mathbf{s}) = \sum_{k=1}^{K_j} w_{jk}(\mathbf{s}) \tilde{Y}_{jk}(\mathbf{s}), \tag{4.2}$$

where the $\tilde{Y}_{jk}(\cdot)$ are orthogonal and stationary, and $w_{jk}(\cdot)$ is a positive kernel weight function such that $w_{jk}(\mathbf{s}) \geq 0$ and $\sum_{k=1}^{K_j} [w_{jk}(\mathbf{s})]^2 = 1$ for all $\mathbf{s} \in \mathcal{D}$ (following Reich et al. (2011)). Define the covariance function of each $\tilde{Y}_{jk}(\cdot)$ to be \tilde{C}_{jk} ; then, the covariance function of $Y(\cdot)$ is

$$\text{Cov}(Y(\mathbf{s}), Y(\mathbf{s}')) \equiv C(\mathbf{s}, \mathbf{s}') = \sum_{k=1}^{K_j} w_{jk}(\mathbf{s}) w_{jk}(\mathbf{s}') \tilde{C}_{jk}(\mathbf{s} - \mathbf{s}'), \tag{4.3}$$

which is a valid nonstationary covariance function. For each \tilde{C}_{jk} , we use an anisotropic version of the parametric Matérn model

$$\tilde{C}_{jk}(\mathbf{s}-\mathbf{s}'; \boldsymbol{\theta}_{jk}) = \frac{\sigma_{jk}^2}{\Gamma(\nu_{jk})2^{\nu_{jk}-1}} [\sqrt{\mathbf{Q}_{jk}(\mathbf{s}-\mathbf{s}')}]^{\nu_{jk}} \mathcal{B}_{\nu_{jk}}(\sqrt{\mathbf{Q}_{jk}(\mathbf{s}-\mathbf{s}')}). \tag{4.4}$$

In (4.4), $\mathcal{B}_{\nu_{jk}}(\cdot)$ denotes the modified Bessel function of the third kind of order ν_{jk} , $\mathbf{Q}_{jk}(\mathbf{s}-\mathbf{s}') = \|\boldsymbol{\Sigma}_{jk}^{-1/2}(\mathbf{s}-\mathbf{s}')\|^2$ is a squared Mahalanobis distance with anisotropy matrix $\boldsymbol{\Sigma}_{jk}$ parameterized according to its spectral decomposition, that is (for $d=2$),

$$\boldsymbol{\Sigma}_{jk} = \begin{bmatrix} \cos(\eta_{jk}) & -\sin(\eta_{jk}) \\ \sin(\eta_{jk}) & \cos(\eta_{jk}) \end{bmatrix} \begin{bmatrix} \phi_{jk}^{(1)} & 0 \\ 0 & \phi_{jk}^{(2)} \end{bmatrix} \begin{bmatrix} \cos(\eta_{jk}) & \sin(\eta_{jk}) \\ -\sin(\eta_{jk}) & \cos(\eta_{jk}) \end{bmatrix}, \tag{4.5}$$

and $\boldsymbol{\theta}_{jk} = (\sigma_{jk}^2, \nu_{jk}, \phi_{jk}^{(1)}, \phi_{jk}^{(2)}, \eta_{jk})$ is a vector of parameters that control the variance, smoothness, and anisotropy of $\tilde{Y}_{jk}(\cdot)$. In the anisotropy matrices (4.5), $\phi_{jk}^{(1)}$ and $\phi_{jk}^{(2)}$ represent directional “ranges” (i.e., inverse decay parameters) and η_{jk} represents an angle of rotation; these parameters allow for locally elliptical correlation patterns.

In general, we might propose a similar model for the mean behavior $\mu(\cdot)$, that is $\mu(\mathbf{s}) = \sum_{k=1}^{K_j} w_{jk}(\mathbf{s}) \tilde{\mu}_{jk}(\mathbf{s})$, where $\tilde{\mu}_{jk}(\mathbf{s})$ can accommodate any mean structure, including (in the case of a linear mean) fully observed covariates and segment-specific intercepts. For log SOC, we considered a variety of segment-specific mean functions involving latitude and longitude, but found that using a constant mean across all segments performed as well as a model with a global mean latitude and longitude coefficients. Furthermore, a model with constant spatial mean performed as well as a model with a different intercept in each segment. Thus, we set $\tilde{\mu}_{jk}(\mathbf{s}) \equiv \mu_j$ for all k .

To complete the specification of our model (4.1), we suppose that the error process $\varepsilon(\mathbf{s})$ is spatially independent and Gaussian. Similar to (4.2), we model $\varepsilon(\mathbf{s}) = \sum_{k=1}^{K_j} w_{jk}(\mathbf{s}) \tilde{\varepsilon}_{jk}(\mathbf{s})$, where each $\tilde{\varepsilon}_{jk}(\mathbf{s}) \stackrel{\text{i.i.d.}}{\sim} N(0, \tau_{jk}^2)$. Thus, we expand the $\boldsymbol{\theta}_{jk}$ vector to include the variance τ_{jk}^2 ; collect all of the variance/covariance parameters across segments into a single vector

$$\boldsymbol{\theta}^{(j)} = \{\theta_{j1}, \dots, \theta_{jK_j}\}.$$

In this paper, the weight functions $w_{jk}(\cdot)$ for partition \mathcal{P}_j are chosen to be indicator functions for segment \mathcal{S}_{jk} , that is, $w_{jk}(\mathbf{s}) = \mathbb{1}(\mathbf{s} \in \mathcal{S}_{jk})$ (following, e.g., Gramacy and Lee (2008)). In this case, for partition \mathcal{P}_j , the process $Y(\cdot)$ and therefore $Z(\cdot)$ is now locally stationary within each \mathcal{S}_{jk} and independent across the \mathcal{S}_{jk} , conditional on \mathcal{P}_j . For the random observed vector $\mathbf{Z} \equiv (Z(\mathbf{s}_1), \dots, Z(\mathbf{s}_n))^T$, the indicator weight function implies a conditional likelihood for \mathbf{Z} ,

$$p(\mathbf{Z} \mid \mu_j, \boldsymbol{\theta}^{(j)}, \mathcal{P}_j) \propto \prod_{k=1}^{K_j} |\mathbf{V}_{jk} + \boldsymbol{\Omega}_{jk}|^{-1/2}$$

(4.6)

$$\times \exp\left\{-\frac{1}{2}(\mathbf{Z}_{jk} - \mu_j \mathbf{1}_{jk})^\top (\mathbf{V}_{jk} + \mathbf{\Omega}_{jk})^{-1} (\mathbf{Z}_{jk} - \mu_j \mathbf{1}_{jk})\right\},$$

which is the product of segment-specific multivariate Gaussian likelihoods. In (4.6), the “ jk ” subscript partitions each term into its partition- and segment-specific components, for example, $\mathbf{Z}_{jk} = \{Z(\mathbf{s}_i): \mathbf{s}_i \in \mathcal{S}_{jk}\}$; $\mathbf{V}_{jk} = \tau_{jk}^2 \mathbf{I}$ captures a partition-/segment-specific measurement error variance; the elements of $\mathbf{\Omega}_{jk}$ come from $\tilde{\mathcal{C}}_{jk}$. Note that the likelihood for \mathbf{Z} in (4.6) is *conditional* on \mathcal{P}_j . As a result, the partition controls the second-order properties of $Z(\cdot)$ by determining independent regions of local stationarity. As mentioned previously, the partition could also specify first-order properties, although for SOC we have not used this property.

