

Spatial confounding under infill asymptotics

David Bolin¹ and Jonas Wallin²

¹Statistics Program, Computer, Electrical and Mathematical Sciences and Engineering Division,
King Abdullah University of Science and Technology (KAUST), Thuwal, Saudi Arabia

²Department of Statistics, Lund University, Lund, Sweden

Abstract

The estimation of regression parameters in spatially referenced data plays a crucial role across various scientific domains. A common approach involves employing an additive regression model to capture the relationship between observations and covariates, accounting for spatial variability not explained by the covariates through a Gaussian random field. While theoretical analyses of such models have predominantly focused on prediction and covariance parameter inference, recent attention has shifted towards understanding the theoretical properties of regression coefficient estimates, particularly in the context of spatial confounding. This article studies the effect of misspecified covariates, in particular when the misspecification changes the smoothness. We analyze the theoretical properties of the generalized least-square estimator under infill asymptotics, and show that the estimator can have counter-intuitive properties. In particular, the estimated regression coefficients can converge to zero as the number of observations increases, despite high correlations between observations and covariates. Perhaps even more surprising, the estimates can diverge to infinity under certain conditions. Through an application to temperature and precipitation data, we show that both behaviors can be observed for real data. Finally, we propose a simple fix to the problem by adding a smoothing step in the regression.

Keywords: spatial regression, generalized least-squares estimator, maximum likelihood, misspecification, estimation

1 Introduction

The estimation of regression parameters, sometimes also referred to as fixed effects, for spatially referenced data is of fundamental importance for practical applications in geostatistics, climate, disease mapping, brain imaging, and many other fields. A simple and commonly used model in these areas is the following additive regression model for observations Y_1, \dots, Y_n at spatial locations s_1, \dots, s_n in some domain D ,

$$Y_i = \sum_{k=1}^K X_k(s_i)\beta_k + \epsilon(s_i) = \mathbf{X}(s_i)\boldsymbol{\beta} + \epsilon(s_i), \quad (1)$$

where X_1, \dots, X_K are covariates, β_1, \dots, β_K regression coefficients, and ϵ is a centered Gaussian random field that is supposed to capture the spatial variation in the data which is not captured by the covariates.

The maximum likelihood estimate of the regression coefficients (assuming that the covariance function of ϵ is known) is given by the generalized least-squares estimator

$$\hat{\beta}_{GLS} = (\mathbf{X}^\top \boldsymbol{\Sigma}^{-1} \mathbf{X})^{-1} \mathbf{X}^\top \boldsymbol{\Sigma}^{-1} \mathbf{Y}, \quad (2)$$

where $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$, \mathbf{X} is an $n \times K$ matrix with entries $X_{ik} = X_k(s_i)$ and $\boldsymbol{\Sigma}$ is an $n \times n$ matrix with elements $\Sigma_{ij} = \text{Cov}(\epsilon(s_i), \epsilon(s_j))$. This estimator is typically preferred over the ordinary least squares estimator for $\boldsymbol{\beta}$ (obtained by replacing $\boldsymbol{\Sigma}$ with the identity matrix). Models of this type have mostly been analyzed theoretically with a focus on prediction or inference of covariance parameters (e.g., [Stein, 1999](#); [Zhang, 2004](#); [Ibragimov and Rozanov, 2012](#); [Kirchner and Bolin, 2022](#)). As noted by [Khan and Calder \(2022\)](#), there has until recently been few attempts to understand the theoretical properties of the regression coefficient estimates in spatial models like this. However, in the last few years several articles have investigated so-called spatial confounding problems (see [Khan and Berrett, 2023](#), for an overview of the literature). [Khan and Berrett \(2023\)](#) note that there is no single definition of spatial confounding, but that one can classify spatial confounding in two broad categories of model analysis confounding and data generation confounding depending on the assumptions that are made. In short, data generation confounding assumes that the model (1) is misspecified and what impact this has on (2). Pioneering work in this area is [Paciorek \(2010\)](#) which studied the case when ϵ and X are generated by a bivariate Gaussian random field. On the other hand, model analysis confounding does not assume a specific misspecification but rather studies the relation between $\boldsymbol{\Sigma}$ and \mathbf{X} and how this effects (2) and how its expected value deviates from the expected value of the OLS estimator. Early work in this area is [Reich et al. \(2006\)](#) which studied the case when $\boldsymbol{\Sigma}$ is generated by a conditionally autoregressive model for disease-mapping models.

Despite the recent interest in spatial confounding, there are some crucial problems with the spatial regression model which seem to have been widely overlooked in the literature, and this is the main focus of this work. Specifically, we are interested in the scenario when the smoothness of the observed process Y is different from an influential covariate X . To model this, we suppose that the true data generating process is

$$Y(s) = SX(s) + \epsilon(s),$$

where ϵ is a Gaussian process, $X(s)$ is some rough covariate, and S is some smoothing operator. Examples of situations like this, where the observed data are smoother than the covariate are not uncommon. One example, which we will study in Section 4 is standardized temperature and precipitation residuals over the midwest region of the United States (Figure 1). If we model the temperature residuals using the precipitation residuals as a covariate, we visually see that the observations are smoother than the covariate, and that there is a clear negative correlation between the observations and the covariate.

Because only $Y(s)$ and $X(s)$ are observed, suppose that we use $X(s)$ as a covariate for $Y(s)$ in a Gaussian process regression

$$Y(s) = X(s)\beta + \epsilon(s).$$

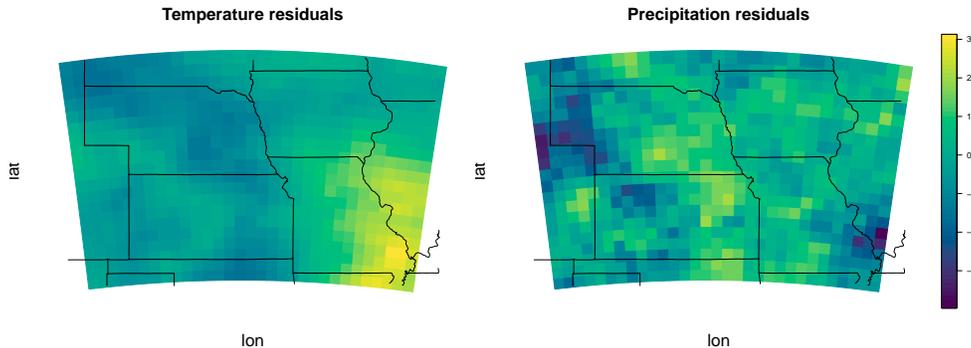


Figure 1: Standardized precipitation and temperature residuals.

In this case, we will show the somewhat counter intuitive result that $\beta \rightarrow 0$ as we get more data, no matter how high the correlation between Y and X is. Thus, using this model will severely underestimate the effect that $X(s)$ has on $Y(s)$. Not surprisingly, a simple solution to this issue is to perform smoothing of the covariate before using it in the spatial regression. However, we will show that one needs to do this with some care since, surprisingly, one can obtain that $|\beta| \rightarrow \infty$ as more data is obtained unless the covariate is smoothed enough.

In the next section, we formalize these problems and provide the main theoretical results, which hold for Gaussian process regression on general compact metric spaces. Thus, our theoretical results are much more general than the common setting of having observations in a bounded domain in \mathbb{R}^d , and are for example also applicable to data on the sphere or other manifolds. We also provide more detailed results for Gaussian random fields with fractional-order covariance operators, where Gaussian Whittle–Matérn fields (Lindgren et al., 2011, 2022) are important special cases, and extend some of the results to observations with measurement noise. Two simulation studies are performed in Section 3 to illustrate the results, and Section 4 contains the application to the temperature and precipitation residuals. We end the article with a discussion in Section 5, where we relate these results to other recent results for spatial confounding. All proofs are given in the appendix of the manuscript.

2 Theoretical setup and main results

In this section, we provide the theoretical setup and the main results. We begin with the most general results and then focus on Gaussian fields with fractional-order covariance operators as a special case, which contains several important examples for which we can obtain refined results.