Still conditional on \mathcal{P}_j , we factor the prior distribution as

$p(\mu_j, \theta^{(j)} | \mathcal{P}_j) = p(\mu_j | \mathcal{P}_j) \times p(\theta^{(j)} | \mathcal{P}_j)$. The prior on μ_j is proper but noninformative and conjugate for the likelihood (4.6), that is, $p(\mu_j | \mathcal{P}_j) = N(0, 100^2)$. The prior for $\theta^{(j)}$ is conditional on the partition since \mathcal{P}_j controls the dimension of $\theta^{(j)}$. Based on the variogram analysis in Section 2, we choose to use a smoothness (ν) which is constant across segments and fixed to be $\nu = 0.5$ (true for all partitions), corresponding to an exponential correlation structure in (4.4). Otherwise,

$$p(\theta^{(j)} | \mathcal{P}_j) = \prod_{k=1}^{K_j} p(\tau_{jk}^2) p(\sigma_{jk}^2) p(\phi_{jk}^{(1)}) p(\phi_{jk}^{(2)}) p(\eta_{jk}) \tag{4.7}$$

(the conditioning on \mathcal{P}_j on the right-hand side is suppressed). Each component of (4.7) is noninformative and proper: for $k = 1, \dots, K_j$,

$$p(\tau_{jk}^2) = \text{Uniform}(0, 100), \quad p(\sigma_{jk}^2) = \text{Uniform}(0, 100),$$

$$p(\phi_{jk}^{(1)}) = \text{Uniform}(0, \sqrt{2}), \quad p(\phi_{jk}^{(2)}) = \text{Uniform}(0, \sqrt{2}),$$

$$p(\eta_{jk}) = \text{Uniform}[0, \pi/2].$$

Uniform priors for the variance parameters τ^2 and σ^2 are used in place of more traditional conjugate inverse-Gamma priors to avoid prior bias in the case where an individual segment contains a small number of observations (following Gelman (2006)). For this analysis, the longitude/latitude coordinates of locations within each cluster were rescaled to lie in $[0,1] \times [0,1]$ to improve mixing of the MCMC; therefore, the upper limits for the priors on the anisotropy matrix eigenvalues, $\phi_{jk}^{(1)}$ and $\phi_{jk}^{(2)}$ (which correspond to squared ranges of

dependence) are set to $\sqrt{2}$, which is the maximum distance on $[0,1] \times [0,1]$. The reasoning behind this choice is that the squared spatial range of the process within a segment is not expected to exceed the size of the segment. The limits on the prior for η_{jk} are set to ensure identifiability (following, e.g., Katzfuss (2013)).

4.2. Model fitting and MCMC.

Conditional on partition \mathcal{P}_j , the posterior distribution is

$$p(\mu_j, \theta^{(j)} \mid \mathcal{P}_j, \mathbf{Z} = \mathbf{z}) \propto p(\mathbf{Z} \mid \mu_j, \theta^{(j)}, \mathcal{P}_j) \cdot p(\mu_j \mid \mathcal{P}_j) \cdot p(\theta^{(j)} \mid \mathcal{P}_j), \tag{4.8}$$

which combines (4.6) and priors defined in Section 4.1. As usual, (4.8) is not available in closed form, and we must resort to Markov chain Monte Carlo (MCMC) methods. Posterior samples from (4.8) are generated using the `nimble` package for R (de Valpine et al. (2017)): the MCMC for \mathcal{P}_j is run for 20,000 total iterations with 10,000 iterations discarded as burn-in. A univariate Metropolis Hastings sampling step is conducted for the overall mean (μ_j), and adaptive random walk samplers are conducted for the segment-specific variance/covariance parameters ($[\tau_{jk}^2, \sigma_{jk}^2, \phi_{jk}^{(1)}, \phi_{jk}^{(2)}, \eta_{jk}]$, sampled as a block).

4.3. Model-averaged posterior prediction.

Recall from Section 1 that our primary goal for using this model is prediction. Define a collection of locations $\{s_1^*, \dots, s_m^*\} \subset \mathcal{D}$ for which we would like to obtain predictions of the corresponding (log) SOC values $\mathbf{Z}^* = (Z(s_1^*), \dots, Z(s_m^*))$. While the statistical model outlined in Section 4.1 is conditional on a single partition \mathcal{P}_j , using a Bayesian framework we can average over the partitions $\{\mathcal{P}_j: j = 1, \dots, 8\}$ from Section 3 to obtain the full posterior predictive distribution for \mathbf{Z}^* conditional on observed $\mathbf{Z}=\mathbf{z}$:

$$p(\mathbf{Z}^* \mid \mathbf{Z} = \mathbf{z}) = \sum_{j=1}^8 \iint p(\mathbf{Z}^*, \mu_j, \theta^{(j)}, \mathcal{P}_j \mid \mathbf{Z} = \mathbf{z}) d\mu_j d\theta^{(j)}. \tag{4.9}$$

The properties of conditional probabilities allow us to re-write (4.9) as

$$p(\mathbf{Z}^* \mid \mathbf{Z} = \mathbf{z}) = \sum_{j=1}^8 \iint p(\mathbf{Z}^* \mid \mu_j, \theta^{(j)}, \mathcal{P}_j, \mathbf{Z} = \mathbf{z}) \times p(\mu_j, \theta^{(j)} \mid \mathcal{P}_j, \mathbf{Z} = \mathbf{z}) \times p(\mathcal{P}_j \mid \mathbf{Z} = \mathbf{z}) d\mu_j d\theta^{(j)}. \tag{4.10}$$

This factorization of (4.9) greatly simplifies posterior prediction, as follows. Since we are using a Gaussian process model, $p(\mathbf{Z}^* | \mu_j, \theta^{(j)}, \mathcal{P}_j, \mathbf{Z} = \mathbf{z})$ is multivariate Gaussian: with

$$\begin{bmatrix} \mathbf{Z} \\ \mathbf{Z}^* \end{bmatrix} \Big| \mu_j, \theta^{(j)}, \mathcal{P}_j \sim N_{n+m} \left(\mu_j \mathbf{1}_{n+m}, \begin{bmatrix} \mathbf{V}_j + \mathbf{\Omega}_j & \mathbf{\Omega}_j^{ZZ^*} \\ \mathbf{\Omega}_j^{Z^*Z} & \mathbf{V}_j^* + \mathbf{\Omega}_j^* \end{bmatrix} \right),$$

where $\mathbf{\Omega}_j^{Z^*Z} \equiv \text{Cov}_j(\mathbf{Z}^*, \mathbf{Z})$, it follows that

$$p(\mathbf{Z}^* | \mu_j, \theta^{(j)}, \mathcal{P}_j, \mathbf{Z} = \mathbf{z}) = N_m(\mu_{\mathbf{Z}^* | \mathbf{z}}^{(j)}, \mathbf{\Sigma}_{\mathbf{Z}^* | \mathbf{z}}^{(j)}), \tag{4.11}$$

where

$$\mu_{\mathbf{Z}^* | \mathbf{z}}^{(j)} = \mu_j \mathbf{1}_m + \mathbf{\Omega}_j^{Z^*Z} (\mathbf{V}_j + \mathbf{\Omega}_j)^{-1} (\mathbf{z} - \mu_j \mathbf{1}_n),$$

$$\mathbf{\Sigma}_{\mathbf{Z}^* | \mathbf{z}}^{(j)} = (\mathbf{V}_j^* + \mathbf{\Omega}_j^*) - \mathbf{\Omega}_j^{Z^*Z} (\mathbf{V}_j + \mathbf{\Omega}_j)^{-1} \mathbf{\Omega}_j^{ZZ^*}.$$

The next component in (4.10) is the conditional posterior distribution (4.8), while the final component of (4.10) is the posterior probability of each partition \mathcal{P}_j conditional on the data. From Bayes' Theorem, the latter is

$$p(\mathcal{P}_j | \mathbf{Z} = \mathbf{z}) = \frac{p(\mathbf{Z} | \mathcal{P}_j) p(\mathcal{P}_j)}{\sum_{i=1}^8 p(\mathbf{Z} | \mathcal{P}_i) p(\mathcal{P}_i)} = \frac{p(\mathbf{Z} | \mathcal{P}_j)}{\sum_{i=1}^8 p(\mathbf{Z} | \mathcal{P}_i)}. \tag{4.12}$$

The last equality follows from our use of a uniform prior over the different partitions, that is, $p(\mathcal{P}_j) = 1/8, j = 1, \dots, 8$, which is appropriate in this setting because the partitions are defined by the covariates and not SOC. In any case, the important quantity here is the marginal likelihood

$$p(\mathbf{Z} | \mathcal{P}_j) = \iint p(\mathbf{Z} | \mu_j, \theta^{(j)}, \mathcal{P}_j) p(\mu_j, \theta^{(j)} | \mathcal{P}_j) d\mu_j d\theta^{(j)}. \tag{4.13}$$

Estimation of the marginal likelihood is a well-known problem in Bayesian analysis (for more details on our approach, see the Supplementary Material, Risser et al. (2019)).

Combining all of the above, (4.10) suggests the following algorithm to sample from $p(\mathbf{Z}^* | \mathbf{Z} = \mathbf{z})$:

1. Draw \mathcal{P}_j according to the $p(\mathcal{P}_j | \mathbf{Z} = \mathbf{z}), j = 1, \dots, 8$.

2. Draw $(\mu_j^*, \theta_*^{(j)})$ from $p(\mu_j, \theta^{(j)} \mid \mathcal{P} = \mathcal{P}_j, \mathbf{Z} = \mathbf{z})$, which is the posterior distribution for the mean and covariance parameters conditional on \mathcal{P}_j (the sampling of which is described in Section 4.2).
3. Draw \mathbf{Z}^* from $p(\mathbf{Z}^* \mid \mu_j = \mu_j^*, \theta^{(j)} = \theta_*^{(j)}, \mathcal{P} = \mathcal{P}_j, \mathbf{Z} = \mathbf{z})$ as defined in (4.11).

5. Results.

5.1. Posterior prediction maps.

Before showing the prediction maps, recall that the sampling algorithm requires estimates of the marginal likelihood. A variety of methods for estimating the marginal likelihoods (4.13), outlined in the Supplementary Material (Risser et al. (2019)), were applied to the eight partitions. The resulting log marginal likelihood estimates (scaled to have mean zero and variance one) and normalized probabilities (4.12) are shown in the Supplementary Material (Risser et al. (2019)). The scaled log marginal likelihood estimates show agreement between all of the estimators except BICM, which we henceforth exclude from consideration. Otherwise, on the probability scale, the estimators collectively give nonzero weight to the same three or four partitions (4, 6, 7, and 8) and none of the methods suggest that the model probabilities should be evenly distributed across all eight partitions. Given this relative agreement and for simplicity, we decided to use the HM estimator for generating posterior predictions of log SOC.

The resulting model-averaged posterior prediction maps and corresponding standard deviations for the Great Lakes region are shown in Figure 6. By averaging over the partitions, which individually consist of independent segments, we are able to obtain a relatively “smooth” posterior mean surface, although artifacts of the individual segments are visible in the posterior standard deviation map. While the standard deviation map is not smooth, we note that this nonsmoothness is what the SOC data select as the best model fits (based on the marginal likelihood estimates), and other marginal likelihood estimators give qualitatively similar results (see Figure 1 in Risser et al. (2019)). Regardless, the standard deviation map characterizes how the variability in predictions of log SOC varies over space; see Section 5.3 for further discussion. Comparing both the predicted (log) SOC surface obtained using our nonstationary Gaussian process model (Figure 6) as well as its standard deviation map to those obtained using other off-the-shelf methods (Figure 3) we can conclude that SOC is much more flexibly modeled with our proposed nonstationary model, as the prediction variances are better described by our model.

Additionally, another major benefit of using our nonstationary method with covariate-partitioning for the SOC data is the significant reduction in computational time relative to fitting a stationary Bayesian spatial model (as well as other related Bayesian models). Recall that the eight candidate partitions are generated ahead of time, and then, for each partition, the nonstationary models can be fit at the same time in parallel. For a dataset of size $n \approx 800$ observations, the maximum time required for an individual partition was 7 minutes (for 20,000 MCMC iterations), and the average time across all partitions was just 3.4 minutes (see Table 1). A traditional stationary model (e.g., SGP; see Section 2.2)

took approximately 20 minutes (computational times given are for an Intel Core i7 3.1 GHz machine with 16 GB memory). This major improvement in computational time is due to the fact that the likelihood conditional on a partition is the product of independent multivariate Gaussian likelihoods (see Section 4.1).

5.2. Model comparison.

Besides the qualitative comparison of the predictive SOC maps generated by our model-averaged, covariate-partitioned nonstationary model (henceforth NSGP) and the other off-the-shelf models, we also compare the predictive performance of the various models quantitatively. To this goal, we first use 10-fold cross validation in which the full data set ($n = 790$) is split into 10 equally sized test data sets with $m = 79$ observations. For each subset, we fit all models using the other nine subsets as training data and obtain predictions at the test locations. The continuous rank probability score (CRPS; Gneiting and Raftery (2007)) is used to evaluate the sharpness and calibration of predictions for the held-out data. Krüger et al. (2016) outline several methods for estimating the CRPS for individual predictions based on the output from MCMC algorithms (i.e., \widehat{CRPS}_t for $t = 1, \dots, m$). Mean scores are calculated for the test set using $\widehat{CRPS} = m^{-1} \sum_{t=1}^m \widehat{CRPS}_t$, and compared across models. Since the conditional predictive cumulative distribution functions (conditional on the true parameter state, i.e., the likelihood) are not available for BART and TGP, we use what Krüger et al. (2016) call the empirical CDF method (ECDF) for calculating CRPS for each model. ECDF is based on samples drawn from the posterior predictive distribution and is a consistent approximation for the CRPS for every predictive distribution with finite mean (Krüger et al. (2016)). The ECDF estimate of CRPS is calculated using the `scoringRules` package for R (Jordan, Krüger and Lerch (2016)).

For the 10-fold cross validation, the CRPS for each fold is the average of univariate CRPS across the test data. Given that our approach targets second-order properties of log SOC, we would also like to evaluate the models with respect to a multivariate criteria, for example, the energy score. However, in other work we have found it difficult to implement the energy score, because deviations in the energy score are challenging to diagnose. Instead, we evaluate the models' predictive capability for spatial *averages* of log SOC. We choose to work with spatial averages because while spatial averages are still univariate variables (and can thus be evaluated using the aforementioned methods), by being functionals of the joint predictive distribution, their second moment depends on the covariance structure of the (log) SOC process. Additionally, spatial averages are more compelling and useful quantities to examine from a soil management perspective. To this end, we fit all models to two additional groups of holdout sets, block and circular. Each block holdout set (12 total) creates a test data set out of a contiguous $\approx 3.3^\circ$ longitude by $\approx 1.7^\circ$ latitude box (the holdout sets range in size from 29 to 79 locations; see the Supplementary Material, Risser et al. (2019)). Each circular holdout set (10 total) creates a test data set using the 29 nearest neighbors of a randomly selected station (so that the holdout sets consist of $29 + 1 = 30$ stations; again see the Supplementary Material, Risser et al. (2019)). For each holdout set (block and circular), the other sets are used as training data to predict the spatial average corresponding to the test data locations. Then, we calculate the univariate CRPS for the spatial average, again using the `scoringRules` package as described previously. We note that the areal evaluation

explores the posterior predictive distribution for the average log SOC value for all available measurement locations in the holdout set, as opposed to the true areal average (which of course is not available for evaluation).