2.1 General results

We let $(\mathcal{X}, d_{\mathcal{X}})$ denote a connected, compact metric space of infinite cardinality and let $\nu_{\mathcal{X}}$ be a strictly positive and finite Borel measure on $(\mathcal{X}, \mathcal{B}(\mathcal{X}))$, where $\mathcal{B}(\mathcal{X})$ denotes the Borel σ -algebra on \mathcal{X} . We further assume that ϵ is a square-integrable and centered Gaussian process $\epsilon: \mathcal{X} \times \Omega \rightarrow \mathbb{R}$ defined on a complete probability space $(\Omega, \mathcal{A}, \mathbb{P})$. The covariance

function of $\epsilon, \varrho: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, is assumed to be (strictly) positive definite and we denote the corresponding covariance operator by \mathcal{C} :

$$\mathcal{C}: L_2(\mathcal{X}, \nu_{\mathcal{X}}) \rightarrow L_2(\mathcal{X}, \nu_{\mathcal{X}}), \quad (\mathcal{C}w)(s) := \int_{\mathcal{X}} \varrho(s, s')w(s') d\nu_{\mathcal{X}}(s'), \quad (3)$$

Here, $L_2(\mathcal{X}, \nu_{\mathcal{X}})$ is the Hilbert space of real-valued, Borel measurable, square-integrable functions on \mathcal{X} , with inner product $(f, g)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} = \int_{\mathcal{X}} f(s)g(s)d\nu_{\mathcal{X}}(s)$ and induced norm $\|f\|_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}^2 = (f, f)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} = \int_{\mathcal{X}} |f(s)|^2 d\nu_{\mathcal{X}}(s)$.

The Gaussian process induces a Gaussian measure $\mu = \mathbf{N}(0, \mathcal{C})$ on the Hilbert space $L_2(\mathcal{X}, \nu_{\mathcal{X}})$, i.e., for every Borel set A we have $\mu(A) = \mathbb{P}(\{\omega \in \Omega : Y(\cdot, \omega) \in A\})$. The Cameron–Martin space $\mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ associated with μ on $L_2(\mathcal{X}, \nu_{\mathcal{X}})$, also known as the reproducing kernel Hilbert space of μ , is a Hilbert space which is defined as the range of $\mathcal{C}^{1/2}$ in $L_2(\mathcal{X}, \nu_{\mathcal{X}})$ equipped with the inner product $(\cdot, \cdot)_{\mathcal{C}} := (\mathcal{C}^{-1/2} \cdot, \mathcal{C}^{-1/2} \cdot)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$ and induced norm $\|\cdot\|_{\mathcal{C}} = \|\mathcal{C}^{-1} \cdot\|_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$. We similarly define $\mathcal{C}^1(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ as the range of \mathcal{C} in $L_2(\mathcal{X}, \nu_{\mathcal{X}})$ equipped with the inner product $(\cdot, \cdot)_{\mathcal{C}^2} := (\mathcal{C}^{-1} \cdot, \mathcal{C}^{-1} \cdot)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$ and corresponding induced norm.

We are now interested in the maximum-likelihood estimation of a model with a misspecified mean value. To make the results more generally applicable than to the setup mentioned in the introduction, we make the following assumption.

Assumption 1. *Suppose that $\{s_i\}_{i \in \mathbb{N}}$ is a sequence of locations in \mathcal{X} which is dense, i.e., each $s \in \mathcal{X}$ is an accumulation point of the sequence, and let $\{Y_i\}_{i \in \mathbb{N}}$ be a corresponding sequence of observations, $Y_i = Y(s_i)$, of the process $Y \sim \mu_{\beta}$, where $\mu_{\beta} = \mathbf{N}(m + \beta \mathcal{S}X, \mathcal{C})$. Here $\mathcal{S}: L_2(\mathcal{X}, \nu_{\mathcal{X}}) \rightarrow L_2(\mathcal{X}, \nu_{\mathcal{X}})$ is some operator, $m \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, and the covariate $X \neq 0$ satisfies $X \in L_2(\mathcal{X}, \nu_{\mathcal{X}})$, $\|X\|_{\infty} < \infty$ and $\|\mathcal{S}X\|_{\infty} < \infty$.*

Our main theorem is the following.

Theorem 1. *Suppose that the data is generated according to Assumption 1 while our candidate model is that $Y \sim \hat{\mu}_{\beta}$, where $\hat{\mu}_{\beta} = \mathbf{N}(\beta X, \mathcal{C})$. Let $\hat{\beta}_n$ denote the maximum likelihood estimator (under the candidate model) based on the n first observations, then:*

- (i) *if $\mathcal{S}X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $(\mathcal{S}X - X) \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, then $\hat{\beta}_n \rightarrow \beta$ \mathbb{P} -a.s.;*
- (ii) *if $\mathcal{S}X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $(\mathcal{S}X - X) \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ then $\hat{\beta}_n \rightarrow 0$ \mathbb{P} -a.s.;*
- (iii) *if $X \in \mathcal{C}^1(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, then $\hat{\beta}_n \rightarrow \beta_{\infty} = (Y, X)_{\mathcal{C}} / \|X\|_{\mathcal{C}}^2$ \mathbb{P} -a.s., where β_{∞} is a random variable with $\mathbf{E}(\beta_{\infty}) = (m + \beta \mathcal{S}X, X)_{\mathcal{C}} / \|X\|_{\mathcal{C}}^2 < \infty$.*

The first two cases of this result can be extended to the situation where measurement noise is included in the observation equation. That is, suppose that the observations are

$$\tilde{Y}_i = Y(s_i) + \epsilon_i, \quad (4)$$

where $\epsilon_i \sim \mathcal{N}(0, \sigma_{\epsilon}^2)$, and $Y(s_i)$ satisfy Assumption 1. We then have the following result.

Corollary 1. *Suppose that \tilde{Y}_i follow (4) while our candidate model is that $Y_i = Y(s_i) + \epsilon_i$, where $Y \sim \hat{\mu}_{\beta}$ with $\hat{\mu}_{\beta} = \mathbf{N}(\beta X, \mathcal{C})$. Let $\hat{\beta}_n$ denote the maximum likelihood estimator (under the candidate model) based on the n first observations, then:*

- (i) if $\mathcal{S}X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $(\mathcal{S}X - X) \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ then $\hat{\beta}_n \rightarrow \beta$ \mathbb{P} -a.s.;
- (ii) if $\mathcal{S}X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $(\mathcal{S}X - X) \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ then $\hat{\beta}_n \rightarrow 0$ \mathbb{P} -a.s..

To understand the implications of these results, let us explore a few important choices for the operator \mathcal{S} in the following corollaries.

Corollary 2. *If $X = \mathcal{S}X + X_0$ for an additive perturbation satisfying $X_0 \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, then $\hat{\beta}_n \rightarrow 0$.*

The following corollary provides the setup mentioned in the introduction.

Corollary 3. *If \mathcal{S} is a smoothing operator such that $\mathcal{S}X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, then $\hat{\beta}_n \rightarrow 0$ if $X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$.*

From these results, we see that we obtain $\hat{\beta}_n \rightarrow 0$ if X is not in the Cameron–Martin space of the covariance operator, whereas $\mathcal{S}X$ is. Thus, to avoid severe bias in these spatial regression problems, it is crucial to know whether or not the covariate, or the smoothed covariate, is in the Cameron–Martin space of the covariance operator. In the following subsection, we provide several examples of these spaces for Gaussian Whittle–Matérn fields.