The CRPS for each model in Table 1 and each group of holdout sets are summarized in Table 2. Our covariate segmentation nonstationary model yields the best CRPS when averaging the univariate CRPS over locations in each fold, but the improvement relative to the other models is modest. The two nonstationary models (NSGP and TGP) perform well for the block holdout sets, although TGP maintains a clear advantage over NSGP. This is not completely surprising, since TGP (which uses rectangular partitions of the domain) is expected to perform well for rectangular holdout sets. When looking at the circular holdout sets, NSGP is preferred to TGP, which again is expected since NSGP uses nonrectangular partitions of the domain. BART outperforms both TGP and NSGP for the circular holdout sets, although recall that we previously decided to exclude BART from consideration based on its unrealistic prediction map.

One limitation of using CRPS to compare models is that it is difficult to evaluate the relative improvement for one model versus another. Clearly, NSGP yields only modest (at best) improvements in CRPS relative to the other fitted models; however, we note that other papers that focus on nonstationary modeling (e.g., Paciorek and Schervish (2006); Fuglstad et al. (2015)) also find very small improvements in out-of-sample evaluation criteria relative to stationary models, even when exploratory analyses indicate the presence of nonstationarity in the data. Fuglstad et al. (2015) note that, in their experience, “nonstationary models do not lead to much difference in the predicted values, [but] that the differences are found in the prediction variances”—this is also noted by Schmidt, Guttorp and O’Hagan (2011) and Vianna Neto, Schmidt and Guttorp (2014), and is certainly true for our application as well. As in Fuglstad et al. (2015), we would like to point out that whether or not we have improved the CRPS is not the only question worth asking: in this case, we argue that our more flexible characterization of the prediction variances yields increased insight into the spatial distribution of SOC (we expound upon this further in the next section).

5.3. Implications for decision making.

Gridded prediction maps and corresponding standard deviations such as those shown in Figures 3 and 6 are extremely important to soil scientists, and (as mentioned in Section 1) are used for a wide variety of purposes, including benchmarking mechanistic models of soil carbon (e.g., Todd-Brown et al. (2014)), identifying target areas for soil restoration projects (e.g., Ryals et al. (2014); Ryals et al. (2015)), and informing carbon sequestration projects that seek to determine the limits of soils’ carbon storing capacity (e.g., Angers et al. (2011); Wiesmeier et al. (2014)). For each of these purposes, it is very important to account for spatial variation in the uncertainty of the resulting predictions: as discussed in Section 5.2, note that the NSGP model can account for this (e.g., right panel of Figure 6) while the stationary (or approximately stationary) models cannot (e.g., bottom row of Figure 3). Thus, while the posterior mean predictions look quite similar (for example, NSGP, SGP, and PP) and the quantitative cross validation results do not yield overwhelming evidence in

support of the nonstationary model, the clear differences in standard deviations are a strong argument for using NSGP.

To emphasize this point, we plot the ratio of standard deviations in Figure 7, specifically showing the posterior standard deviations for SGP, PP, and TGP divided by the corresponding quantity for NSGP. In this plot, areas with a ratio of greater than one correspond to regions where NSGP provides much more precise predictions relative to the other models; areas with a ratio of less than one represent regions where posterior standard deviations from the other models are smaller. Using the results from the nonstationary model, a practitioner seeking to identify target areas for a new soil restoration project may want to focus their efforts in the Ohio River Valley (Kentucky, southern Indiana, and southern Illinois), as this is an area where both the soil carbon predictions are low and the uncertainty associated with these predictions is small. Alternatively, while soil carbon appears to be low along the Mississippi River in Wisconsin, the uncertainty in these measurements is much larger. Neither SGP nor PP provide this information.

6. Discussion.

In this paper, we have proposed a novel method for obtaining predictions of soil organic carbon at unobserved locations using a covariate-driven nonstationary spatial Gaussian process model. Our approach uses covariate partitioning to divide the geographic space according to the within-segment distribution of the covariates (here, land use-land cover and drainage class), so that the covariates indirectly inform the modeled first- and second-order properties of SOC as well as the resulting predictions. This approach extends existing covariate-driven nonstationary approaches in that the information a covariate provides on SOC can be used for prediction even if the covariate is not fully observed over the spatial domain. Bayesian model averaging accounts for uncertainty in the partitioning, and while our approach yields only modest improvements in terms of out-of-sample evaluation criteria, it provides a more appropriate characterization of the spatial variation in prediction uncertainty. Furthermore, while many nonstationary spatial approaches are computationally intensive, our approach results in computational times that are significantly faster than the time required to fit a corresponding second-order stationary model.

Of course, there are limitations to our nonstationary model. For example, the marginal likelihood estimation seems to be fairly sensitive to individual likelihood values, resulting in a few nonzero and many nearly zero posterior model probabilities. Regardless of which marginal likelihood method was used (see the Supplementary Material, Risser et al. (2019)), the estimates always put nearly all the weight on a single model. In a general setting, this may not be problematic; however, in our approach, we rely on model averaging to produce a scientifically meaningful (i.e., smooth) fitted surface, and we are well aware that the individual models, which specify independence across segments, are not on their own appropriate for modeling an environmental process like SOC. In light of the variability in the marginal likelihood estimates, we might also consider a second approach to model averaging, which could involve identifying a subset of the partitions that are in some sense “good” (as indicated by the data) and uniformly weighting over the reduced set. The

resulting surface would likely be more scientifically meaningful while also using the data to indicate which of the partitions provide a better fit to the observed SOC.

Finally, an additional extension that might address the nonsmooth mean predictions in Figure 6 would be to use different weighting functions for the partition-specific likelihood in (4.6). Recall that we use indicator weight functions; more generally, we might use weight functions without sharp boundaries. In this case, the transition between segments for a particular partition would be smoother, resulting in a more realistic specification for the individual models. In fact, we implemented such an approach, using the bivariate Gaussian densities from (3.1) as the unnormalized weights combined with the nearest neighbor Gaussian process likelihood (NNGP; Datta et al. (2016)) to speed up the computation. Unfortunately, such an approach did not improve the CRPS criteria relative to NSGP with indicator weights, and the computational times for fitting the individual partitions increased to as much as 2 hours (even with the NNGP likelihood; without NNGP the computational times were much longer, exceeding 10 hours for an individual partition). In general, the primary benefit of using indicator weight functions is computational: the product nature of (4.6) greatly reduces the computational time for fitting each model. The decision to use a different weight function should be balanced with the computational limitations of model fitting for a general data set.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments.

The authors would like to acknowledge Dr. Katerina Georgiou, a soil scientist at the Lawrence Berkeley National Laboratory, for helpful discussion regarding soil organic carbon, as well as the reviewers and Associate Editor for their feedback and suggestions for improvement.

REFERENCES

- ANGERS DA, ARROUAYS D, SABY NPA and WALTER C. (2011). Estimating and mapping the carbon saturation deficit of French agricultural topsoils. *Soil Use Manage.* 27 448–452.
- BANERJEE S, GELFAND AE, FINLEY AO and SANG H. (2008). Gaussian predictive process models for large spatial data sets. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* 70 825–848. MR2523906
- BATJES NH. (1996). Total carbon and nitrogen in the soils of the world. *Eur. J. Soil Sci.* 47 151–163. DOI:10.1111/j.1365-2389.1996.tb01386.x.
- BEAUDETTE DE and SKOVLIN JM. (2015). *soilDB*: Soil database interface. R package version 1.5–2.
- BLISS NB, WALTMAN SW, WEST LT, NEALE A and MEHAFFEY M. (2014). Distribution of soil organic carbon in the conterminous United States. In *Carbon Soil* (Hartemink AE and McSweeney K, eds.). *Progress in Soil Science* 85–93. Springer, Berlin.
- BONATO V, BALADANDAYUTHAPANI V, BROOM BM, SULMAN EP, ALDAPE KD and DO K-A. (2011). Bayesian ensemble methods for survival prediction in gene expression data. *Bioinformatics* 27 359–367. [PubMed: 21148161]
- CALDER CA. (2008). A dynamic process convolution approach to modeling ambient particulate matter concentrations. *Environmetrics* 19 39–48. MR2416543
- CHIPMAN HA, GEORGE EI and MCCULLOCH RE. (2010). BART: Bayesian additive regression trees. *Ann. Appl. Stat.* 4 266–298. MR2758172

- DATTA A, BANERJEE S, FINLEY AO and GELFAND AE. (2016). Hierarchical nearest-neighbor Gaussian process models for large geostatistical datasets. *J. Amer. Statist. Assoc.* 111 800–812. MR3538706
- DE VALPINE P, TUREK D, PACIOREK CJ, ANDERSON-BERGMAN C, TEMPLE LANG D and BODIK R. (2017). Programming with models: Writing statistical algorithms for general model structures with NIMBLE. *J. Comput. Graph. Statist.* 26 403–413. MR3640196
- DING J, BASHASHATI A, ROTH A, OLOUMI A, TSE K, ZENG T, HAFFARI G, HIRST M, MARRA MA, CONDON A, APARICIO S and SHAH SP. (2012). Feature-based classifiers for somatic mutation detection in tumour-normal paired sequencing data. *Bioinformatics* 28 167–175. [PubMed: 22084253]
- FINLEY AO, BANERJEE S and CARLIN BP. (2007). spBayes: An R package for univariate and multivariate hierarchical point-referenced spatial models. *J. Stat. Softw.* 19 1. [PubMed: 21494410]
- FINLEY AO, BANERJEE S and GELFAND AE. (2013). spBayes for large univariate and multivariate point-referenced spatio-temporal data models. Preprint. Available at arXiv:1310.8192.
- FUENTES M. (2001). A high frequency kriging approach for non-stationary environmental processes. *Environmetrics* 12 469–483.
- FUGLSTAD G-A, SIMPSON D, LINDGREN F and RUE H. (2015). Does non-stationary spatial data always require non-stationary random fields? *Spat. Stat.* 14 505–531. MR3431054
- GELMAN A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Anal.* 1 515–533. MR2221284
- GNEITING T and RAFTERY AE. (2007). Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102 359–378. MR2345548
- GOIDTS E and VAN WESEMAEL B. (2007). Regional assessment of soil organic carbon changes under agriculture in Southern Belgium, 1955–2005. *Geoderma* 141 341–354.
- GRAMACY RB. (2007). tgp: An R package for Bayesian nonstationary, semiparametric nonlinear regression and design by treed Gaussian process models. *J. Stat. Softw.* 19 1–46. [PubMed: 21494410]
- GRAMACY RB and LEE HKH. (2008). Bayesian treed Gaussian process models with an application to computer modeling. *J. Amer. Statist. Assoc.* 103 1119–1130. MR2528830
- GREEN DP and KERN HL. (2012). Modeling heterogeneous treatment effects in survey experiments with Bayesian additive regression trees. *Public Opin. Q.* 76 491–511.
- HOETING JA, MADIGAN D, RAFTERY AE and VOLINSKY CT. (1999). Bayesian model averaging: A tutorial. *Statist. Sci.* 14 382–417. MR1765176
- INGEBRIGTSEN R, LINDGREN F and STEINSLAND I. (2014). Spatial models with explanatory variables in the dependence structure. *Spat. Stat.* 8 20–38. MR3326819
- JOBBÁGY EG and JACKSON RB. (2000). The vertical distribution of soil organic carbon and its relationship to climate and vegetation. *Ecol. Appl.* 10.
- JORDAN A, KRÜGER F and LERCH S. (2016). scoringRules: Scoring rules for parametric and simulated distribution forecasts. R package version 0.9.
- KATZFUSS M. (2013). Bayesian nonstationary spatial modeling for very large datasets. *Environmetrics* 24 189–200. MR3067342
- KRÜGER F, LERCH S, THORARINSDOTTIR T and GNEITING T. (2016). Probabilistic forecasting and comparative model assessment based on Markov chain Monte Carlo output. Preprint. Available at arXiv:1608.06802.
- LEISCH F. (2004). FlexMix: A general framework for finite mixture models and latent class regression in R. *J. Stat. Softw.* 11 1–18. DOI:10.18637/jss.v011.i08.
- MINASNY B, MCBRATNEY AB, MENDONCA-SANTOS ML, ODEH IOA and GUYON B. (2006). Prediction and digital mapping of soil carbon storage in the Lower Namoi Valley. *Soil Res.* 44 233–244.
- MISHRA U, LAL R, SLATER B, CALHOUN F, LIU D and VAN MEIRVENNE M. (2009). Predicting soil organic carbon stock using profile depth distribution functions and ordinary kriging. *Soil Sci. Soc. Amer. J.* 73 614–621.