2.2 Fractional-order covariance operators

In this section, we consider an application of the results of the previous section to Gaussian random fields with fractional order covariance operators, and provide some additional results not covered in the general case. For this, we need some additional notation. Suppose that $A : \mathcal{D}(A) \subsetneq L_2(\mathcal{X}, \nu_{\mathcal{X}}) \rightarrow L_2(\mathcal{X}, \nu_{\mathcal{X}})$ is a densely defined, self-adjoint and positive linear operator with a compact inverse (see [Bolin and Kirchner, 2023](#), for details). In this case, A has a set of eigenvectors $\{e_j\}_{j \in \mathbb{N}}$ and corresponding positive eigenvalues $\{\lambda_j\}_{j \in \mathbb{N}}$ which can be taken in a non-decreasing order and which only accumulate at ∞ . We suppose that there exists constants $c, C > 0$ and $\eta > 0$ such that

$$cj^\eta \leq \lambda_j \leq Cj^\eta, \quad \forall j \in \mathbb{N}. \quad (5)$$

For $\alpha > 0$, the action of the fractional-order operator $A^\alpha : \mathcal{D}(A^\alpha) \subsetneq L_2(\mathcal{X}, \nu_{\mathcal{X}}) \rightarrow L_2(\mathcal{X}, \nu_{\mathcal{X}})$ is defined in the spectral sense as

$$A^\alpha f := \sum_{j=1}^{\infty} \lambda_j^\alpha (f, e_j)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} e_j,$$

for $f \in \mathcal{D}(A^\alpha)$. For $s > 0$, the domain of A^s is defined as $\dot{H}_A^{2s} := \mathcal{D}(A^s)$, where

$$\mathcal{D}(A^s) = \{f \in L_2(\mathcal{X}, \nu_{\mathcal{X}}) : \|A^s f\|_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}^2 = \sum_{j=1}^{\infty} \lambda_j^{2s} (f, e_j)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}^2 < \infty\}.$$

We let \dot{H}_A^{-s} denote the dual space of \dot{H}_A^s and note that for $s > 0$, \dot{H}_A^s is a separable Hilbert space with inner product $(f, g)_{s,A} := (A^{s/2}f, A^{s/2}g)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$. Now, suppose that the covariance operator \mathcal{C} of the Gaussian process is $\mathcal{C} = A^{-\alpha}$ for some $\alpha > 0$, then the Cameron–Martin space is $\mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}})) = \dot{H}_A^\alpha$. This allows us to quantify the requirements regarding the Cameron–Martin space through \dot{H}_A^α , and we now provide some explicit examples of these for the popular Whittle–Matérn fields and generalized Whittle–Matérn fields on various spatial domains.

Example 1. Suppose that $\mathcal{X} = D \subset \mathbb{R}^d$ is connected, bounded and open domain with Lipschitz boundary. Choosing $A = \kappa^2 - \nabla \cdot (a\nabla)$ for a bounded function $\kappa : \bar{D} \rightarrow \mathbb{R}$ and a symmetric Lipschitz and uniformly positive definite function $a : \bar{D} \rightarrow \mathbb{R}^d$, and endowing the operator with homogeneous Dirichlet or Neumann boundary conditions, results in a generalized Whittle–Matérn field (see, e.g. [Bolin and Kirchner, 2020](#)). With a as the identity function and κ as a positive constant, we obtain the standard Whittle–Matérn fields as a special case. For these fields, one can characterize the Cameron–Martin spaces in terms of standard fractional Sobolev spaces $H^\alpha(\mathcal{X})$, see, e.g., [Bolin and Kirchner \(2023\)](#). By these results, $f \in H^\alpha(D)$ is required to have $f \in \dot{H}_A^\alpha$, and we therefore have an explicit smoothness requirement on the functions in the Cameron–Martin space.

Example 2. Suppose that $\mathcal{X} = M$ is a closed, connected, orientable, smooth and compact surface in \mathbb{R}^3 . Choosing $A = \kappa^2 - \nabla_M \cdot (a\nabla_M)\Delta_{\mathcal{X}}$, where $\nabla_M \cdot$ is the surface divergence and ∇_M the surface gradient, results in a generalized Whittle–Matérn field on the surface. The characterization of the Cameron–Martin space can then be done in terms of standard fractional Sobolev spaces $H^\alpha(M)$ ([Herrmann et al., 2020](#)). These results can also be extended to more general Riemannian manifolds, and in particular are applicable to Gaussian Whittle–Matérn fields on manifolds by considering $A = \kappa^2 - \Delta_M$, where $\kappa > 0$ is constant and Δ_M is the Laplace–Beltrami operator ([Borovitskiy et al., 2020](#); [Harbrecht et al., 2021](#)).

Example 3. Suppose that $\mathcal{X} = \Gamma$ is a compact metric graph, such as a linear network. Choosing $A = \kappa^2 - \Delta_\Gamma$, where $\kappa > 0$ and Δ_Γ is the Kirchhoff–Laplacian results in a Whittle–Matérn field on Γ for which [Bolin et al. \(2023\)](#) provide a representation of the Cameron–Martin spaces. [Bolin et al. \(2023\)](#) extends this to generalized Whittle–Matérn fields on Γ , obtained by considering $A = \kappa^2 - d(a d)$, where κ is a bounded function, a is a positive Lipschitz function, and d is an operator acting as the derivative on the edges of Γ .

We assume that the operator \mathcal{S} is of the form $\mathcal{S} = A_S^\gamma$, where A_S is an operator that diagonalizes with respect to the eigenvectors $\{e_j\}_{j \in \mathbb{N}}$, with corresponding positive eigenvalues $\{\lambda_{S,j}\}$ satisfying the same asymptotic growth as A :

$$c_s j^\eta \leq \lambda_{S,j} \leq C_s j^\eta, \quad \forall j \in \mathbb{N},$$

for some constants $c_s, C_s > 0$. Then, from [Theorem 1](#) it follows that:

Corollary 4. Suppose that the data is generated according to [Assumption 1](#), where $\mathcal{C} = A^{-\alpha}$ for $\alpha > 0$ and that $\mathcal{S} = A_S^\gamma$ for $\gamma \in \mathbb{R}$. Further suppose that $X \in \dot{H}_A^p$, for some $p > 0$, and that $X \notin \dot{H}_A^{\tilde{p}}$ for any $\tilde{p} > p$. The candidate model is $Y \sim \hat{\mu}_\beta$, where $\hat{\mu}_\beta = \mathbf{N}(\beta X, \mathcal{C})$. Let $\hat{\beta}_n$ denote the maximum likelihood estimator based on the n first observations, then:

$$\hat{\beta}_n \rightarrow \begin{cases} \beta & \text{if } p < \alpha \text{ and } \gamma = 0, \\ 0 & \text{if } p < \alpha \text{ and } 2\gamma \leq p - \alpha, \\ \beta_\infty = (Y, X)_\mathcal{C} / \|X\|_\mathcal{C}^2 & \text{if } p \geq 2\alpha \text{ and } \gamma < 0, \end{cases} \quad (6)$$

\mathbb{P} -a.s., where β_∞ is a random variable with $\mathbf{E}(\beta_\infty) = \beta(\mathcal{S}X, X)_\mathcal{C} / \|X\|_\mathcal{C}^2 < \infty$.

We can note that some combinations of p, α and γ are missing in the result. It is more complicated to get a general result in these cases as the behavior appear to depend on how the observations are taken. One situation where we can obtain a complete characterization is if the observations are eigenbasis observations, which is a common assumption is several theoretical investigations of maximum likelihood estimators, (see, e.g., [Cialenco, 2018](#)).

Proposition 1. *Suppose that $\mathcal{C} = A^{-\alpha}$ for $\alpha > 0$ and that $\mathcal{S} = A_S^\gamma$ for $\gamma \in \mathbb{R}$. Further suppose that $X \in \dot{H}_A^p$ for $p > 0$, and that $X \notin \dot{H}_A^{\tilde{p}}$ for $\tilde{p} > p$. The candidate model is $Y \sim \hat{\mu}_\beta$, where $\hat{\mu}_\beta = \mathbf{N}(\beta X, \mathcal{C})$. Let $\hat{\beta}_n$ denote the maximum likelihood estimator based on the n observations $Y_i = (Y, e_i)_{L_2(\mathcal{X}, \nu_X)}$ for $i = 1, \dots, n$ then*

$$\hat{\beta}_n \rightarrow \begin{cases} \beta & \text{if } p < \alpha \text{ and } \gamma = 0, \\ 0 & \text{if } p < \alpha \text{ and } \gamma < 0, \\ \text{sign}(\beta)\infty & \text{if } (p < \alpha \text{ and } \gamma > 0) \text{ or } (p \geq \alpha \text{ and } \gamma > p - \alpha), \\ \beta_\infty = (Y, X)_\mathcal{C} / \|X\|_\mathcal{C}^2 & \text{if } p \geq \alpha \text{ and } \gamma \leq p - \alpha, \end{cases}$$

\mathbb{P} -a.s., where β_∞ is finite with $\mathbf{E}(\beta_\infty) = \beta(\mathcal{S}X, X)_\mathcal{C} / \|X\|_\mathcal{C}^2$.