- MÜLLER P, QUINTANA F and ROSNER GL. (2011). A product partition model with regression on covariates. *J. Comput. Graph. Statist.* 20 260–278. MR2816548
- PACIOREK CJ and SCHERVISH MJ. (2006). Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics* 17 483–506. MR2240939 [PubMed: 18163157]
- POST WM and KWON KC. (2000). Soil carbon sequestration and land-use change: Processes and potential. *Glob. Change Biol.* 6 317–327.
- POST WM, EMANUEL WR, ZINKE PJ and STANGENBERGER A. (1982). Soil carbon pools and world life zones. *Nature* 298 156–159.
- REICH BJ, EIDSVIK J, GUINDANI M, NAIL AJ and SCHMIDT AM. (2011). A class of covariate-dependent spatiotemporal covariance functions for the analysis of daily ozone concentration. *Ann. Appl. Stat.* 5 2425–2447. MR2907121
- RISSER MD and CALDER CA. (2015). Regression-based covariance functions for nonstationary spatial modeling. *Environmetrics* 26 284–297. MR3340964
- RISSER MD, CALDER CA, BERROCAL VJ and BERRETT C. (2019). Supplement to “Nonstationary spatial prediction of soil organic carbon: Implications for stock assessment decision making.” DOI:10.1214/18-AOAS1204SUPP.
- RYALS R, KAISER M, TORN MS, BERHE AA and SILVER WL. (2014). Impacts of organic matter amendments on carbon and nitrogen dynamics in grassland soils. *Soil Biol. Biochem.* 68 52–61.
- RYALS R, HARTMAN MD, PARTON WJ, DELONGE MS and SILVER WL. (2015). Long-term climate change mitigation potential with organic matter management on grasslands. *Ecol. Appl.* 25 531–545. [PubMed: 26263673]
- SCHMIDT AM, GUTTORG P and O’HAGAN A. (2011). Considering covariates in the covariance structure of spatial processes. *Environmetrics* 22 487–500. MR2843403
- SIMBAHAN GC, DOBERMANN A, GOOVAERTS P, PING J and HADDIX ML. (2006). Fine-resolution mapping of soil organic carbon based on multivariate secondary data. *Geoderma* 132 471–489.
- SLEUTEL S, DE NEVE S, HOFMAN G, BOECKX P, BEHEYDT D, VAN CLEEMPUT O, MESTDAGH I, LOOTENS P, CARLIER L, VAN CAMP N, VERBEECK H, VANDE WALLE I, SAMSON R, LUST N and LEMEURE R. (2003). Carbon stock changes and carbon sequestration potential of Flemish cropland soils. *Glob. Change Biol.* 9 1193–1203.
- TODD-BROWN KEO, RANDERSON JT, HOPKINS F, ARORA V, HAJIMA T, JONES C, SHEVLIKOVA E, TJIPUTRA J, VOLODIN E, WU T, ZHANG Q and ALLISON SD. (2014). Changes in soil organic carbon storage predicted by Earth system models during the 21st century. *Biogeosciences* 11 2341–2356.
- VIANNA NETO JH, SCHMIDT AM and GUTTORG P. (2014). Accounting for spatially varying directional effects in spatial covariance structures. *J. R. Stat. Soc. Ser. C. Appl. Stat.* 63 103–122. MR3148271
- WIESMEIER M, HÜBNER R, SPÖRLEIN P, GEUSS U, HANGEN E, REISCHL A, SCHILLING B, LÜTZOW M and KÖGEL-KNABNER I. (2014). Carbon sequestration potential of soils in southeast Germany derived from stable soil organic carbon saturation. *Glob. Change Biol.* 20 653–665.
- WILLS S, SEQUEIRA C, LOECKE T, BENHAM E, FERGUSON R, SCHEFFE K and WEST L. (2013). Rapid Carbon Assessment (RaCA) Methodology Sampling and Initial Summary. United States Department of Agriculture, Natural Resources Conservation Service. Available online at http://www.nrcs.usda.gov/Internet/FSE_DOCUMENTS/nrcs142p2_052841.pdf.

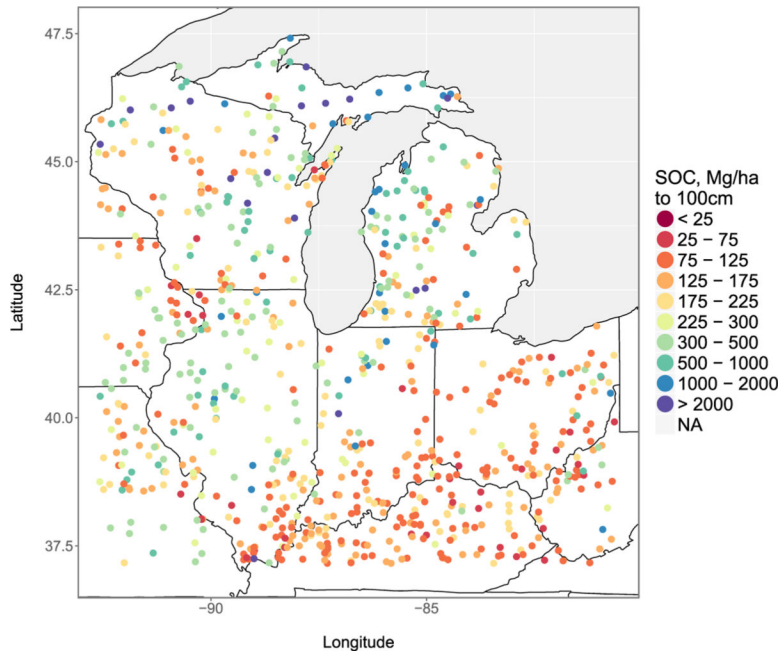




Fig. 1. Great Lakes subset of the soil organic carbon stocks data set.


(a) Entire region



(b) Fitted exponential semivariogram



(c) Arbitrary subregions



(d) Fitted exponential semivariograms

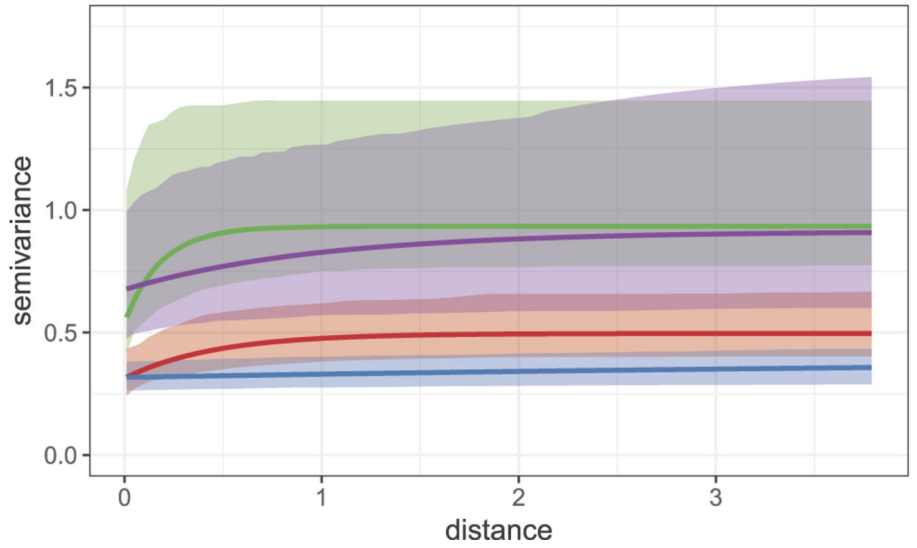


Fig. 2. Fitted exponential semivariograms for the entire Great Lakes region (top) and for four arbitrary subregions (bottom), with 95% confidence bands.

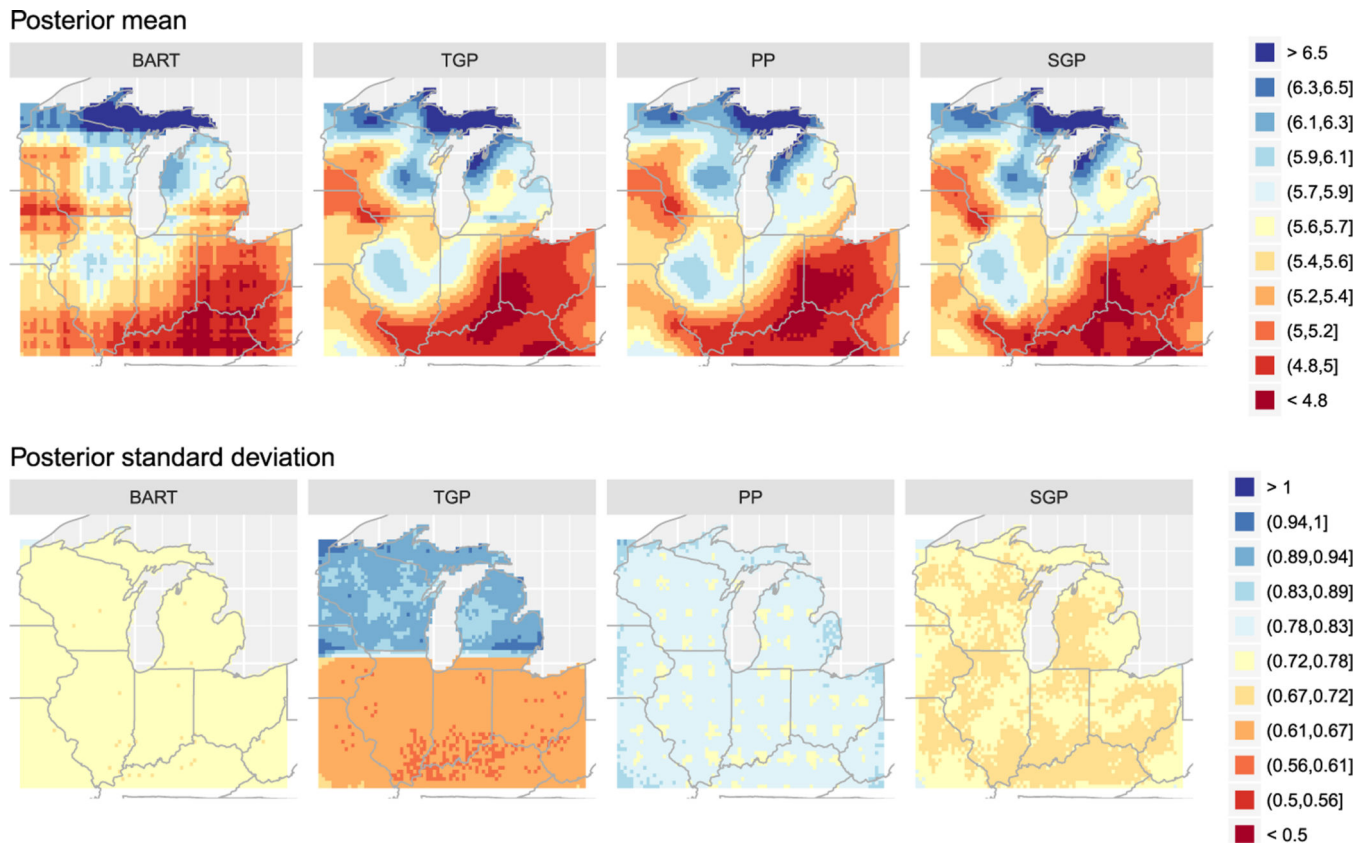


Fig. 3. Posterior mean predictions (top) and corresponding posterior standard deviations (bottom) for log SOC in the Great Lakes region, using the models BART (left), TGP (left-center), PP (right-center), and SGP (right). (Note: both color bars have the same limits as in Figure 6.)

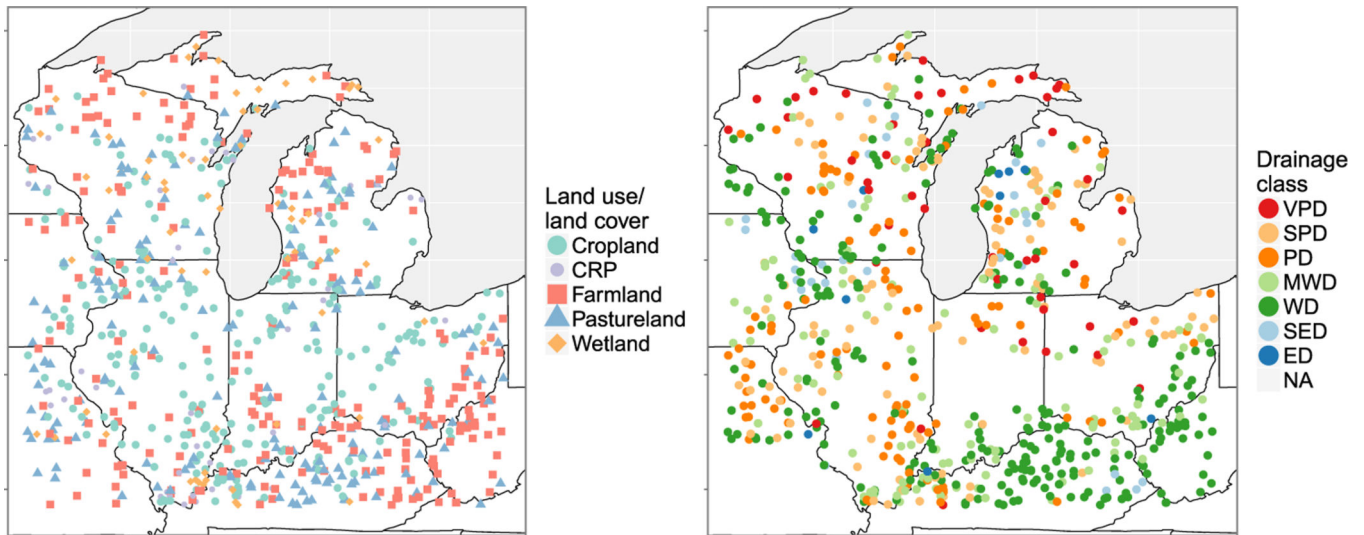


Fig. 4. The Great Lakes region subset of the Rapid Carbon Assessment (RaCA) land use-land cover classes (left) and drainage classes. For land use, CRP refers to a Conservation Reserve Program cropland site. For drainage class, the labels are as follows: VPD = very poorly drained, SPD = somewhat poorly drained, PD = poorly drained, MWD = moderately well drained, WD = well drained, SED = somewhat excessively drained, ED = excessively drained.

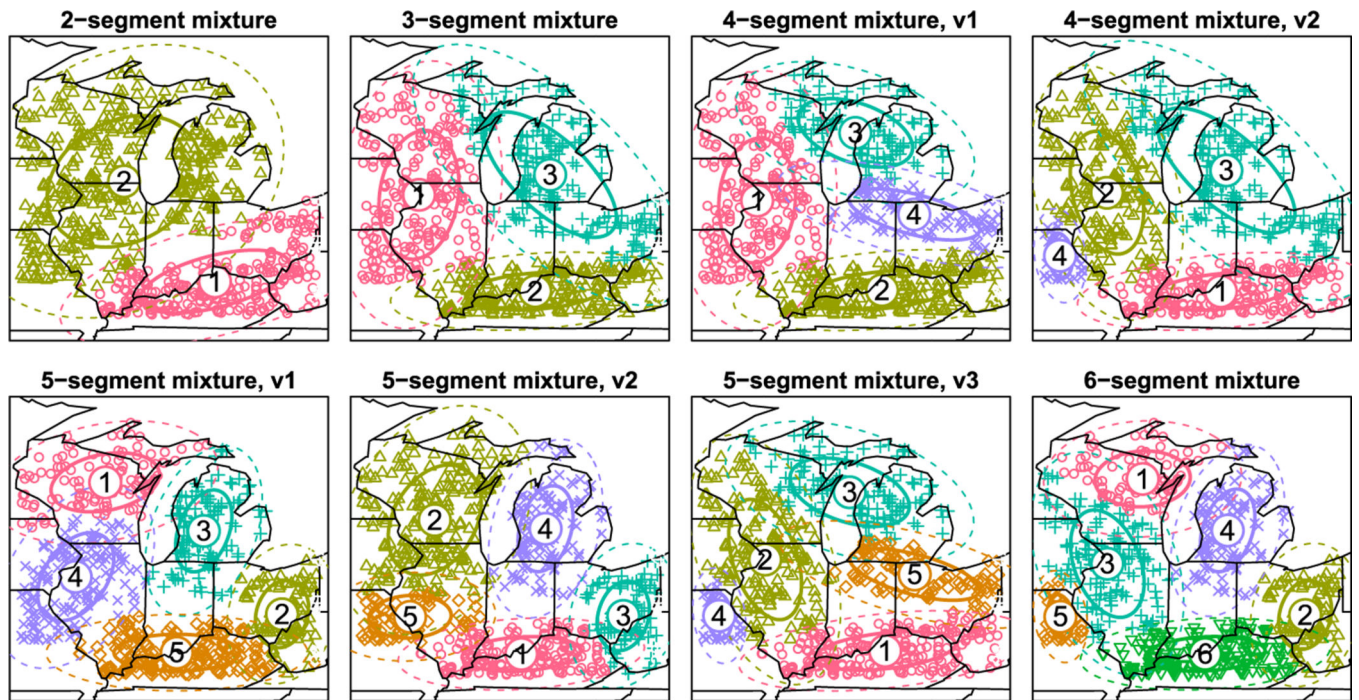


Fig. 5. Finite mixture models used to generate the candidate partitions. The plotted points represent locations with a nonmissing observation of both the drainage class and land use variables.