This result should carry over to any reasonable fill in limit sequence of observations such that $\mathbf{X}^\top \Sigma^{-1} \mathbf{Y}$ and $\mathbf{X}^\top \Sigma^{-1} \mathbf{X}$ in (2) converges to $(Y, X)_\mathcal{C} < \infty$ and $\|X\|_\mathcal{C}^2 < \infty$ as $n \rightarrow \infty$. Thus, we conjecture that for most practical situations, the results obtained for eigenbasis observations hold also for the point observations, in particular that $|\hat{\beta}_n| \rightarrow \infty$ if $p < \alpha$ and $\gamma > 0$ or if $p \geq \alpha$ and $\gamma > p - \alpha$. This is also what we observe in Section 3.2 for point observations.

3 Simulation experiments

To illustrate the results of the previous section, we here perform a few simulation experiments. In the first experiment we take a time series approach and use a lowess smoother (Cleveland, 1979) to illustrate the effect of the results. In the second experiment, we consider a setup similar to the spatial data application in Section 4.

3.1 A time series example

Suppose that we are given data $Y_i = \mathcal{S}X(s_i) + Z(s_i)$ where $s_i \in [0, 10]$, and $Z(s)$ is a centered Gaussian process with a Matérn covariance with $\kappa = \nu = 1$ and $\sigma = 0.1$. We generate $\mathcal{S}X$ by first simulating a realization $X(s)$ of a centered Gaussian process with a Matérn covariance with $\kappa = \nu = 1$ and $\sigma = 0.4$, and then smoothing it with a lowess smoother with smoother span 0.1. An example of the process and covariates can be seen in Figure 2.

As a first experiment, we assume that we know the covariance function of Z and estimate two models

$$\text{Model 1 : } Y_i = \beta X(s_i) + Z(s_i), \quad \text{Model 2 : } Y_i = \beta \hat{\mathcal{S}}X(s_i) + Z(s_i),$$

where the first model uses the unmodified, rough, covariate X , and the second model uses $\hat{\mathcal{S}}X$ which is a smoothed version of X , obtained by a different lowess smoother with smoother span 0.2.

For both models, we estimate β using maximum likelihood using an increasing amount of observations, regularly spaced in the interval $[0, 1]$. The resulting estimates can be seen in the columns for Experiment 1 in Table 1, where the values are obtained as averages of the results for 10 repetitions of the experiment.

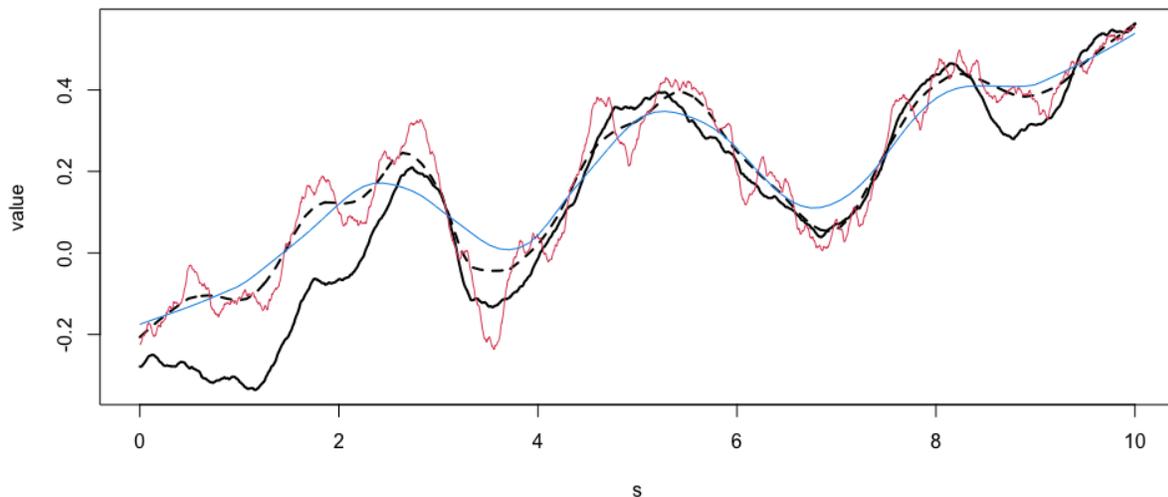


Figure 2: Simulation of $Y(s) = \mathcal{S}X(s) + Z(s)$ (black solid), $\mathcal{S}X(s)$ (black dashed), the non-smoothed covariate $X(s)$ (red), and the smoothed covariate $\hat{\mathcal{S}}X(s)$ (blue).

n	Experiment 1		Experiment 2		Experiment 3	
	Model 1	Model 2	Model 1	Model 2	Model 1	Model 2
128	0.108	1.059	0.027	1.043	0.026	1.006
256	0.052	1.049	0.018	1.081	0.042	1.023
512	0.026	1.046	0.011	1.094	0.034	1.04
1024	0.014	1.048	0.009	1.097	0.034	1.039

Table 1: Estimates of the regression coefficient β for the three simulation experiments. Model 1 has a rough covariate whereas Model 2 has a smooth covariate. Each number is an average over 10 different datasets.

As a second experiment, we redo the same simulations, but estimate the parameters κ, σ, ν of the covariance function jointly with β when performing the maximum likelihood estimation. The results are shown in the columns for Experiment 2 in Table 1. Finally, as a third experiment, we extend the model for Z by adding a nugget effect. That is, we use a covariance function $r(h) = \rho(h) + \sigma_e^2 1(h=0)$, where σ_e^2 is the nugget representing the variance of measurement error. We generate the data Y_i assuming this covariance for Z , where we choose $\sigma_e = 0.01$, and estimate all parameters $\kappa, \sigma, \nu, \sigma_e$ of the covariance function jointly with β when performing the maximum likelihood estimation. The results of this experiment are shown in the columns for Experiment 3 in Table 1.

In all three experiments, the estimate of β tends to zero as the number of observations increases for the first model, whereas the estimate is stable for the second model, despite that it uses the wrong covariate. This result is expected from the results of the previous section, but might seem counterintuitive given that X has a high correlation with the observed data (it was, for example, 0.951 on average for the first experiment).

3.2 A spatial example

We now illustrate the results in an example related to the application in the next section. We assume that ϵ in (1) is a Gaussian Matérn field with range $\kappa = 0.4$, variance $\sigma^2 = 1.3^2$, and smoothness $\nu = 2$ on the domain in Figure 1. These parameter values were chosen so that they are similar to parameters for the application. We further assume that X also is a Gaussian Matérn field, independent of ϵ , with the same parameters κ and σ , and a different smoothness parameter ν_X . Finally, we assume that \mathcal{S} is chosen so that $\mathcal{S}X$ also is a Gaussian Matérn field with the same parameters κ and σ and a different smoothness parameter $\nu_{\mathcal{S}X}$.

We consider 620 observation locations as in the Figure 1 for the application. If we let \mathbf{Y} denote the observations at these locations, the model specifically is

$$\mathbf{Y} = \mathcal{S}\mathbf{X} + \Sigma_\epsilon^{1/2}\mathbf{Z}_1,$$

where $\mathcal{S}\mathbf{X} = \Sigma_{\mathcal{S}X}^{1/2}\mathbf{Z}_2$ and $\mathbf{X} = \Sigma_X^{1/2}\mathbf{Z}_2$. Here, \mathbf{Z}_i are vectors with independent standard Gaussian variables and $\Sigma_\epsilon, \Sigma_X$ and $\Sigma_{\mathcal{S}X}$ denote the covariance matrices corresponding to Matérn covariance functions with parameters $\kappa = 0.4$, $\sigma = 1.3$ and smoothness $\nu = 2, \nu_X$, and $\nu_{\mathcal{S}X}$ respectively. For $\nu_{\mathcal{S}X}$ we consider the two cases $\nu_{\mathcal{S}X} = \nu$ and $\nu_{\mathcal{S}X} = \nu - 0.5$. In the first, the smoothness of ϵ and $\mathcal{S}\mathbf{X}$ are equal which is likely the situation one is in when estimating the smoothness from data. In the second case, $\mathcal{S}\mathbf{X}$ is rougher than ϵ , which could occur if the smoothness of ϵ is kept fixed (e.g., if an exponential or a squared exponential covariance is used).

We subsample \mathbf{Y} , \mathbf{X} and $\mathcal{S}\mathbf{X}$ by sampling n out of the 620 observation locations uniformly at random and compute $\hat{\beta}_n$ using (2) where n is the number of observations in the sample. These estimates are computed for different values of $\nu_X \in [0.5, 10]$.

In the notation of Section 2.2, the smoothness parameters correspond to $\alpha = \nu + 1$, $p = \nu_X$ and $\gamma = (\nu_X - \nu_{\mathcal{S}X})/2$. Thus, if Proposition 1 also would hold for point observations, we expect that

$$\text{case 1: } \hat{\beta}_n \rightarrow \begin{cases} 0 & \text{if } \nu_X < 2, \\ \beta & \text{if } \nu_X = 2, \\ \infty & \text{if } 2 < \nu_X < 3, \\ \beta_\infty & \text{if } \nu_X \geq 3, \end{cases} \quad \text{case 2: } \hat{\beta}_n \rightarrow \begin{cases} 0 & \text{if } \nu_X < 1.5, \\ \beta & \text{if } \nu_X = 1.5, \\ \infty & \text{if } 1.5 < \nu_X < 3, \\ \beta_\infty & \text{if } \nu_X \geq 3, \end{cases}$$

with $\mathbf{E}(\beta_\infty) = \beta(\mathcal{S}X, X)_c / \|X\|_c < \infty$.

We repeat the experiment 400 times, and the result is seen in Figure 3. One can observe the result from Proposition 1 emerging as the sample size grows. In particular when $\nu_X < 2$ for the first case, and $\nu_X < 1.5$ for the second, the estimator goes to zero. When $\nu_X \in (2, 3)$ for the first case, and $\nu_X \in (1.5, 3)$ in the second, there is a positive bias. As the sample size is not that large in this example, it is not clear if this bias will be unbounded if n was increased further, but the bias is clearly large enough that it could have serious implications for the interpretation of the results.

4 An application to reanalysis data

To illustrate the theory on real data we take a data set of precipitation and temperature reanalysis. The variables are the average summer (June, July and August) temperature,

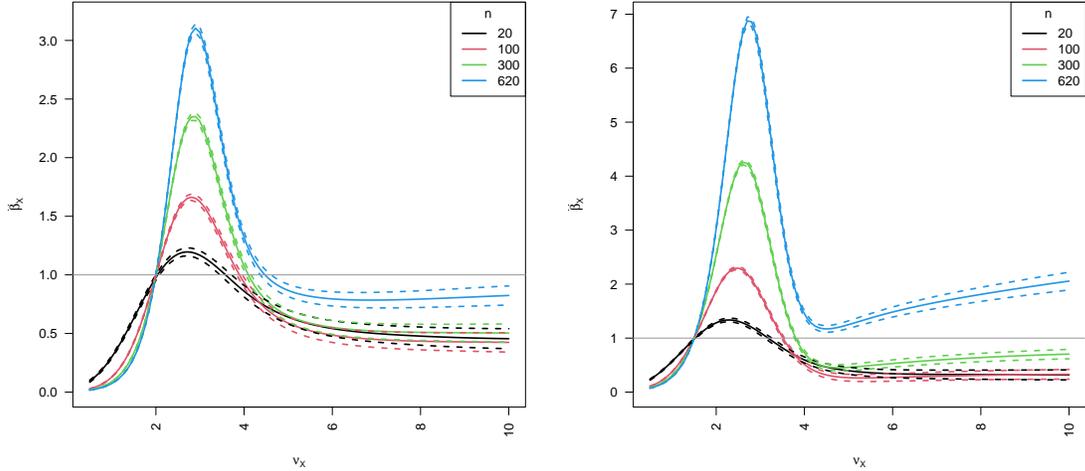


Figure 3: The effect of smoothing the covariate for the cases $\nu_{SX} = \nu$ (left) and $\nu_{SX} = \nu - 0.5$ (right). The solid line is the mean effect of 400 simulations, and the dashed lines represent Monte Carlo 95% confidence bands.

$n(space)$	Temperature model		Precipitation model	
	$\beta_P(CI95\%)$	RMSE	$\beta_T(CI95\%)$	RMSE
15	-0.21 (-0.26, -0.16)	0.19	-0.55 (-0.67, -0.44)	0.41
50	-0.18 (-0.20, -0.16)	0.18	-0.73 (-0.81, -0.65)	0.42
100	-0.15 (-0.16, -0.14)	0.17	-0.80 (-0.86, -0.73)	0.42
250	-0.11 (-0.12, -0.11)	0.17	-1.07 (-1.13, -1.01)	0.43
620	-0.09 (-0.09, -0.09)	0.17	-1.57 (-1.60, -1.50)	0.46

Table 2: Results for modelling temperature with precipitation as covariate and precipitation with temperature as covariate. The table shows the estimated regression coefficients, their 95% confidence intervals, and the root mean square prediction error for the prediction set.

$T(s, t)$, and the cube-root precipitation, $P(s, t)$, over the Midwest United states for 24 years (1981-2004). To simplify interpretation we standardize the data for each year. The data for one year is shown in Figure 1. The empirical correlation between the two variables is -0.67 . This data was studied earlier in [Genton and Kleiber \(2015\)](#) for examining the joint distribution through various multivariate models, and they in particular studied a linear coregionalization model,

$$\begin{aligned}
 T(s, t) &= a_{11}Z_1(s, t), \\
 P(s, t) &= a_{12}Z_1(s, t) + a_{22}Z_2(s, t),
 \end{aligned}$$

where Z_1, Z_2 are independent Gaussian random fields. The formulation is motivated by “We opt for this formulation since temperature is expected to be smoother than precipitation,...”. Therefore the data seems ideal to study in our context as differentiability of the field determines the Cameron–Martin space of the Matérn covariance.

In order to study how the estimators behave as functions of sample size, we subsample n of the 620 spatial locations, keeping the data from all 24 years for the subsampled locations.

$n(space)$	Temperature model		Precipitation model	
	$\beta_{SP}(CI95\%)$	RMSE	$\beta_{ST}(CI95\%)$	RMSE
15	-0.32 (-0.44, -0.19)	0.17	-0.37 (-0.50, -0.25)	0.42
50	-0.19 (-0.33, -0.06)	0.17	-0.48 (-0.59, -0.38)	0.42
100	-0.19 (-0.33, -0.04)	0.17	-0.55 (-0.65, -0.46)	0.42
250	-0.31 (-0.46, -0.17)	0.17	-0.58 (-0.68, -0.49)	0.42
620	-0.36 (-0.53, -0.19)	0.17	-0.58 (-0.67, -0.49)	0.42

Table 3: Results for modelling temperature with smoothed precipitation as covariate and precipitation with smoothed temperature as covariate. The table shows the estimated regression coefficients, their 95% confidence intervals, and the root mean square prediction error for the prediction set.

We explore both predicting temperature using precipitation as a covariate,

$$T(s, t) = \beta_0 + \beta_P P(s, t) + Z_T(s, t; \theta_T),$$

and predicting precipitation using temperature as a covariate,

$$P(s, t) = \beta_0 + \beta_T T(s, t) + Z_P(s, t; \theta_P).$$

Here Z_T and Z_P are fitted using a Matérn covariance with parameters estimated from the subsampled data. We estimate the regression coefficients and covariance parameters based on the subsampled dataset. Then, after fitting the parameters we split the data set into a prediction set of 36 locations, which is predicted given the remaining 584 observations. The result is shown in Table 2. The coefficients change as the theory suggests, and one can note the absurd implication of using this model for predicting precipitation using temperature: As the data is standardized, the yearly variance of both pressure and temperature is one, while the predictor $\beta_T T(s, t)$ has a variance of 1.57^2 . Also note that the prediction error is slightly decreasing with n for temperature, even though the coefficient goes to zero, this is inline with theory as the unsmoothed precipitation cannot contribute to the interpolation.

Finally, we fit the same models but using smoothed covariates. These smoothed covariates are obtained by first estimating the covariance of the covariates, i.e., we estimate the models $P(s, t) = \beta_0 + Z_P(s, t; \theta_P^{fixed})$ and $T(s, t) = \beta_0 + Z_T(s, t; \theta_T^{fixed})$. Then we multiply the observed precipitation \mathbf{P} and temperature \mathbf{T} by some powers of the fitted covariance matrices, yielding $\mathcal{S}\mathbf{P} = \Sigma_{\theta_P^{fixed}}^3 \mathbf{P}$ and $\mathcal{S}\mathbf{T} = \Sigma_{\theta_T^{fixed}}^{0.5} \mathbf{T}$. These choices are likely not optimal, but the results are sufficient to ensure that the covariates are smooth enough as can be seen in the stable estimates in Table 3. Note also that there is no large difference in prediction error compared to the results in Table 2, which means that the problem cannot be identified through cross validation.

5 Discussion

We have investigated the implications of misspecifying the smoothness of covariates in regression models within an infill asymptotic framework, and in particular showed the importance of ensuring that the covariates do not exhibit rougher characteristics than the

underlying field. If the covariates are too rough, the corresponding regression coefficients will converge to zero no matter how high the correlation is with the observed data. This main result is formally established in Theorem 1.

On the other hand, if a misspecified covariate possess roughness equal to or slightly smoother than the field, the corresponding regression coefficient may diverge to infinity under certain scenarios. We were unable to establish this result solely based on the assumption of observations filling the domain; indeed, we suspect that the result need not be true under such weak conditions. Instead, we demonstrated the behavior utilizing observations of eigenvectors, as detailed in Proposition 1. However, both simulated data from random locations (see Section 3.2) and real-world data (see Section 4) indicate that this behavior holds for point observations at random locations, at least for stationary Gaussian random fields with Matérn covariances.

We have not encountered any previous works studying misspecified covariates under infill asymptotics. There are, however, related works worth commenting on. For example in Gilbert et al. (2023), the regression coefficient is defined as identifiable if it is less smooth than ϵ , underscoring the significance of smoothness. While their definition of identifiability aligns with the ability to perfectly separate noise from regression coefficients, we prefer to denote this as a lack of consistent estimability, to avoid confusion with the classical definition of identifiability (refer to Definition 5.2 in Lehmann and Casella (2006)). Another pertinent work is Dupont et al. (2023), which conducts a non-asymptotic analysis of spatial confounding effects and highlights the impact of smoothing. On page 8, they state, “the expression in Corollary 3.2 shows firstly that bias in the spatial model is directly linked to the smoothing in the model, as without smoothing the bias would not arise.” Finally, a recent method proposed by Dupont et al. (2022) suggests removing the smooth component of X using a thin plate spline. This recommendation appears contrary to our findings. However, the authors clarify in their discussion that their methods are designed for covariates with non-spatial components, implying a measurement error component within X . Consequently, their approach would be suitable in scenarios where $\mathcal{S}X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ in our notation.

An intriguing avenue for future research is to explore general assumptions on the observations under which the regression coefficients tend to infinity under misspecified smoothness. Additionally, investigating whether these properties persist when substituting the Gaussian field ϵ in (1) with a more general Lévy random field is an interesting topic for future research.

Acknowledgments

The authors thanks Peter Diggle and Claudio Fronterre for helpful discussions and suggestions of relevant articles to study. Further the authors would like to thank Alexandre Simas for discussion of the results and help with certain proofs.

A Collected proofs

Throughout the proofs we will use representations of Gaussian fields in terms of Gaussian white noise \mathcal{W} on $L_2(\mathcal{X}, \nu_{\mathcal{X}})$, which can be represented as a family of centered Gaussian

random variables $\{\mathcal{W}(h) : h \in L_2(\mathcal{X}, \nu_{\mathcal{X}})\}$ such that

$$\forall f, g \in L_2(\mathcal{X}, \nu_{\mathcal{X}}), \quad \mathbb{E}(\mathcal{W}(f)\mathcal{W}(g)) = (f, g)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}.$$

A Gaussian random field ϵ with covariance operator \mathcal{C} can then be represented as $\epsilon = \mathcal{C}^{1/2}\mathcal{W}$.

A.1 General results

Proof of Theorem 1. Let f_n^β denote the density of Y_1, \dots, Y_n under μ_β , and let $\hat{f}_n(\beta)$ denote the corresponding density under $\hat{\mu}_\beta$. By the Feldman–Hajek theorem (Da Prato and Zabczyk, 2014, Theorem 2.25), we have that μ_β and $\hat{\mu}_{\hat{\beta}}$ are equivalent if

$$m + \beta \mathcal{S}X - \hat{\beta}X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}})), \quad (7)$$

and they are orthogonal otherwise.

We start by proving (i). Since $\mathcal{S}X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $\mathcal{S}X - X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, we have that $X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and (7) therefore shows that $\hat{\mu}_{\hat{\beta}}$ is equivalent to μ_β if $\beta = \hat{\beta}$ and that the measures are orthogonal whenever $\hat{\beta} \neq \beta$. Because of this, we by Gikhman and Skorokhod (2004, Theorem 1, p.442) and Da Prato and Zabczyk (2014, Proposition 2.26) have that

$$\lim_{n \rightarrow \infty} \log \frac{f_n^\beta}{\hat{f}_n(\hat{\beta})} = \begin{cases} -\infty & \hat{\beta} \neq \beta, \\ \log(c_0) & \hat{\beta} = \beta, \end{cases} \quad (8)$$

with probability 1, where $c_0 > 0$ is the Radon–Nikodym derivative

$$\frac{d\mu_\beta}{d\hat{\mu}_\beta}(y) = \exp \left((y - m - \beta \mathcal{S}X, \beta (\mathcal{S}X - X) + m)_\mathcal{C} - \frac{1}{2} \|\beta (\mathcal{S}X - X) + m\|_\mathcal{C}^2 \right), \quad (9)$$

for $\hat{\beta} = \beta$. When $\hat{\beta} = \beta$, $\|m + \beta \mathcal{S}X - \hat{\beta}X\|_\mathcal{C}$ is finite because $\mathcal{S}X - X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $m \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$. Further, for $\hat{\beta} = \beta$ we have that $m + \beta \mathcal{S}X - \hat{\beta}X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ so that we can write $m + \beta \mathcal{S}X - \hat{\beta}X = \mathcal{C}^{1/2}f$ for $f \in L_2(\mathcal{X}, \nu_{\mathcal{X}})$. Using this and that $y = m + \beta \mathcal{S}X + \epsilon$, where $\epsilon = \mathcal{C}^{1/2}\mathcal{W}$, we obtain that

$$\begin{aligned} (y - \hat{\beta}X, m + \beta \mathcal{S}X - \hat{\beta}X)_\mathcal{C} &= (\epsilon, m + \beta \mathcal{S}X - \hat{\beta}X)_\mathcal{C} + \|m + \beta \mathcal{S}X - \hat{\beta}X\|_\mathcal{C} \\ &= (\mathcal{C}^{1/2}\mathcal{W}, \mathcal{C}^{-1}(m + \beta \mathcal{S}X - \hat{\beta}X))_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} + \|m + \beta \mathcal{S}X - \hat{\beta}X\|_\mathcal{C} \\ &= (\mathcal{C}^{1/2}\mathcal{W}, \mathcal{C}^{-1/2}f)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} + \|m + \beta \mathcal{S}X - \hat{\beta}X\|_\mathcal{C} \\ &= \mathcal{W}(f) + \|m + \beta \mathcal{S}X - \hat{\beta}X\|_\mathcal{C} < \infty \quad \mathbb{P}\text{-a.s.} \end{aligned}$$

Thus, $c_0 < \infty$ \mathbb{P} -a.s. To prove the claim we need show that $\hat{\beta}_n$ is not a divergent sequence. It is enough to show that for a fixed $\epsilon > 0$, there exists an N such that $\rho_n(\hat{\beta}) = \log f_n - \log \hat{f}_n(\hat{\beta}) < \log(c_0) - 1$ for all $|\beta - \hat{\beta}| > \epsilon$ whenever $n > N$. To that end, fix a β^* with $|\beta^* - \beta| > \epsilon$ and note that by (8), there is an N such that $\rho_n(\beta^*) < \log(c_0) - 1$. Next, it is clear that $\hat{\beta} \mapsto \log \hat{f}_n(\hat{\beta})$ is strictly concave, and we therefore have that $\rho_n(\hat{\beta}) < \log(c_0) - 1$ for all $\hat{\beta}$ with $|\beta - \hat{\beta}| > |\beta^* - \hat{\beta}|$. This finishes the proof since ϵ was arbitrary.

We now prove (ii). Since $\mathcal{S}X \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ and $(\mathcal{S}X - X) \notin \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$, (7) shows that $\hat{\mu}_{\hat{\beta}}$ is orthogonal to μ_{β} whenever $\hat{\beta} \neq 0$. We have that

$$\lim_{n \rightarrow \infty} \log \frac{f_n^\beta}{\hat{f}_n(\hat{\beta})} = \begin{cases} -\infty & \hat{\beta} \neq 0, \\ \log(c_0) & \hat{\beta} = 0, \end{cases} \quad (10)$$

with probability 1, where $c_0 > 0$ is the Radon–Nikodym derivative (9), which is finite because $\hat{\beta} = 0$. The remaining proof is now identical to that of (i).

To prove (iii) note that $\lim_{n \rightarrow \infty} \log \frac{f_n^\beta}{\hat{f}_n(\hat{\beta})}$ will be equal (with probability 1) to the Radon–Nikodym derivative (9). The maximum likelihood estimator of $\hat{\beta}$ is the maximizer of this expression, which is given by

$$\beta^* = \arg \min_{\hat{\beta}} \left\{ (Y, \hat{\beta}X)_C - \frac{1}{2} \|\hat{\beta}X\|_C^2 \right\} = \frac{(Y, X)_C}{\|X\|_C^2}.$$

Now will show that this a random variable with finite variance and expectation. Since $X \in \mathcal{C}^{1/2}(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ there exists an $\tilde{X} \in L_2(\mathcal{X}, \nu_{\mathcal{X}})$ such that $X = \mathcal{C}^{1/2}\tilde{X}$. Furthermore, Y can be expressed as $m + \beta\mathcal{S}X + \mathcal{C}^{1/2}\mathcal{W}$ where \mathcal{W} is Gaussian white noise. Thus,

$$(Y, X)_C = (m + \beta\mathcal{S}X, X)_C + (\mathcal{C}^{1/2}\mathcal{W}, X)_C. \quad (11)$$

Here $(\mathcal{C}^{1/2}\mathcal{W}, X)_C = (\mathcal{C}^{1/2}\mathcal{W}, \mathcal{C}^{-1/2}\tilde{X})_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} = \mathcal{W}(\tilde{X}) < \infty$, \mathbb{P} -a.s., since $\tilde{X} \in L_2(\mathcal{X}, \nu_{\mathcal{X}})$ and \mathcal{W} is Gaussian white noise on $L_2(\mathcal{X}, \nu_{\mathcal{X}})$. Finally, since $X \in \mathcal{C}^1(L_2(\mathcal{X}, \nu_{\mathcal{X}}))$ there exists $\bar{X} \in L_2(\mathcal{X}, \nu_{\mathcal{X}})$ such that $X = \mathcal{C}\bar{X}$ and $(m + \beta\mathcal{S}X, X)_C = (m + \beta\mathcal{S}X, \bar{X})_{L_2(\mathcal{X}, \nu_{\mathcal{X}})} < \infty$. \square

Proof of Corollary 1. This is an immediate consequence of the proof of Theorem 1 and Stein (1999, Theorem 6 (Chapter 4)). \square

A.2 Fractional-order covariance operators

Proof of Corollary 4. Let $\dot{H}_S^{2\gamma}$ denote the domain of A_S^γ and note that we have the equivalence $\dot{H}_A^0 \cong L_2(\mathcal{X}, \nu_{\mathcal{X}}) \cong H_S^0$. By Bolin et al. (2020, Lemma 2.1), A_S^γ can be extended to an isometric isomorphism $A_S^\gamma : \dot{H}_S^s \rightarrow \dot{H}_S^{s-2\gamma}$ for $s \in \mathbb{R}$. Further, because $\lambda_{S,j}$ and λ_j have the same asymptotic growth, we have that $\dot{H}_A^s \cong \dot{H}_S^s$ for $s \in \mathbb{R}$. Thus, we have that $\mathcal{S}X \in \dot{H}_A^{p-2\gamma}$.

Using the same notation as in the proof of Theorem 1, we have that μ_{β} and $\hat{\mu}_{\hat{\beta}}$ are equivalent if and only if $\beta\mathcal{S}X - \hat{\beta}X \in \dot{H}_A^\alpha$. If $\gamma = 0$ and $p < \alpha$, we have that $\mathcal{S}X \notin \dot{H}_A^\alpha$ and $\mathcal{S}X - X = 0 \in \dot{H}_A^\alpha$. Therefore, the result follows from Theorem 1 (i). If $p < \alpha$ and $2\gamma < p - \alpha$, we have that $\mathcal{S}X \in \dot{H}_A^{p-2\gamma} \subseteq \dot{H}_A^\alpha$ and $\mathcal{S}X - X \notin \dot{H}_A^\alpha$. Therefore, the result follows from Theorem 1 (ii) as γ is negative so that $\mathcal{S} : L_2(\mathcal{X}, \nu_{\mathcal{X}}) \rightarrow L_2(\mathcal{X}, \nu_{\mathcal{X}})$. Finally, if $p \geq 2\alpha$ and $\gamma \leq 0$, we have that $X \in \dot{H}_A^{2\alpha}$ and the result follows from Theorem 1 (iii). \square

Proof of Proposition 1. When the observations are eigenfunction observations, the likelihood is given by $l(\beta) = -\frac{1}{2} \sum_{i=1}^n \lambda_i^\alpha (Y_i - X_i\beta)^2$, where $X_i = (X, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$, and thus

$$\hat{\beta}_n = \frac{\sum_{i=1}^n \lambda_i^\alpha Y_i X_i}{\sum_{i=1}^n \lambda_i^\alpha X_i^2}.$$

First note that $Y_i = (\mathcal{S}X + \epsilon, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}$, where ϵ can be written as $\epsilon = \mathcal{C}^{1/2}\mathcal{W}$ for Gaussian white noise \mathcal{W} on $L_2(\mathcal{X}, \nu_{\mathcal{X}})$. Thus,

$$\hat{\beta}_n = \beta \frac{\sum_{i=1}^n \lambda_i^\alpha X_i (\mathcal{S}X, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} + \frac{\sum_{i=1}^n \lambda_i^\alpha X_i (\mathcal{C}^{1/2}\mathcal{W}, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}}{\sum_{i=1}^n \lambda_i^\alpha X_i^2}. \quad (12)$$

The second term is a random variable with

$$\mathbb{V} \left[\frac{\sum_{i=1}^n \lambda_i^\alpha X_i (\mathcal{C}^{1/2}\mathcal{W}, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} \right] = \mathbb{V} \left[\frac{\sum_{i=1}^n \lambda_i^{\alpha/2} X_i \xi_i}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} \right] = \frac{\sum_{i=1}^n \lambda_i^\alpha X_i^2}{(\sum_{i=1}^n \lambda_i^\alpha X_i^2)^2} = B_n.$$

Consider first the case $p < \alpha$: then, $\sum_{i=1}^\infty \lambda_i^\alpha X_i^2 = \infty$ and hence $B_n \rightarrow 0$. Thus, the asymptotic value of $\hat{\beta}_n$ depends solely on the first term in (12):

$$\beta \frac{\sum_{i=1}^n \lambda_i^\alpha X_i (\mathcal{S}X, e_i)_{L_2(\mathcal{X}, \nu_{\mathcal{X}})}}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} = \beta \frac{\sum_{i=1}^n \lambda_i^\alpha \lambda_{S,i}^\gamma X_i^2}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} := \beta A_n.$$

If $\gamma = 0$, we clearly have that $A_n \rightarrow 1$, so $\hat{\beta} \rightarrow \beta$. Further,

$$A_n \leq \frac{\lambda_{S,n}^\gamma \sum_{i=1}^n \lambda_i^\alpha X_i^2}{\sum_{i=1}^n \lambda_i^\alpha X_i^2} = \lambda_{S,n}^\gamma.$$

Thus, $A_n \rightarrow 0$ if $\gamma < 0$, so in this case $\hat{\beta}_n \rightarrow 0$. Now, take $m < n$, we then have that

$$A_n \geq \frac{\sum_{i=1}^m \lambda_i^\alpha \lambda_{S,i}^\gamma X_i^2 + \lambda_{S,m}^\gamma \sum_{i=m}^n \lambda_i^\alpha X_i^2}{\sum_{i=1}^m \lambda_i^\alpha X_i^2 + \sum_{i=m}^n \lambda_i^\alpha X_i^2} = \frac{\frac{\sum_{i=1}^m \lambda_i^\alpha \lambda_{S,i}^\gamma X_i^2}{\sum_{i=m}^n \lambda_i^\alpha X_i^2} + \lambda_{S,m}^\gamma}{\frac{\sum_{i=1}^m \lambda_i^\alpha X_i^2}{\sum_{i=m}^n \lambda_i^\alpha X_i^2} + 1} \rightarrow \lambda_{S,m}^\gamma$$

as $n \rightarrow \infty$ because all X_i are bounded and $\sum_{i=1}^\infty \lambda_i^\alpha X_i^2 = \infty$. Since this holds for any $m < n$, we can conclude that A_n diverges if $\gamma > 0$. Now consider the case $p \geq \alpha$. In this case, $\|X\|_{\mathcal{C}}^2 = \sum_{i=1}^\infty \lambda_i^\alpha X_i^2 < \infty$. Further,

$$\sum_{i=1}^n \lambda_i^\alpha \lambda_{S,i}^\gamma X_i^2 \leq C_s \sum_{i=1}^n \lambda_i^{\alpha+\gamma} X_i^2 \rightarrow \|X\|_{A, \alpha+\gamma}^2 < \infty$$

if $\alpha + \gamma \leq p$. Thus, in this case, $\hat{\beta}_n \rightarrow (Y, X)_{\mathcal{C}} / \|X\|_{\mathcal{C}}^2 < \infty$ \mathbb{P} -a.s.. If, on the other hand, $\alpha + \gamma > p$, $\sum_{i=1}^n \lambda_i^\alpha \lambda_{S,i}^\gamma X_i^2 \rightarrow \infty$ and thus, $\hat{\beta}_n \rightarrow \text{sign}(\beta)\infty$. \square

References

- Bolin, D. and K. Kirchner (2020). The rational SPDE approach for Gaussian random fields with general smoothness. *J. Comp. Graph. Stat.* 29(2), 274–285.
- Bolin, D. and K. Kirchner (2023). Equivalence of measures and asymptotically optimal linear prediction for Gaussian random fields with fractional-order covariance operators. *Bernoulli* 29, 1476–1504.

- Bolin, D., K. Kirchner, and M. Kovács (2020). Numerical solution of fractional elliptic stochastic PDEs with spatial white noise. *IMA J. Numer. Anal.* 40(2), 1051–1073.
- Bolin, D., M. Kovács, V. Kumar, and A. B. Simas (2023). Regularity and numerical approximation of fractional elliptic differential equations on compact metric graphs. *Math. Comp.*. In press.
- Bolin, D., A. B. Simas, and J. Wallin (2023). Gaussian Whittle-Matérn fields on metric graphs. *Bernoulli*. In press.
- Borovitskiy, V., A. Terenin, P. Mostowsky, and M. Deisenroth (2020). Matérn Gaussian processes on Riemannian manifolds. In H. Larochelle, M. Ranzato, R. Hadsell, M. Balcan, and H. Lin (Eds.), *Advances in Neural Information Processing Systems*, Volume 33, pp. 12426–12437. Curran Associates, Inc.
- Cialenco, I. (2018). Statistical inference for SPDEs: an overview. *Statistical Inference for Stochastic Processes* 21(2), 309–329.
- Cleveland, W. S. (1979). Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association* 74(368), 829–836.
- Da Prato, G. and J. Zabczyk (2014). *Stochastic Equations in Infinite Dimensions* (Second ed.), Volume 152 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge.
- Dupont, E., I. Marques, and T. Kneib (2023). Demystifying spatial confounding. *arXiv preprint arXiv:2309.16861*.
- Dupont, E., S. N. Wood, and N. H. Augustin (2022). Spatial+: a novel approach to spatial confounding. *Biometrics* 78(4), 1279–1290.
- Genton, M. G. and W. Kleiber (2015). Cross-covariance functions for multivariate geostatistics. *Statistical Science* 30(2), 147 – 163.
- Gikhman, I. I. and A. V. Skorokhod (2004). *The theory of stochastic processes. I*. Classics in Mathematics. Springer-Verlag, Berlin. Translated from the Russian by S. Kotz, Reprint of the 1974 edition.
- Gilbert, B., E. L. Ogburn, and A. Datta (2023). Consistency of common spatial estimators under spatial confounding. *arXiv preprint arXiv:2308.12181*.
- Harbrecht, H., L. Herrmann, K. Kirchner, and C. Schwab (2021). Multilevel approximation of Gaussian random fields: Covariance compression, estimation and spatial prediction. *arXiv preprint arXiv:2103.04424*.
- Herrmann, L., K. Kirchner, and C. Schwab (2020). Multilevel approximation of Gaussian random fields: fast simulation. *Math. Models Methods Appl. Sci.* 30(1), 181–223.
- Ibragimov, I. A. and Y. A. Rozanov (2012). *Gaussian random processes*, Volume 9. Springer Science & Business Media.

- Khan, K. and C. Berrett (2023). Re-thinking spatial confounding in spatial linear mixed models. *arXiv preprint arXiv:2301.05743*.
- Khan, K. and C. A. Calder (2022). Restricted spatial regression methods: Implications for inference. *Journal of the American Statistical Association* 117(537), 482–494.
- Kirchner, K. and D. Bolin (2022). Necessary and sufficient conditions for asymptotically optimal linear prediction of random fields on compact metric spaces. *The Annals of Statistics* 50(2), 1038–1065.
- Lehmann, E. L. and G. Casella (2006). *Theory of point estimation*. Springer Science & Business Media.
- Lindgren, F., D. Bolin, and H. Rue (2022). The SPDE approach for Gaussian and non-Gaussian fields: 10 years and still running. *Spat. Stat.* 50, Paper No. 100599.
- Lindgren, F., H. Rue, and J. Lindström (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 73(4), 423–498. With discussion and a reply by the authors.
- Paciorek, C. J. (2010). The importance of scale for spatial-confounding bias and precision of spatial regression estimators. *Statistical science: a review journal of the Institute of Mathematical Statistics* 25(1), 107.
- Reich, B. J., J. S. Hodges, and V. Zadnik (2006). Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models. *Biometrics* 62(4), 1197–1206.
- Stein, M. L. (1999). *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics. Springer-Verlag, New York.
- Zhang, H. (2004). Inconsistent estimation and asymptotically equal interpolations in model-based geostatistics. *J. Amer. Statist. Assoc.* 99(465), 250–261.