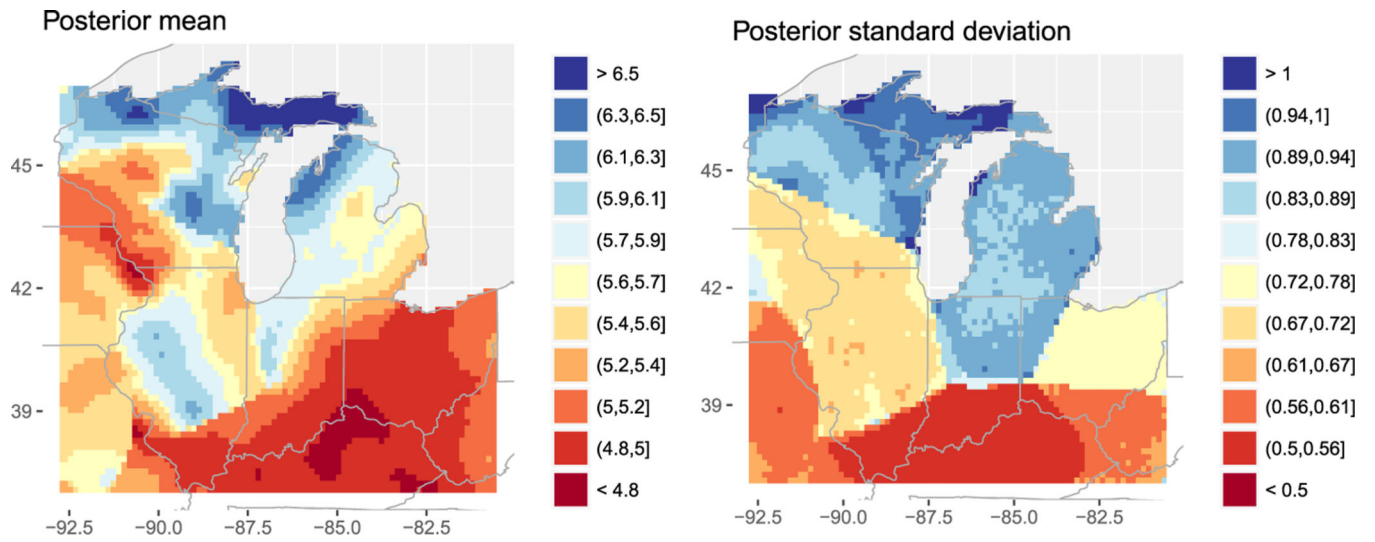


Fig. 6. Model-averaged posterior mean predictions (left) and corresponding posterior standard deviations (right) for log SOC (units = logMg/ha to 100 cm).

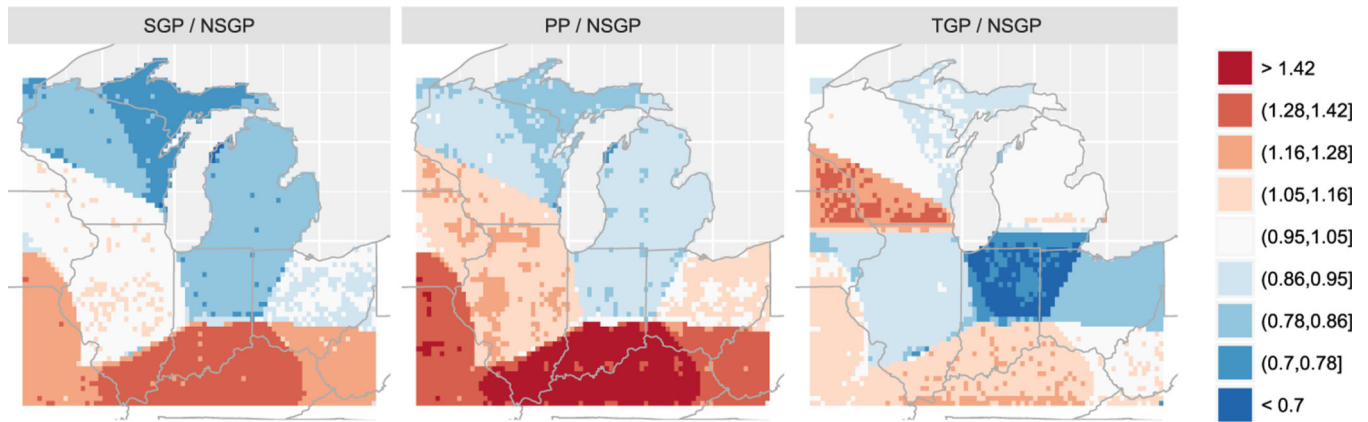


Fig. 7. Comparison of the posterior standard deviations across SGP, PP, TGP, and NSGP. The plotted colors represent the ratio of standard deviations for SGP (left), PP (center), and TGP (right) divided by the standard deviations for NSGP. Areas with a ratio of less than one correspond to regions where NS-GP provides more precise predictions relative to the other models; areas with a ratio of less than one represent regions where the other models likely underestimate the standard deviations.

Table 1

A summary of the models fit to the SOC data. (GP indicates a Gaussian process; lat/lon refers to latitude/longitude.) All spatial models use an exponential correlation function. Computational times correspond to fitting the model to the full data set (n= 790)

Label	Details	Mean function	R package	Computational time
BART [‡]	Bayesian additive regression trees	Lat/lon as inputs	BayesTree	12.4 minutes [‡]
TGP [‡]	Bayesian treed GP	Lat/lon for inputs and mean	tgp	13.1 minutes [‡]
PP [‡]	Bayesian Predictive Process ($r = 103$ knots)	Constant	spBayes	2.4 minutes [‡]
SGP [‡]	Bayesian stationary GP	Lat/lon	spBayes	19.4 minutes [‡]
NSGP	Bayesian nonstationary GP	Constant	n/a	7.0 (max), 3.4 (avg.) minutes [‡]

[‡]Time given for an Intel Core i7 3.1 GHz machine (16 GB memory) for 20,000 MCMC iterations.

[‡]Methods using available off-the-shelf software.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 2

Mean continuous rank probability scores (averaged over holdout sets) for each group of holdout sets, with the standard deviation in scores in parenthesis. The best model for each holdout set is in bold

Holdout	NSGP	SGP	BART [†]	TGP	PP
10-fold	0.391 (0.026)	0.396 (0.025)	0.407 (0.026)	0.399 (0.025)	0.398 (0.023)
Block	0.168 (0.142)	0.185 (0.140)	0.168 (0.089)	0.136 (0.125)	0.188 (0.162)
Circular	0.178 (0.157)	0.229 (0.192)	0.151 (0.068)	0.189 (0.155)	0.223 (0.201)

[†]BART is excluded from consideration based on its unrealistic prediction map.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript