RESEARCH ARTICLE

Methods in Ecology and Evolution | BRITISH ECOLOGICAL SOCIETY

Spatial confounding in joint species distribution models

Francis K. C. Hui[1] [iD] | Quan Vu[1] | Mevin B. Hooten[2]

[1]Research School of Finance, Actuarial Studies & Statistics, The Australian National University, Canberra, Australian Capital Territory, Australia

[2]Department of Statistics and Data Sciences, The University of Texas at Austin, Austin, Texas, USA

Correspondence
Francis K. C. Hui
Email: fhui28@gmail.com

Abstract

1. Joint species distribution models (JSDMs) are a popular method for analysing multivariate abundance data, with important applications such as uncovering how species communities are driven by environmental processes, model-based ordination to visualise community composition patterns across sites and variance partitioning to quantify the relative contributions of different processes in shaping a species community.

2. One issue that has received relatively little attention in the study of joint species distributions is that of spatial confounding: when one or more of the environmental predictors exhibit spatial correlation, and spatially structured random effects such as spatial factors are also included in the model, then these two components may be collinear with each other.

3. Through a combination of simulations and case studies, we show that if not managed properly, spatial confounding can result in misleading inference on covariate effects in a spatially structured JSDM, along with difficulties in interpreting ordination results and incorrect attribution of variation to environmental processes in a species community.

4. We present one approach to treat spatial confounding called restricted spatial factor analysis, which is designed to ensure that the covariate effects retain their full explanatory power, and ordinations constructed using the spatial factors explain species covariation beyond that accounted for by the measured predictors. We encourage ecologists to consider the inferences they seek to make from spatially structured JSDMs and to ensure that the covariate effects and ordinations they estimate and interpret are aligned with their scientific questions of interest.

KEYWORDS
collinearity, community ecology, covariate effects, latent variables models, ordination, spatio-temporal modelling

1 | INTRODUCTION

Joint species distribution models (JSDMs, Pollock et al., 2014) have emerged as a powerful statistical framework for analysing multivariate abundance data. Some of their primary applications include learning how species communities are driven by environmental and habitat predictors, model-based ordination to visualise patterns in community composition and quantifying the relative

contributions of environmental versus biotic effects in shaping species assemblages (Ovaskainen & Abrego, 2020; Schliep et al., 2018; Thorson et al., 2016; van der Veen et al., 2023; Warton et al., 2015).

This article focuses on applying JSDMs to spatial multivariate abundance data that arise when the sites are spatially indexed (e.g. by longitude and latitude). Indeed, the majority of data collected in community ecology are spatially, and possibly also temporally, indexed. One motivating example we consider relates to presence-absence records for 45 butterfly species collected at over 2500 locations across Great Britain (Asher et al., 2001). Along with species records at each location, we have three environmental covariates: (1) the mean annual number of growing degree days above 5 degrees Celsius, (2) the percentage of each grid cell covered by broadleaved woodland and (3) the percentage of each grid cell covered by coniferous woodland. An important feature of these covariates, which is also characteristic of many spatial multivariate abundance datasets, is that they exhibit a noticeable spatial pattern (e.g. the number of growing degree days presents a strong latitudinal gradient; see fig. 2 in Ovaskainen et al. (2016) and the application results from Section 5).

Spatial multivariate abundance data often display evidence of spatial autocorrelation both within a species and between species. To account for such dependence, spatially structured JSDMs have been developed, with the most commonly used instance of these being spatial factor analysis (SFA), also known as spatial latent variable models. Such models rose to prominence with the work of Thorson et al. (2015) and Ovaskainen et al. (2016), and since then several variations of SFA have been developed (e.g. Hui et al., 2023; Shirota et al., 2019; Thorson et al., 2019; Tikhonov et al., 2020). In SFA, a small number of spatially structured latent factors are included to account for residual spatial correlations both within and between species. These spatial factors are often interpreted as a set of unobserved, spatially structured covariates, with the corresponding loadings representing the species-specific responses to these. By making the latent factors spatially structured, as opposed to the original formulation of JSDMs in Pollock et al. (2014) and Warton et al. (2015) where the factors were assumed to be independent across sites, it allows two species to be correlated even if they are observed at different spatial locations.

While SFA and spatially structured JSDMs more generally have grown in popularity for modelling spatial multivariate abundance data, one issue that remains underappreciated in this context is that of spatial confounding. Spatial confounding (a term coined by Hodges & Reich, 2010) refers to the phenomenon where spatially structured random effects, in this case the spatial factors, become collinear with one or more covariates included in the model that are also spatially varying (e.g. the number of growing degree days in the Great Britain butterfly data). Analogous to multicollinearity in other regression settings, spatial confounding has the potential to produce misleading interpretations and inferences in spatially structured JSDMs.

In the case of a single response (e.g. a single species distribution model), spatial confounding is a well-studied but unsettled issue in the literature and continues to remain an active area of research and discussion. We refer the reader to Hanks et al. (2015) and Mäkinen et al. (2022) for overviews of spatial confounding in the context of single response regression models; Khan and Berrett (2023) and Dupont et al. (2023) for more general statistical reviews; and Paciorek (2010), Hughes and Haran (2013), Hefley et al. (2017), Thaden and Kneib (2018), Khan and Calder (2022), Zimmerman and Ver Hoef (2022), Dupont et al. (2022) and Hui and Bondell (2022) among others for research into the consequences of, and methods for, alleviating spatial confounding.

For SFA and spatially structured JSDMs, spatial confounding is a much less studied issue. Shirota et al. (2019) noted the presence of spatial confounding and its implications for the interpretation of covariate effects in their JSDM developments. Lany et al. (2020) compared various types of stacked and spatial JSDMs, arguing for the use of non-spatial or independent JSDMs for evaluating hypotheses about species-environment relationships (see also Khan & Berrett, 2023; Mielke et al., 2020; Zimmerman & Ver Hoef, 2022). Most recently, Van Ee et al. (2022) discussed the related problem of community confounding in JSDMs, where (even) in a non-spatial setting, the latent factors and covariate effects can compete with each other to explain species responses.

In this article, we highlight the issue of spatial confounding in spatially structured JSDMs; particularly in the context of SFA. Using a combination of simulations and case studies, we demonstrate that, when present, spatial confounding means the regression coefficients may only capture part of the explanatory power of spatially structured covariates. This so-called unpartitioned effect (Hui & Bondell, 2022) may or may not be what practitioners are actually interested in and without careful consideration may be misinterpreted. We further show that spatial confounding has implications for other common types of inferences drawn from SFA not examined so far in the literature, including misleading quantification of the relative contributions of covariate effects versus spatial factors in driving covariation in a species community and ambiguous results arising from model-based ordinations formed using the spatial factors.

To treat spatial confounding, we investigate the use of restricted spatial factor analysis (RSFA), building on the ideas of Hodges and Reich (2010), Shirota et al. (2019) and Van Ee et al. (2022). RSFA constrains the spatial factors to be uncorrelated with the covariates, including the intercept. This ensures that the covariates maintain their full explanatory power, estimating what we refer to as a partitioned effect instead. Put another way, RSFA constructs the spatial factors to more authentically represent the residual spatial covariation both within and between species, beyond what can be explained by the (spatially structured) covariates. In doing so, this also seeks to resolve ambiguity in the interpretation of the results of variance partitioning and model-based ordination from a spatially structured JSDM.

It is important to emphasise that RSFA is only one approach to address spatial confounding. There are other approaches to treat spatial confounding, which focus on capturing a different covariate effect to RSFA, say, through a different adjustment on the response

and/or the spatially structured covariates (e.g. Dupont et al., 2022; Marques et al., 2022; Prates et al., 2019; Thaden & Kneib, 2018), although to our knowledge none of these have been adapted to the multivariate abundance data setting. We do not view RSFA as a panacea for all spatial multivariate abundance data analysis. Instead, we aim to raise attention to the issue of spatial confounding in JSDMs: ecologists should integrate the question of whether to adjust for spatial confounding into part of their model-building process (through RSFA or by other means), ensuring that the models and methods they use are consistent with the effects, inferences and ordinations they want to perform and interpret.

The remainder of this article is structured as follows. We begin in Section 2 by reviewing spatial factor analysis and introducing the idea of restricted spatial factor analysis. In Section 3, we discuss the impacts of spatial confounding on several key applications of spatially structured JSDMs, contrasting the results produced from SFA and RSFA. In Section 4, we present results from a numerical study demonstrating these impacts in finite samples and how they can vary in severity depending on the strength of collinearity between the observed and unobserved covariates, and on the relative spatial scale of both. We follow this up with two case studies in Section 5 using the Great Britain butterfly data along with a dataset comprising count records of zooplankton communities from the Southern Ocean Continuous Plankton Recorder (SOCPR) survey (Hosie et al., 2003), to illustrate the differing results or sometimes lack thereof that can arise from applying SFA and RSFA in the presence of spatial confounding. Finally, we conclude with some general remarks in Section 6. The data from both case studies are publicly available, while the code for reproducing the simulation study and both case studies are archived at Hui et al. (2024), noting we used maximum approximate likelihood estimation via template model builder (TMB, Kristensen et al., 2016) to fit all the models in this article.

2 | SPATIALLY STRUCTURED JSDMs

Let y_{ij} denote the recorded response for species $j = 1, \ldots, J$ at site $i = 1, \ldots, N$, where each site is indexed by a spatial coordinate (e.g. latitude and longitude). In addition to the response, let $\mathbf{x}_i = (x_{i1}, \ldots, x_{iP})^\top$ denote a set of P covariates measured at site i, where $x_{i1} = 1$ to represent an intercept term.

2.1 | Spatial factor analysis

We first review spatial factor analysis (SFA), otherwise known as a spatial latent variable model (Ovaskainen et al., 2016; Thorson et al., 2015). As discussed in Section 1, a number of variations to SFA have been proposed, and the developments on spatial confounding below will carry over to these variations. For ease of presentation, however, we focus on the established form of SFA as follows. Let $\mathbf{u}_i = (u_{i1}, \ldots, u_{iK})^\top$ denote a vector of $K \ll J$ latent factors for site i. For model-based ordination, it is common to select

$K = 1$ to 3 to facilitate visualisation (e.g. van der Veen et al., 2023; Warton et al., 2015), although for other applications of SFA, this may be larger. Conditional on $\mathbf{u}_i$, the responses y_{ij} are assumed to be independent observations from a specified distribution, examples of which include the Bernoulli distribution for presence-absence responses, the negative binomial or zero-inflated/hurdle count distributions for overdispersed count responses and the Tweedie distribution for biomass responses (Ovaskainen & Abrego, 2020; Stoklosa et al., 2022). The mean of the response, denoted here as μ_{ij}, is modelled as $g(\mu_{ij}) = \eta_{ij} = \mathbf{x}_i^\top \boldsymbol{\beta}_j + \mathbf{u}_i^\top \boldsymbol{\lambda}_j$, where $g(\cdot)$ is a known link function (e.g. the logit link for presence-absence responses), $\boldsymbol{\beta}_j$ denote species-specific regression coefficients corresponding to the covariates, and $\boldsymbol{\lambda}_j$ denote species-specific loadings for the latent factors. Note the $\boldsymbol{\beta}_j$ are sometimes treated as random effects that are shrunk toward a common distribution (or equivalently, assigned a hierarchical prior distribution, Pollock et al., 2014; Ovaskainen et al., 2016; Ovaskainen & Abrego, 2020); this is useful to borrow strength across species when modelling environmental relationships and also stabilises model fitting because it induces a form of regularisation.

To facilitate the developments below, we write the SFA model in matrix form as follows. We let $\mathbf{X}$ denote the $N \times P$ matrix of covariates formed by stacking the $\mathbf{x}_i$ as row vectors and similarly define $\mathbf{U}$ as the $N \times K$ latent factor matrix formed from stacking the latent factors $\mathbf{u}_i$. Next, we let $\mathbf{B}$ denote the $J \times P$ matrix of regression coefficients formed from stacking the $\boldsymbol{\beta}_j$, and $\boldsymbol{\Lambda}$ denote the $J \times K$ loading matrix formed from stacking the loadings $\boldsymbol{\lambda}_j$. We can then write the SFA model as

$$g(\boldsymbol{\mu}) = \mathbf{X}\mathbf{B}^\top + \mathbf{U}\boldsymbol{\Lambda}^\top, \tag{1}$$

where $\boldsymbol{\mu}$ denotes an $N \times P$ mean matrix with the (i, j)th element equal to μ_{ij} and $g(\cdot)$ is applied element-wise to $\boldsymbol{\mu}$.

As opposed to standard, independent factor analytic models where the elements of $\mathbf{U}$ are assumed to be independent of each other (Warton et al., 2015), the latent factors in SFA are assumed to be spatially structured. Let $(u_{1k}, \ldots, u_{Nk})^\top$ denote the N-dimensional vector representing the kth column of $\mathbf{U}$. Then in SFA, this vector is commonly modelled as a multivariate normal distribution with a zero mean vector and a covariance matrix parametrised by a spatial correlation function. As an example, consider two sites i and i', and suppose the geographic distance between them is denoted by $d_{ii'}$. Then, we can specify an exponential spatial correlation function (e.g. Ovaskainen & Abrego, 2020; Ovaskainen et al., 2016), such that $\text{Cor}(u_{ik}, u_{i'k}) = \exp(-d_{ii'}/\theta_k)$, where $\theta_k > 0$ is a scale parameter governing the strength of the spatial correlation for the kth factor. A more flexible structure replaces the exponential form by a Matérn correlation function (e.g. Thorson et al., 2015), while further extensions of this were reviewed in Section 1.

2.2 | Restricted spatial factor analysis

In SFA, spatial confounding can arise when the spatially structured latent factors exhibit some degree of collinearity with one or more

measured/observed covariates. This can occur despite the fact the spatial factors are assumed to be independent of the covariates: for a given dataset, any two columns of **U** and **X** may exhibit non-negligible finite sample correlations with each other (Hanks et al., 2015; Mäkinen et al., 2022). We present examples of this in both the simulation study and our two case studies on the Great Britain butterfly and zooplankton SOCPR survey data; in particular, from a fitted SFA we can compute the empirical correlation between each of the measured covariates and predicted spatial factors as one measure of the degree of spatial confounding present in spatial multivariate abundance data (see also Hefley et al., 2017; Yuan et al., 2017; Schulz et al., 2020, for related ideas). It is also worth pointing out the finite sample correlations between the columns **X** and **U**, and hence the presence and extent of spatial confounding, is not guaranteed to diminish as the number of sites sampled N increases; see Appendix A for further discussion of this point.

Next, we introduce one approach to manage spatial confounding, which we refer to as restricted spatial factor analysis (RSFA). Put simply, RSFA treats spatial confounding by constructing the spatial factors such that they are always uncorrelated with the measured predictors. To introduce this approach in detail, we first review the concept of projection and residual projection matrices. Let $\mathbf{P} = \mathbf{X}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top$ denote the $N \times N$ projection matrix formed from the matrix of covariates **X**. Statistically, **P** projects a N-dimensional vector onto the space spanned by the covariates (i.e. the column space of **X**). For example, given the ordinary least squares estimates $\hat{\alpha} = (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{y}$ for some vector of responses **y**, we can write the corresponding fitted values as $\hat{\mathbf{y}} = \mathbf{X}\hat{\alpha} = \mathbf{P}\mathbf{y}$ (i.e. it transforms the responses to fitted values). Next, let $\mathbf{I}_N$ denote the $N \times N$ identity matrix. Then, the residual projection matrix of **X** is given by $\mathbf{P}_\perp = (\mathbf{I}_N - \mathbf{P})$, named as such because it transforms the responses in a linear regression model to residuals, $\hat{\mathbf{r}} = \mathbf{y} - \mathbf{X}\hat{\beta} = (\mathbf{I}_N - \mathbf{P})\mathbf{y} = \mathbf{P}_\perp\mathbf{y}$. Importantly, it is straightforward to show $\mathbf{X}^\top\hat{\mathbf{r}} = \mathbf{X}^\top\mathbf{P}_\perp\mathbf{y} = \mathbf{0}$, meaning the residuals are uncorrelated with the covariates.

In RSFA, we constrain the spatial factors to be uncorrelated with (also referred to as orthogonality) the measured predictors by using the residual projection matrix defined above. Specifically, left-multiplying the matrix **U** from Equation (1) by $\mathbf{P}_\perp$, we obtain

$$g(\boldsymbol{\mu}) = \mathbf{X}(\mathbf{B}^*)^\top + \mathbf{P}_\perp\mathbf{U}\boldsymbol{\Lambda}^\top = \mathbf{X}(\mathbf{B}^*)^\top + \mathbf{U}^*\boldsymbol{\Lambda}^\top, \qquad (2)$$

where **U*** refers to an $N \times K$ matrix of restricted spatial factors, and $\mathbf{B}^* = (\boldsymbol{\beta}_1^*, \dots, \boldsymbol{\beta}_J^*)$ denotes the $J \times P$ matrix of regression coefficients which RSFA estimates. Note the dimensions of **B** and **B*** are both $J \times P$, and likewise the dimensions of **U** and **U*** are both $J \times K$. More importantly, we have $\mathbf{X}^\top\mathbf{U}^* = \mathbf{X}^\top\mathbf{P}_\perp\mathbf{U} = \mathbf{0}$, and so the spatial factors and covariates are restricted to be uncorrelated in RSFA. This exemplifies how RSFA treats spatial confounding: whereas in SFA the columns of **U** and **X** may be correlated with each other, in RSFA they are constrained to be uncorrelated. The loading matrix $\boldsymbol{\Lambda}$ is the same in both SFA and RSFA.

As an aside, note that it is not necessary to construct the residual projection matrix and RSFA based on all the covariates. Instead,

we may choose to group covariates (columns of **X**) into those we desire to be uncorrelated of the spatial factors, and the remainder of which no restriction is placed on them. We can then construct $\mathbf{P}_\perp$ and hence the restricted spatial factors **U*** in Equation (2) based only on the former group, and doing so treats spatial confounding based only on a subset of the measured covariates. For ease of presentation however, we focus on the case where all covariates are constrained as presented in Equation (2).

At first glance, it appears as if RSFA and SFA are distinct models. In fact, Equations (1) and (2) are reparametrisations of each other. That is, we can start with SFA and then add and subtract the quantity $\mathbf{P}\mathbf{U}\boldsymbol{\Lambda}^\top$ to the right-hand side of (1) to obtain

$$\begin{aligned} g(\boldsymbol{\mu}) &= \mathbf{X}\mathbf{B}^\top + \mathbf{P}\mathbf{U}\boldsymbol{\Lambda}^\top + (\mathbf{I}_N - \mathbf{P})\mathbf{U}\boldsymbol{\Lambda}^\top \\ &= \mathbf{X}\left(\mathbf{B}^\top + (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{U}\boldsymbol{\Lambda}^\top\right) + \mathbf{P}_\perp\mathbf{U}\boldsymbol{\Lambda}^\top \\ &= \mathbf{X}(\mathbf{B}^*)^\top + \mathbf{U}^*\boldsymbol{\Lambda}^\top, \end{aligned} \qquad (3)$$

where $(\mathbf{B}^*)^\top = \mathbf{B}^\top + (\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{U}\boldsymbol{\Lambda}^\top$ and the last line is exactly the RSFA model in Equation (2). One immediate consequence of the above result is that RSFA is straightforward to implement. For instance, in both our simulation and case studies where all models were estimated using maximum approximate likelihood estimation via TMB (Kristensen et al., 2016), RSFA requires only a few additional lines of code in addition to that used to fit SFA to extract the estimates of **B*** and **U*** (see also Hanks et al., 2015, for estimation using Markov Chain Monte Carlo).

Because SFA and RSFA are reparameterisations of each other, this implies that Equations (1) and (2) will produce the same fitted values and maximised likelihood (or deviance) values. This is analogous to multicollinearity in standard regression settings, where fitted values and point predictions are not affected by the degree of collinearity between measured covariates. More importantly, as noted by Hanks et al. (2015), it means we cannot apply standard inference tools such as residual diagnostics or information criteria to select between SFA and RSFA. We do not view the inability to use model selection tools as a drawback however; instead, it compels practitioners to proactively think about the model assumptions and interpretations underlying the construction of their spatial JSDM.

To conclude this section, we highlight that because spatial confounding can arise when one or more spatial factors are correlated with the measured covariates, then a naive but straightforward alternative to RSFA for managing spatial confounding is to forgo the use of spatially structured factors in the first place, instead favouring independent factor analysis as in Warton et al. (2015). Indeed, due to the possibility of spatial confounding, Lany et al. (2020) made a case for use of the independent factor analysis if the goal of the analysis is specifically to assess hypotheses about species-environment relationships, and even if the multivariate abundance data are known/found to be spatially structured. On the other hand, there is well-established research advocating for the use of spatial random effects whenever spatial correlation is detected in the residuals of the data, to ensure that valid hypothesis tests and uncertainty quantification can be performed

(Dormann et al., 2007). Moreover, Hui et al. (2022) showed that SFA almost always performs better than assuming independence in the latent factors when predicting community composition at new spatial locations, In Section 4, we also demonstrate using simulations that independent factor analysis performs consistently worse than both SFA and RSFA for recovering a spatially correlated unobserved covariate.

3 | IMPACTS OF SPATIAL CONFOUNDING

In this section, we discuss the impacts of spatial confounding for multivariate abundance data by comparing the inferences arising from SFA versus RSFA models in Equations (1) and (2), respectively. We focus on three predominant usages of spatially structured JSDMs in ecology: Quantifying the effects of measured covariates on community composition, assessing the relative contributions of measured covariates versus residual processes in driving spatial covariation in the community and model-based residual ordination. While there are other reasons why spatially structured JSDMs are applied (e.g. for recovering species networks and interspecies associations; Astarloa et al., 2019; Haak et al., 2020), many of these are related and will be impacted by spatial confounding similarly to the three inferences detailed above.

3.1 | Covariate effects

The covariate effects **B** and **B*** from the SFA and RSFA models respectively can differ substantially in the presence of spatial confounding. In SFA, the potential collinearity between the covariates **X** and spatial factors **U** in Equation (1) means part of the explanatory power of the former may be absorbed by the inclusion of the latter. This is analogous to standard regression settings, where, upon including a new covariate that is collinear with one or more covariates already in the model, this 'soaks up' part of the explanatory power the already included covariates had, and the interpretation of all covariate effects must be adjusted accordingly (Graham, 2003). With this in mind, we can interpret **B** as an unpartitioned covariate effect because the spatial factors **U** are not decomposed into components relating to, and residual to, the covariates **X**.

By contrast, for RSFA in Equation (3), there is an explicit partitioning of the spatial factors into two components: **PU**, which is the projection of the spatial factors onto the space spanned by the covariates, and **P$_\perp$U**, which is the residual component of the spatial factors that cannot be explained by the covariates. The first part is combined with **B** to form the so-called partitioned covariate effect **B***. The idea of estimating a partitioned effect is not new, for instance, the concept of partial regression plots in linear models (Velleman & Welsch, 1981) is based on the same idea. More importantly, for understanding the effect of the covariates on species responses, the partitioned effect (re)incorporates the explanatory power that was lost due to the spatial factors absorbing this. That is, we manage the

collinearity arising from spatial confounding by moving the portion of residual spatial correlation (both within and between species) that can be explained by the covariates, back into the covariate component itself. What remains is the component **U*****Λ**$^\top$, which is a more authentic representation of the residual covariation in the model: it is the component left over accounting for the full explanatory power of the covariates.

To summarise, when seeking to quantify how species distributions are related to physical and environmental habitat, spatial confounding can cause part of what the covariates can explain about the response to be inadvertently contained in these spatial factors. RSFA estimates partitioned coefficients **B*** which attempts to correct this, and thereby quantify the full association between the covariates and the species response. As shown in Hanks et al. (2015) and Khan and Calder (2022) among others, as well as in Section 4, if we are interested in one type of effect (either partitioned or unpartitioned) but estimate and make inference on the other, then it may result in biased estimates and erroneous inference.

3.2 | Between-species correlations and variance partitioning

In Section 2.1, we show that by imbuing the latent factors in Equation (1) with an explicit spatial structure, SFA allows for spatial correlations both within and across species. For instance, assuming an exponential spatial correlation function, the residual covariance between species $j, j' = 1, \ldots, J$ at spatial coordinates $i, i' = 1, \ldots, N$ can be shown to be

$\text{Cov}(\eta_{ij}, \eta_{i'j'}) = \sum_{k=1}^{K} \lambda_{jk} \lambda_{j'k} \exp(-d_{ii'}/\theta_k)$, where $d_{ii'}$ is the geographic distance between locations i, i', and $\theta_k > 0$ is a scale parameter governing the strength of the spatial correlation for the kth factor. Compare this now to what arises from RSFA, where from Equation (2) we can show that $\text{Cov}(\eta_{ij}, \eta_{i'j'}) = \sum_{k=1}^{K} \lambda_{jk} \lambda_{j'k} \text{Cov}(U_{ik}^*, U_{i'k}^*) = [\mathbf{P}_\perp]_i^\top \left(\sum_{k=1}^{K} \lambda_{jk} \lambda_{j'k} \mathbf{R}_k\right) [\mathbf{P}_\perp]_{i'}$, where $[\mathbf{P}_\perp]_i$ denotes the ith row in the residual projection matrix $\mathbf{P}_\perp$, and $\mathbf{R}_k$ is an $N \times N$ matrix whose elements are given by $\exp(-d_{ii'}/\theta_k)$. Upon converting both quantities to correlations, we obtain:

$$\text{SFA}: \text{Cor}(\eta_{ij}, \eta_{i'j'}) = \frac{\sum_{k=1}^{K} \lambda_{jk} \lambda_{j'k} \exp(-d_{ii'}/\theta_k)}{\sqrt{\sum_{k=1}^{K} \lambda_{jk}^2} \sqrt{\sum_{k=1}^{K} \lambda_{j'k}^2}}, \quad (4a)$$

$$\text{RSFA}: \text{Cor}(\eta_{ij}, \eta_{i'j'}) = \frac{[\mathbf{P}_\perp]_i^\top \left(\sum_{k=1}^{K} \lambda_{jk} \lambda_{j'k} \mathbf{R}_k\right) [\mathbf{P}_\perp]_{i'}}{\sqrt{[\mathbf{P}_\perp]_i^\top \left(\sum_{k=1}^{K} \lambda_{jk}^2 \mathbf{R}_k\right) [\mathbf{P}_\perp]_i} \sqrt{[\mathbf{P}_\perp]_{i'}^\top \left(\sum_{k=1}^{K} \lambda_{j'k}^2 \mathbf{R}_k\right) [\mathbf{P}_\perp]_{i'}}}. \quad (4b)$$

When spatial confounding is present, Equation (4a) may no longer accurately represent the residual between-species spatial correlation because it has absorbed a portion of the between-species covariation attributable to the measured spatially structured covariates. That is, the estimated values from using (4a) may instead resemble between-species correlations arising from the measured

covariate component of SFA, $\mathbf{XB}^\top$. Arguably, a more accurate representation of the residual spatial correlation between species instead arises from Equation (4b). Although straightforward to compute, this formula is noticeably more difficult to interpret: The residual correlation between two species at any two different spatial locations in RSFA depends on both the measured covariates at all N spatial locations and the entire spatial correlation matrix $\mathbf{R}_k$. Even when $i = i'$, it is straightforward to see the residual correlation between any two species at the same site does not change with spatial location in SFA. But in RSFA, the residual correlation at any given site depends on all other sites through the measured covariates $\mathbf{X}$, and the pairwise distances $d_{ii'}$ contained in $\mathbf{R}_k$.

The fact that Equation (4b) depends on the residual projection matrix precisely captures the concept of what the *residual* spatial correlation between two species means. We can draw an analogy with linear regression, where we recall the residuals are given as $\hat{\mathbf{r}} = \mathbf{P}_\perp \mathbf{y}$. If the responses are assumed to be independent observations with error variance σ^2, then the covariance between any two residuals is given by $\mathrm{Cov}(\hat{r}_i, \hat{r}_{i'}) = \sigma^2 [\mathbf{P}_\perp]_i^\top [\mathbf{P}_\perp]_{i'}$, which is a special case of Equation (4b). This analogy suggests RSFA presents a more authentic calculation of the residual spatial correlation both within and between species in the presence of spatial confounding: it is the correlation remaining after accounting for the full explanatory power of the covariates.

The differing calculations of the residual spatial correlations between SFA and RSFA have implications when it comes to variance partitioning, which is a common form of inference when applying JSDMs (Hui et al., 2023; Ovaskainen et al., 2016). That is, we can construct measures for the relative proportion of variation explained by the measured covariates versus the spatial factors (i.e. the residual variation) for SFA, this variance partitioning is typically done in terms of the unpartitioned effects $\mathbf{XB}^\top$ and $\mathbf{U\Lambda}^\top$. However, in the presence of spatial confounding, such variance partitioning may be misleading because a portion of the variation that can be explained by the measured covariates will instead be attributed to the variation explained by the spatial factors. This is exemplified by the Great Britain butterfly data analysis where, in their application, Ovaskainen et al. (2016) noted that the overall amount of variation attributed to the measured covariates actually reduced when spatial factors were included in the model. We also observed this in our analysis of the same dataset in Section 5. This is emblematic of how spatial confounding can result in unexpected changes in variance partitioning.

In RSFA, the variance partitioning is performed in terms of the partitioned effects $\mathbf{X}(\mathbf{B}^*)^\top$ and $\mathbf{U}^*\mathbf{\Lambda}^\top$. By doing so, the variation attributed to the former quantity seeks to capture the full extent to which the variation in the response can be explained by the measured covariates, including the spatial variation in the response that can be captured by the spatial variation in the predictors, while the proportion of variation attributed to $\mathbf{U}^*\mathbf{\Lambda}^\top$ better reflects the actual component of variation in the species community, which cannot be attributed to measured covariates.

3.3 | Residual ordination

One key application of factor analytic models in community ecology is residual ordination, that is a low-dimensional graphical representation of residual between-species covariation (van der Veen et al., 2023; Warton et al., 2015). For SFA, this involves constructing a scatterplot of the estimated spatial factors (columns of $\mathbf{U}$), which may be interpreted as a set of unmeasured covariates or community composition gradients.

In the presence of spatial confounding, such residual ordinations from SFA need to be carefully interpreted. Consistent with the discussions in Sections 3.1 and 3.2, because the spatial factors may absorb some of the explanatory power of the measured covariates, then an ordination formed from the columns $\mathbf{U}$ in Equation (1) may not reflect the notion of a residual ordination per se. Instead, the ordination represents a combination of the spatial covariation within and between species not accounted for by the measured covariates, plus a component that is attributable to the measured covariates themselves. Put another way, the estimated community composition gradients from SFA may resemble (exhibit some correlation with) the measured covariates themselves; see the simulations and case studies in Sections 4 and 5, respectively, for empirical evidence of this. This means an ordination of them can, due to spatial confounding, bear similarity to simply plotting the measured covariates as a function of their spatial locations.

If the statistical analysis aims to construct an ordination that more authentically represents spatial covariation in addition to that explained by measured covariates, then one should consider fitting RSFA and plotting the restricted spatial factors $\mathbf{U}^*$. From Equation (3), these spatial factors are, by construction, a term that is uncorrelated with the measured covariates, and represents what remains after accounting for the full explanatory power of the covariates (as quantified by $\mathbf{B}^*$). In turn, ordination of the $\mathbf{U}^*$ more accurately reflects the idea of a residual ordination: a set of community composition axes representing species covariation not attributable to the included covariates.

3.4 | Degree of impacts

In the case of a single response, extensive research has been done on the relationship between the degree of spatial confounding and the relative spatial scales of the measured versus the unobserved covariates (which are approximated by the spatially structured random effects; see Mäkinen et al., 2022; Paciorek, 2010) and references therein. With multivariate abundance data and SFA, this relationship is more complex because each spatial factor (column of $\mathbf{U}$) acts as a separate random effect typically with its own spatial scale. Indeed, in many applications of SFA, the aim is to capture the residual spatial correlation between and within each species at multiple spatial scales (Ovaskainen et al., 2016; Thorson et al., 2016). As such, the possibility of spatial confounding arising is expected to be even greater with multivariate abundance data and spatially structured JSDMs.

The degree of spatial confounding for a given covariate (column of **X**) will also depend on how its spatial scale compares in relation to the estimated spatial scales of all the latent factors. In particular, we show via simulation in the next section that spatial confounding will generally be more prominent in SFA when the measured covariates are comparably coarse-scaled and spatially smooth (i.e. the spatial range of the observed covariate is larger than any unobserved predictors that the spatial factors are trying to capture). Conversely, if the measured covariates are fine-scaled (i.e. their spatial range is less than that of any unobserved predictors that the spatial factors are trying to capture) then spatial confounding is less likely to arise and/or have less impact on estimation and inference when applying SFA (this can arise, for instance, by using downscaled microclimate co-variates, Lembrechts et al., 2019). Note similar conclusions regarding the importance of scale have been obtained previously in the context of a single response regression model by Paciorek (2010) and Mäkinen et al. (2022).

4 | SIMULATION STUDY

We performed a simulation study to investigate the impacts of spatial confounding discussed above, comparing the performance of independence factor analysis, SFA in Equation (1) and RSFA in Equation (2) in terms of estimation and inference on covariate effects and residual ordination. Note we only consider models involving latent factors: as mentioned previously, some have recommended the use of models that do not explicitly account for residual between-species correlations, that is, stacked species distribution models, when the aim is purely to quantify the effects of measured covariates. However, given that stacked species distribution models do not facilitate inferences such as model-based ordinations, then we do not include such a model for comparison here.

The full details of the simulation design are provided in Appendix B, but to summarise, we adapted the design of Dupont et al. (2022) and simulated spatial multivariate abundance data from a model with $N = 400$ spatial locations, $J = 10$ species, one measured covariate and two unmeasured covariates. The single measured covariate was constructed using the form $\mathbf{x}_{obs} = 0.5\mathbf{f} + \boldsymbol{\epsilon}$, where $\mathbf{f}$ denotes a spatial field formed via a set of spatial basis functions (Hui et al., 2023; Tzeng & Huang, 2018) and accompanying normally distributed weights, and the elements of $\boldsymbol{\epsilon} = (\epsilon_1, \ldots, \epsilon_N)^\top$ are independent fine-scaled normally distributed error terms. Meanwhile, the two unmeasured covariates were constructed using the form $\mathbf{x}_{unobs,1} = -C_0\mathbf{f} - \bar{\mathbf{f}}_1$ and $\mathbf{x}_{unobs,2} = -C_0\mathbf{f} + \bar{\mathbf{f}}_2$, where C_0 is a constant and $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$ are spatial fields built similarly to, but independently of, $\mathbf{f}$ and of each other. We let $\mathcal{X}$ denote the $N \times 4$ matrix where the first column is a vector of ones representing an intercept term, while columns two to four comprise $\mathbf{x}_{obs}$, $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$, respectively. We then constructed the mean model $g(\boldsymbol{\mu}) = \mathcal{X}\mathcal{B}^\top$ where $\mathcal{B}$ is an $J \times 4$ matrix of species-specific regression coefficients whose elements are generated from a normal distribution with mean and variance given in Appendix B.

Finally, we simulated spatial multivariate abundance data with $J = 10$ species using this mean model and considering three response types: (1) presence-absence records; (2) overdispersed counts; and (3) biomass (non-negative continuous) responses.

In the above set-up, the constant C_0 controls the degree of collinearity between the measured and unmeasured covariates and hence the severity of spatial confounding in the data generation process. We set $C_0 = 0, 0.5, 1$, where $C_0 = 0$ corresponds to negligible spatial confounding. Furthermore, to assess how the impact of spatial confounding varies depending on the spatial scale of $\mathbf{x}_{obs}$ relative to $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$, we examined two settings depending on whether $\mathbf{f}$ was a coarse-scaled spatial field and $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$ are both fine-scaled spatial fields, or vice versa; see Appendix B for details. For each combination of response type, the value of C_0 and whether $\mathbf{x}_{obs}$ was a coarse- or fine-scaled spatial covariate, we simulated 200 datasets.

For each simulated dataset, we assumed only $\mathbf{x}_{obs}$ (i.e. the first two columns of $\mathcal{X}$) were observed and proceeded to fit independent factor analysis, SFA and RSFA assuming $K = 2$ latent factors. In this way, when $C_0 \neq 0$ the true effect of the measured covariate (i.e. its full explanatory power) is defined as the effect given by $\mathbf{x}_{obs}$ and a corresponding portion given by $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$ (due to them sharing the spatial field $\mathbf{f}$). Analogously, the true residual spatial covariation is given by the portion of $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$ not attributable to the measured covariate (i.e. the portion due only to $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$). For each fitted model, we assessed performance by comparing the estimated matrix of regression coefficients to the above-defined true covariate effect and by comparing the two predicted latent factors to the above-defined true residual variation.

We provide results for the case of presence-absence responses, while results for overdispersed counts and biomass responses present similar trends to those seen below and are given in Appendix B. When $C_0 = 0$ and there is negligible spatial confounding, all three methods perform well in terms of absolute bias, root mean squared error (RMSE) and coverage probability of corresponding 95% confidence intervals for the true covariate effects (Figure 1). However, as C_0 and the degree of spatial confounding increased, the performance of SFA noticeably deteriorates because it fails to capture the full explanatory power of the measured covariate, instead estimating the unpartitioned effect **B** (consistent with the discussion in Section 3.1). By contrast, both independent factor analysis and RSFA (which estimates **B***) are more successful at quantifying and performing inference on the full association between the covariates and the species response. When comparing the left and right columns of Figure 1, we see that the consequences of spatial confounding are more severe for SFA when the measured covariate is coarse-scaled (i.e. spatially smooth), relative to the unmeasured covariates; this is consistent with the discussion in Section 3.4. In Appendix B, we also present results on the average width of the confidence intervals, from which we note that RSFA tends to have generally narrower intervals compared to independent factor analysis, while the intervals for SFA are typically the widest even though they have the worst coverage probability as C_0 increased.

Spatial confounding also had an impact on the recovery of the true residual spatial covariation, which directly affects variance partitioning
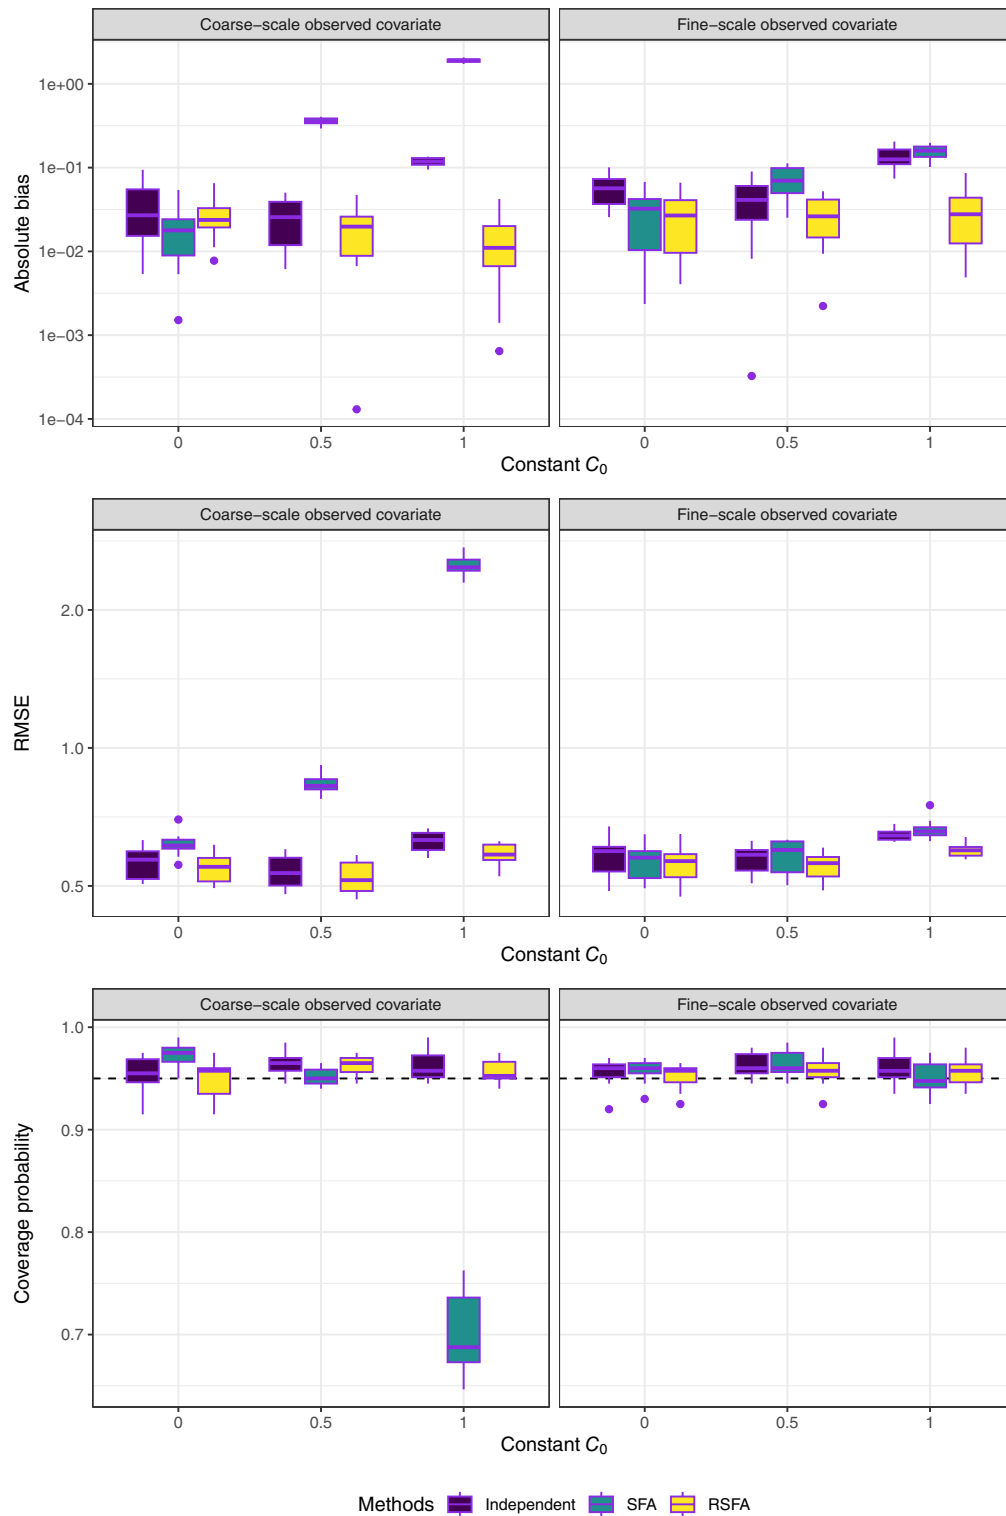
FIGURE 1 Simulation results for the estimated regression coefficients, comparing independent factor analysis, SFA and RSFA in the case of presence-absence responses. The left and right columns correspond to $\mathbf{x}_{obs}$ being a coarse- and fine-scaled spatial field relative to $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$, while performance was assessed in terms of absolute bias (top row), root mean squared error (RMSE; middle row) and empirical coverage probability of 95% confidence intervals across 200 simulated datasets. In each panel, the performance of each method is plotted against C_0, where each method is represented by a different (coloured) boxplot constructed based on $J = 10$ species.

and residual ordination. The recovery of the spatial variation was assessed by computing the Procrustes error between the predicted latent factors and the true residual spatial variation (as given by $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$

), where the Procrustes error can be thought of as the sum of squared differences between the predicted latent factors and the true residual spatial variation, after rotating and scaling the former so that it is

as similar in shape as possible to the latter (see Borcard et al., 2011, for technical details). The impact was most noticeable when the measured covariate was spatially smooth and when $C_0 = 1$, where we see RSFA (which estimates $\mathbf{U}^*$) tended to perform better in terms of the Procrustes error (left half of Figure 2). When C_0 and the degree of spatial confounding was small, SFA and RSFA performed similarly in terms of estimating the latent factors. Likewise, and similar to Figure 1, when the measured covariate was fine-scaled compared to the unmeasured covariates, the impact of spatial confounding was greatly reduced. Independent factor analysis consistently performed the worst in terms of recovering the residual spatial variation; this is attributed to the fact that it fails to explicitly account for spatial correlations in the residual terms and so suffers from model misspecification. Finally, as mentioned in Section 2.2, one simple method of assessing the degree of spatial confounding in a given dataset is based on computing the correlations between the measured covariates and predicted factors. This is exemplified in the right half of Figure 2, which plots the maximum Pearson correlation between $\mathbf{x}_{obs}$ and $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$. Consistent with the results seen above on its impacts, the extent of spatial confounding was most severe when the measured covariate was spatially smooth and $C_0 = 1$: the median correlation between the predicted latent factors from SFA and the observed covariate exceeded 0.8. Interestingly, when $C_0 = 0$ there were still non-negligible correlations between the latent factors from SFA and the observed covariate: this is due to finite sample correlations that may still arise between the spatial fields $\mathbf{f}$ and $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$ as discussed at the beginning of Section 2.2. The correlations were substantially reduced for SFA when the measured covariate was spatially more fine-scaled than unobserved covariates. In contrast, both independent factor analysis and RSFA exhibited negligible correlations between the latent factors and the measured covariates, where the latter occurs by construction as seen in the form of (2).

Overall, this simulation study improves our understanding of the impacts of spatial confounding on JSDMs: if the aim of the analysis is to perform inference on the full explanatory effect of response due to measured covariates and have the latent factors more accurately represent residual covariation not explained by these measured covariates, then RSFA is typically more successful at achieving this compared to SFA and independent factor analysis. However, the degree of impact of spatial confounding varies greatly depending on the relative spatial scales of measured versus unmeasured covariates (see also Mäkinen et al., 2022; Paciorek, 2010), and the extent to which the latter contains a portion of the spatial variation that can be captured by the latter. Finally, we performed additional simulations at the larger sample size of $N = 1000$, with results (not shown) overall exhibiting very similar trends to those at $N = 400$ and demonstrating that the consequences of spatial confounding in the simulation design do not diminish even with more data.

5 | CASE STUDY—GREAT BRITAIN BUTTERFLY DATA

We provide two case studies illustrating the impacts of spatial confounding on statistical analysis in multivariate abundance data, comparing the results obtained from SFA to those obtained using
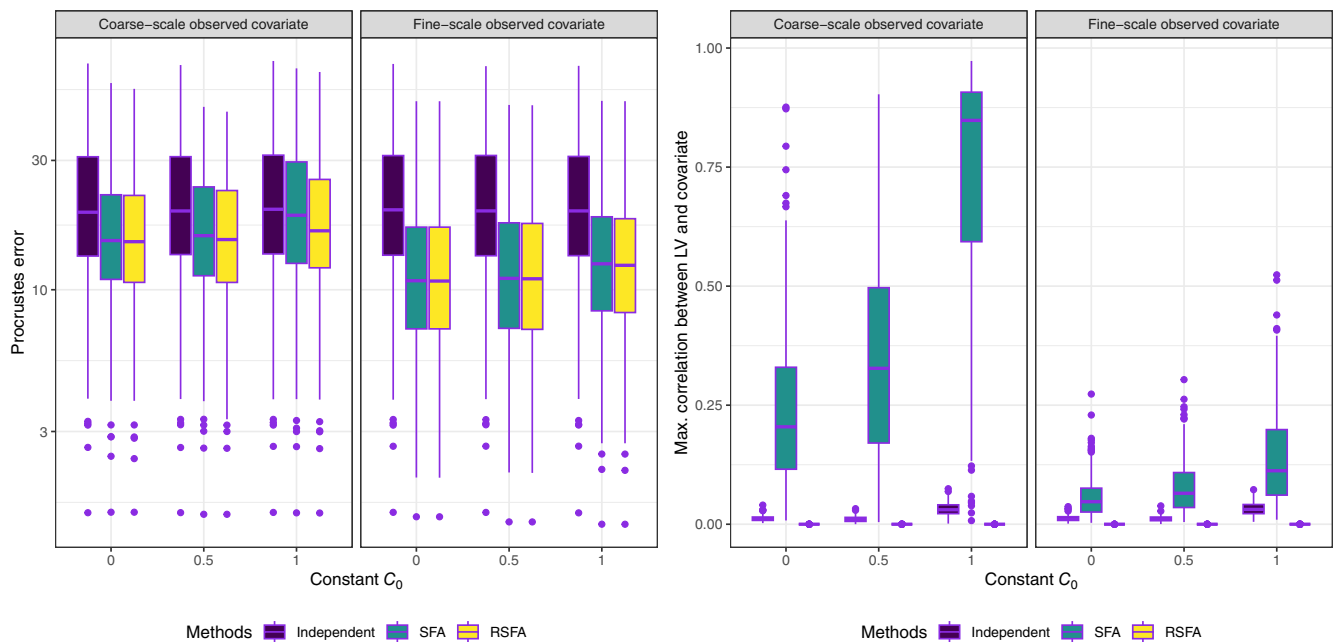
FIGURE 2 Simulation results for the predicted latent factors, comparing independent factor analysis, SFA and RSFA in the case of presence-absence responses. Performance was assessed in terms of Procrustes error between the predicted latent factors and true residual spatial variation (effectively scaled versions of $\bar{\mathbf{f}}_1$ and $\bar{\mathbf{f}}_2$; left half), and the maximum Pearson correlation between the two predicted latent variables and the $\mathbf{x}_{obs}$. Within each half, the left and right columns of each panel correspond to $\mathbf{x}_{obs}$ being a coarse- and fine-scaled spatial field relative to $\mathbf{x}_{unobs,1}$ and $\mathbf{x}_{unobs,2}$, while in each panel, the performance of each method is plotted against C_0, where each method is represented by a different (coloured) boxplot constructed based on the 200 simulated datasets.

an independent factor analysis (i.e. using the non-spatial latent variable model Warton et al., 2015) and RSFA. Indeed, results from the simulation study in the preceding section suggest the importance of performing sensitivity analysis between SFA, RSFA and independent factor analysis, to assess the severity of spatial confounding. For brevity, in this section we focus only on the Great Britain butterfly presence-absence data; results from our application to count records of zooplankton from the SOCPR survey are presented in Appendix D.

The Great Britain butterfly dataset we consider in this case study comprises presence-absence records of $J = 45$ species collected at $N = 2609$ sites spanning Great Britain. The recorded prevalence across the 45 species ranged from 0.067 to 0.963, with an average of 0.487. More information about the data collection process can be found at Asher et al. (2001), while Ovaskainen et al. (2016) used it to demonstrate an application of SFA. The data were subsequently used and made publicly available as part of Wilkinson (2018).

Figure 3 presents spatial maps of three covariates available for this dataset (all standardised to have mean zero and variance one). Note the prominent smooth latitude gradient of growing degree days (GDD; left panel) and to a lesser extent the percentage of coniferous woodland (right panel). Based on Section 3.4 and the empirical results from Section 4, this suggests that spatial confounding may pose a problem when applying SFA, and care needs to be taken when interpreting results from such a fitted model. Figure 4 presented the estimated regression coefficients based on applying SFA, RSFA and independent factor analysis, all assuming linear effects of the three covariates and $K = 2$ factors. Across all three predictors and 45 species, RSFA and independent factor analysis produced broadly similar estimates of the covariate effects, while the estimates from SFA

differ greatly from these two models and tended to be attenuated toward zero. We stress that such attenuation is not always expected to occur when spatial confounding is present Paciorek (2010) and Khan and Berrett (2023), but the fact the estimated values differ greatly between **B** in SFA and **B*** in RSFA suggests that, consistent with the findings from Section 4, the former is being impacted by the collinearity between the measured covariates and the spatial factors and that it is important to decide whether the interest lies in the unpartitioned versus the partitioned effects.

The impact flows into inference on the covariate effects, where in Table 1 we observe that, across all three predictors, the proportion of species judged to have statistically significant effects (based on whether the corresponding 95% confidence intervals excludes zero or not) is considerably lower for SFA compared to RSFA and independent factor analysis; see also Appendix C for caterpillar plots of the confidence intervals. This is consistent with the discussion in Section 3.1, where part of the explanatory power of the measured covariates is being absorbed by the spatial factors in the presence of spatial confounding. The partitioned effects in RSFA attempt to treat this: by reincorporating the full explanatory power of the covariates, this generally leads to a greater number of statistically significant effects when performing inference on **B*** versus **B**. Equivalently, we can view as being that if the focus is on partitioned effects, then SFA can fail to provide valid inference on this; see also the results on coverage probability in Figure 1 of the simulation study.

When we examine variance partitioning from the three models on a per-species basis, we observe that independent factor analysis and RSFA are, again, more similar to each other compared to the results from SFA (Figure 5). For instance, in SFA the three covariates explain less than 25% of the variation across most species, compared
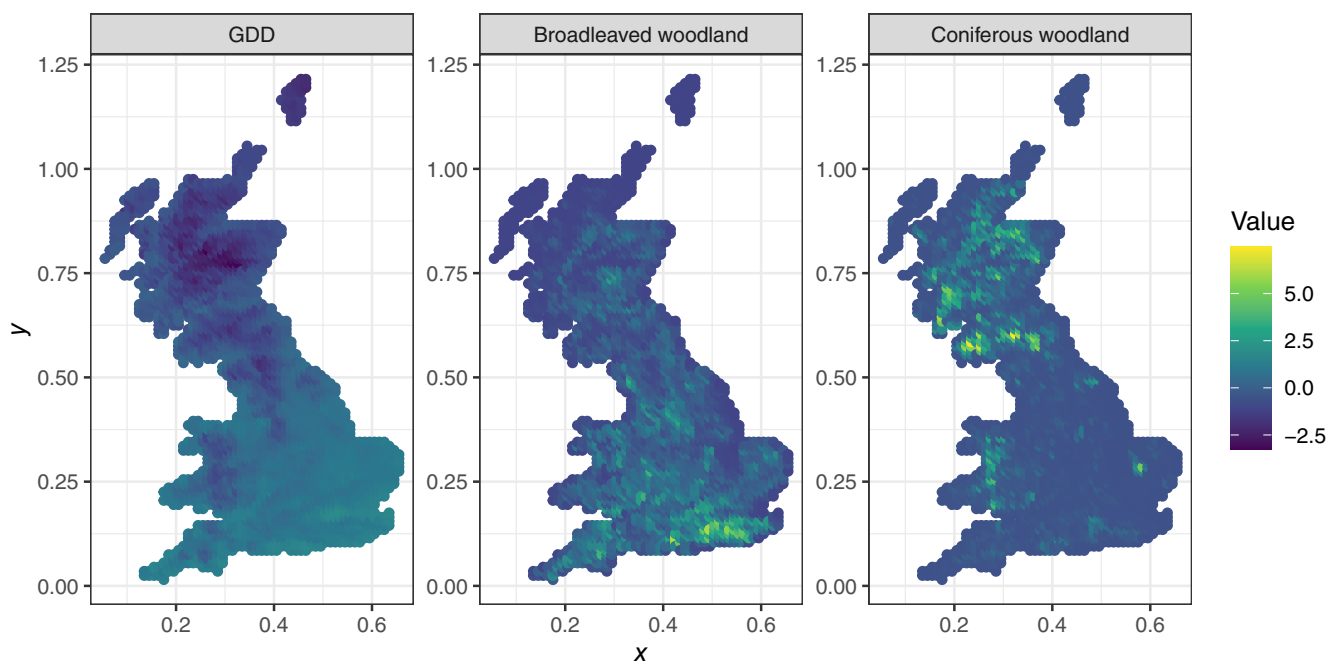

FIGURE 3 Spatial plots of the three covariates (from left to right: Mean annual number of growing degree days above 5°C, percentage of broadleaved woodland, percentage of coniferous woodland) in the Great Britain butterfly data.
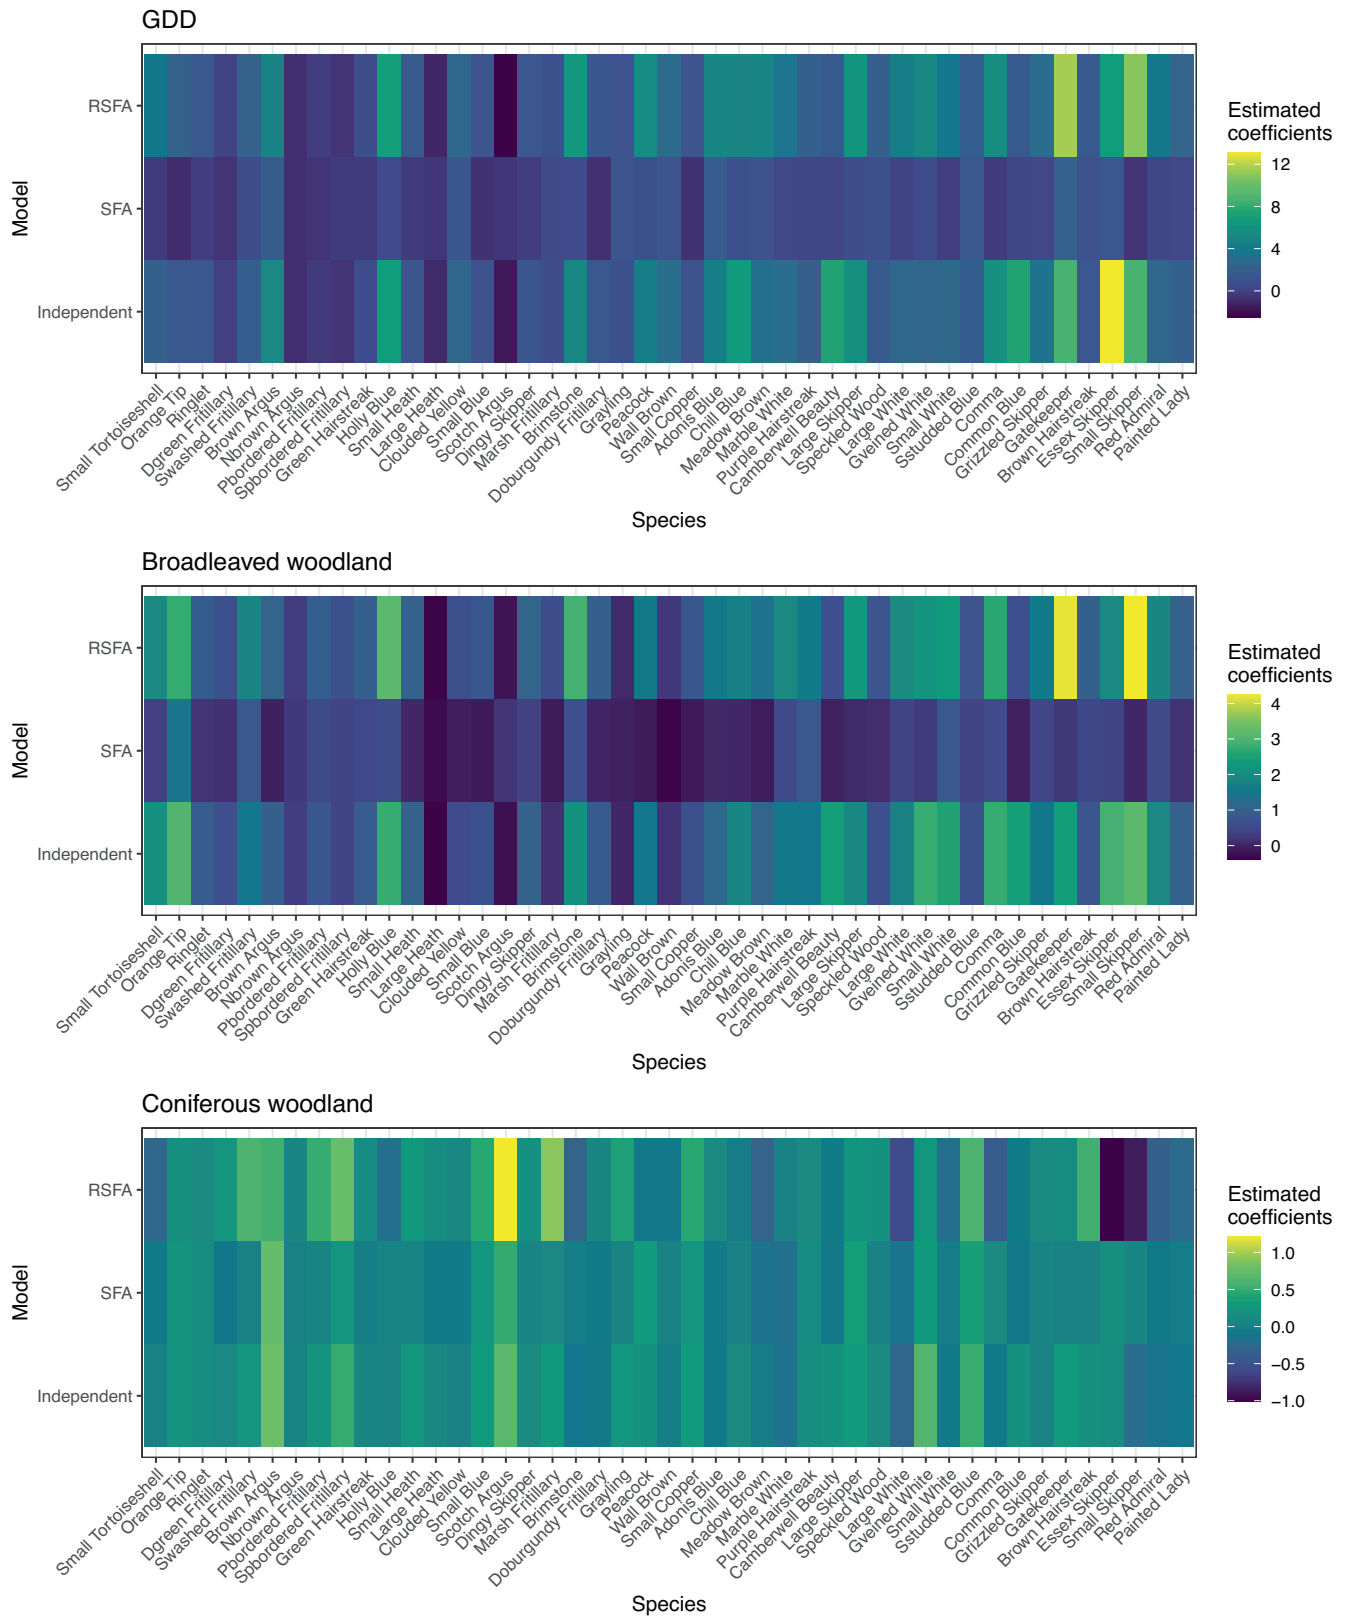
FIGURE 4 Tile plots of the $J = 45$ estimated species regression coefficients and three measured covariates (mean annual number of growing degree days above 5°C, percentage of broadleaved woodland, percentage of coniferous woodland) in the Great Britain butterfly data, estimated using independent factor analysis, SFA and RSFA.

TABLE 1 Proportion of statistically significant results (based on whether 95% Wald confidence intervals exclude zero or not) for each of the three measured covariates (mean annual number of growing degree days above 5°C, percentage of broadleaved woodland, percentage of coniferous woodland) and three models fitted to the Great Britain butterfly data.

Covariate	Model	Prop. of significant results
GDD	Independent	0.978
	SFA	0.733
	RSFA	0.956
Broadleaved woodland	Independent	0.956
	SFA	0.511
	RSFA	0.956
Coniferous woodland	Independent	0.378
	SFA	0.244
	RSFA	0.667

with the independent factor analysis where the covariates tend to explain more than 50% of the variation. This result is in agreement with the expected consequences of spatial confounding highlighted below Equations (4a) and (4b) when it comes to the true residual spatial covariation and its relative role in explaining species responses. RSFA is partly able to recover the variance partitioning results seen in the independent factor analysis, although we note the restricted factors tend to explain slightly more in RSFA compared to their counterparts in independent factor analysis. This may be attributed to the $\mathbf{U}^*$ retaining some spatial structure due to their construction from $\mathbf{U}$ (see Equation 3). In doing so, they can more successfully capture residual spatial correlations within and between species that cannot be accounted for by the measured covariates.

Finally, Figure 6 presents the estimated values of the two latent factors as a function of spatial location. Both estimated spatial factors from SFA (columns of $\mathbf{U}$) exhibit a noticeable latitude gradient reminiscent of that of growing degree days, in agreement with the expected impacts of spatial confounding on residual ordination (see Section 3.3) and the simulation results from Figure 2. This is confirmed when we compute the Pearson correlation between each covariate and the predicted spatial factors, which shows that both spatial factors from SFA are strongly correlated with growing degree days (correlations of 0.819 and −0.635, respectively). There is also some correlation with the two other covariates, although to a lesser extent; see Appendix C for complete results including residual ordination plots where points are coloured by the covariate values. By contrast, the estimated factors from independent factor analysis present close to zero correlation with the three covariates, while the corresponding correlations from RSFA are exactly zero, as expected. In fact, if we compare the spatial maps of the estimated factors from independent factor analysis and RSFA (left and right columns of Figure 6), we observe that the first factor exhibits similar trends between the two models. Note however that the factors from RSFA

are spatially smoother than from independent factor analysis: This is consistent with the results from Figure 5 and can be explained by the fact that while the columns of $\mathbf{U}^*$ are restricted to be orthogonal to $\mathbf{X}$, they still retain spatial structure due to their construction from $\mathbf{U}$; see also Figure 2 in the simulation study where independent factor analysis performed poorly due to their failure to explicitly account for spatial correlation in the residual covariation.

In Appendix D, we present a second case study comparing independent factor analysis, SFA and RSFA on count records of $J = 15$ zooplankton species recorded across $N = 579$ locations using vessels traversing the Indian sector of the Southern Ocean. Overall, the results of this application again demonstrate the impacts of spatial confounding on quantifying covariate effects, variance partitioning and residual ordination, although the differences between SFA and RSFA are less dramatic compared with the results seen above for the butterfly data.

6 | DISCUSSION

Through a simulation study and two case studies, we have demonstrated the consequences of spatial confounding on three key applications of joint species distribution modelling, namely inferences and interpretations on covariates effects, quantifying between-species correlations and variance partitioning, and model-based residual ordination. We presented one approach to treating spatial confounding called restricted spatial factor analysis, which seeks to purposefully adjust the spatial factors so that they remain uncorrelated with the covariates. Simulation results suggest RSFA is relatively successful at achieving this and thus is arguably more in alignment with the intention of why the spatial random effects were included originally. Note although we fit the SFA and RSFA models via maximum approximate likelihood estimation using TMB (Kristensen et al., 2016), which is available at Hui et al. (2024), there are a number of existing JSDM packages in the literature employing either likelihood-based or Bayesian estimation that could be used to obtain results for RSFA from an already fitted SFA model (e.g. post-hoc processing of MCMC samples). This should facilitate assessments of spatial confounding and sensitivity analysis between SFA and RSFA models (although see the work of Hefley et al., 2017; Mäkinen et al., 2022; Marques & Wiemann, 2023, regarding the role of the prior distribution/hierarchical structure on the severity of spatial confounding).

In Section 3.3, we focused on the impacts of spatial confounding on residual ordination and its interpretation. However, these impacts also extend to the case of model-based unconstrained ordination (another major usage of JSDMs, Warton et al., 2015; van der Veen et al., 2023) when $\mathbf{X}$ contains only an intercept term and $\mathbf{B}$ reduces a set of species-specific intercepts. In this setting, spatial confounding can cause the unconstrained ordination from SFA to have a non-zero location, while RSFA makes an adjustment to ensure the unconstrained ordination is centred around the origin (consistent with how we expect the latent gradients representing community
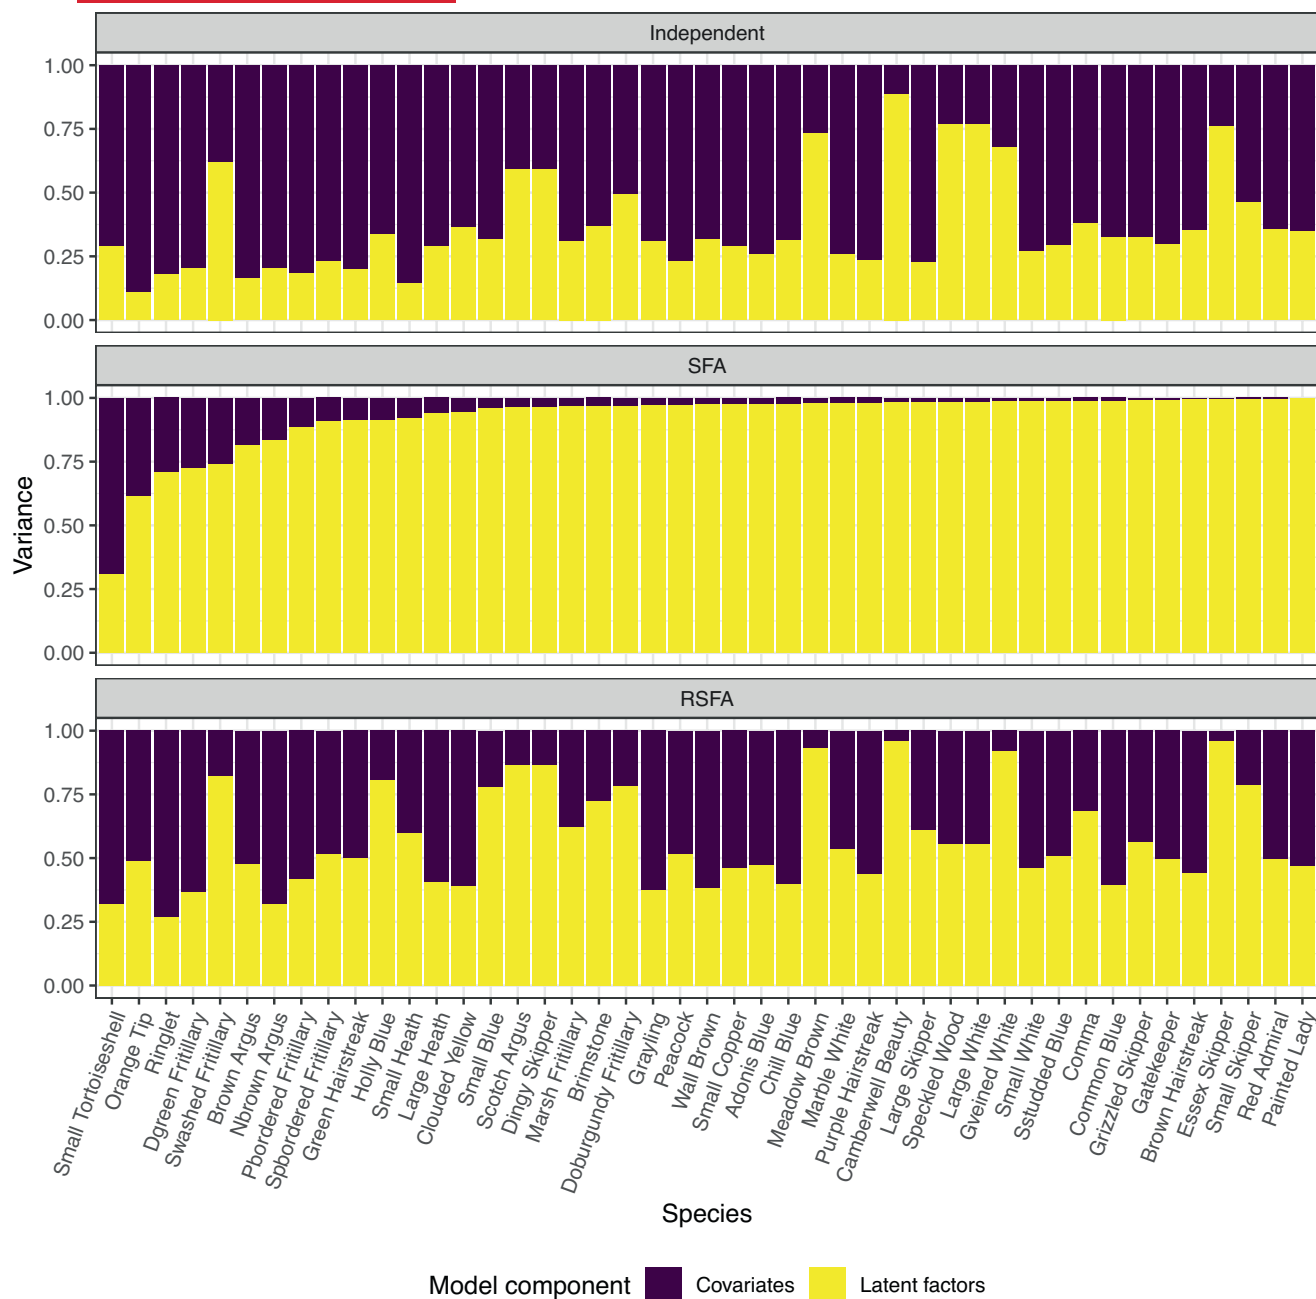
FIGURE 5 Variance partitioning for the $J = 45$ species in the Great Britain butterfly data, using independent factor analysis, SFA and RSFA. The species are ordered by the amount of variation explained by the three measured covariates in SFA.

composition to be set up). We illustrate an example of this using the Great Britain butterfly data in Appendix E.

As discussed in Section 1, RSFA is by no means the only method to treat spatial confounding. For single response models, some authors have argued against 'restricted-like' approaches due to potential issues with inflated Type I errors when assessing covariate effects (Khan & Calder, 2022; Zimmerman & Ver Hoef, 2022), although our simulation design did not find particularly strong evidence of this, and there are potential remedies to this drawback based on adjusting the standard errors (Hanks et al., 2015; Hui & Bondell, 2022). One reason why we view RSFA as being an appropriate practical solution to spatial confounding is that it still facilitates many of the

statistical inferences ecologists want to obtain from fitting JSDMs with little additional burden: As noted at the beginning of Section 4, stacked species distribution models do not allow for model-based ordination or variance partitioning. Also, while they have yet to be adapted to spatial multivariate abundance data, we expect that other methods to treat spatial confounding such as the approaches of Thaden and Kneib (2018), Dupont et al. (2022) and Marques et al. (2022) will encounter challenges when it comes to producing/interpreting ordinations, between-species correlations and variance partitioning, and/or require the ecologist to model the spatial structure of $\mathbf{x}_i$ as well (which is not something we typically want to undertake). More broadly, as we discussed below Equation (3), because
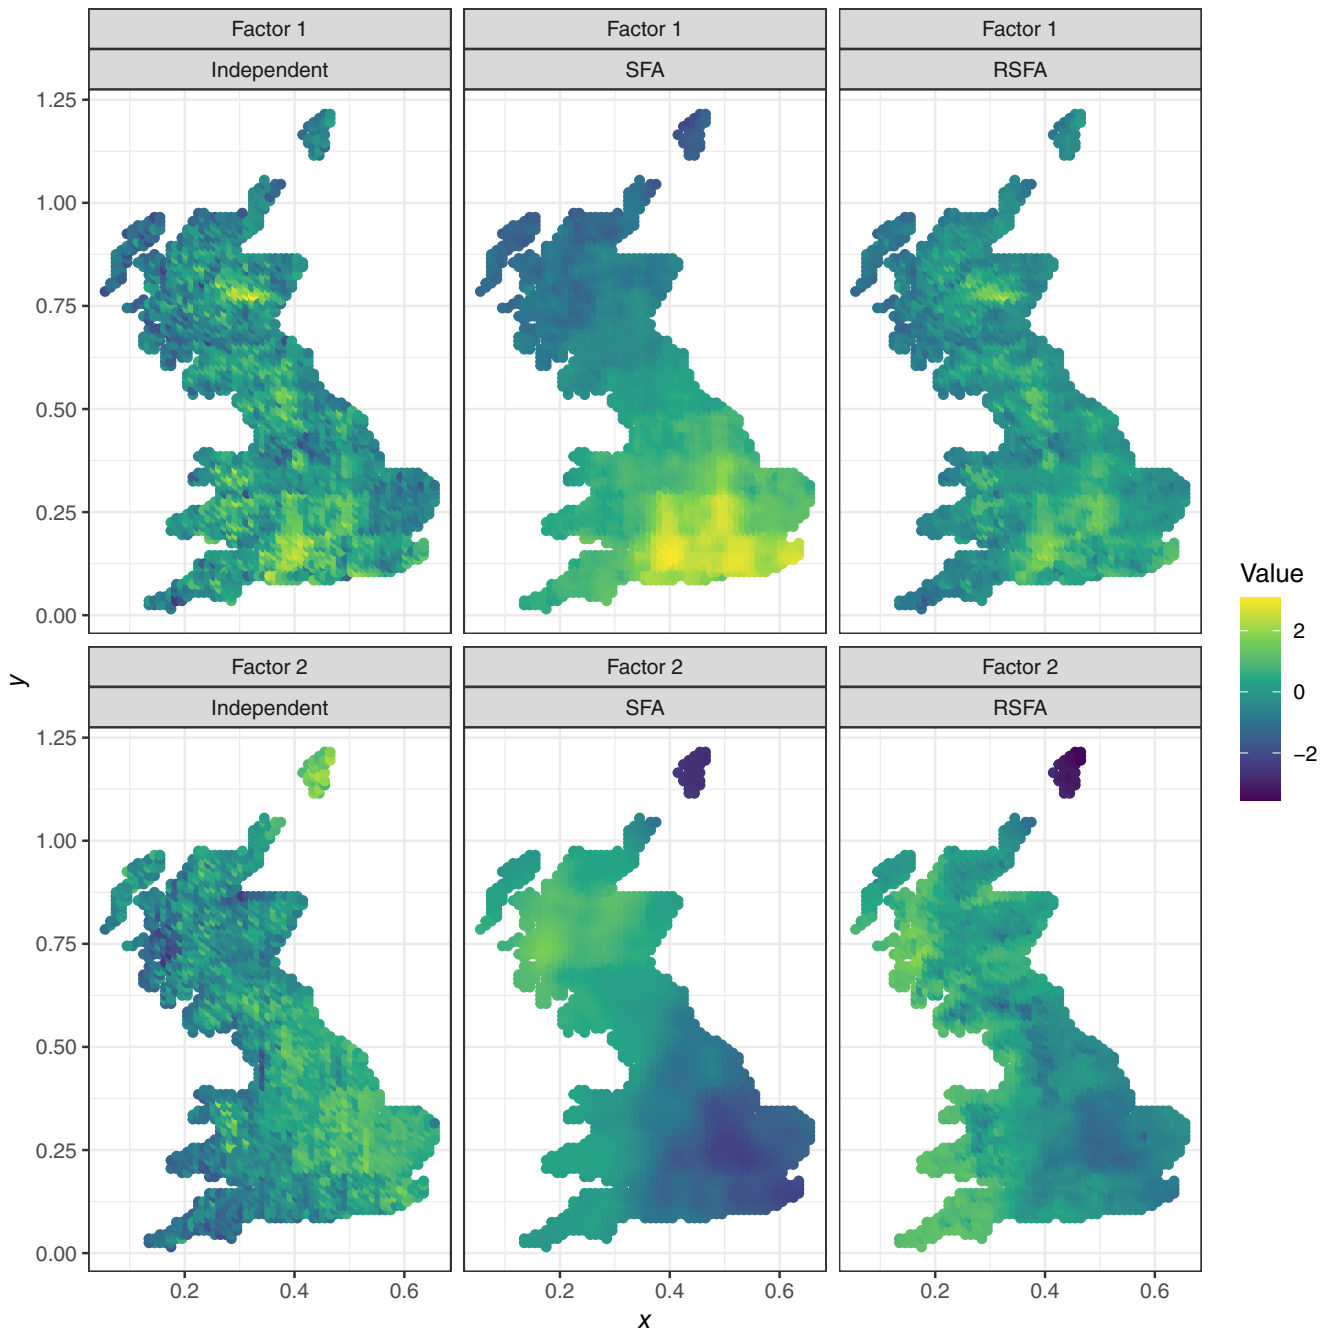
FIGURE 6 Spatial plots of the $K = 2$ latent factors for the Great Britain butterfly data, estimated using independent factor analysis, SFA and RSFA.

spatial confounding is analogous to correlation between measured covariates, then similar to multicollinearity as a whole it is not always evident how to treat spatial confounding. Thus, currently there is a lack of universal rules that can guide the management of spatial confounding. This is especially relevant here given RSFA treats spatial confounding by imposing orthogonality between the spatial factors and measured covariates, which is a contentious constraint (see also the discussion by Papadogeorgou, 2022).

Finally, in many applications of spatially structured JSDMs, practitioners want to incorporate more than just linear effects of the measured covariates to account for complex non-linear species-environmental relationships or incorporation of additional information such as traits (e.g. Niku et al., 2021). The presence, impact and way of treating (through RSFA) spatial confounding we have discussed in this article still applies in these settings. However, additional research is needed given many of these more flexible methods usually also involve some degree of penalisation to avoid overfitting/oversmoothing; how spatial confounding impacts the choice of smoothing and tuning parameter selection is much less understood (see Dupont et al., 2022, for a connected discussion). Related to this,

Hoeting et al. ([2006](#)) and Hoeting ([2009](#)) among others have shown that ignoring spatial correlation can have detrimental consequences (e.g. underfitting and higher prediction errors) on the variable selection of covariates. However, the impact of spatial confounding on the performance of formal model selection tools in JSDMs remains under-explored: (see for instance the results from Appendices [A](#) and [B](#), as well as Hefley et al., [2017](#); Yang et al., [2023](#), for related discussions in the context of a single response).

AUTHOR CONTRIBUTIONS

Francis K. C. Hui conceived the idea, implemented the methodology and performed the case studies. Francis K. C. Hui and Quan Vu designed and implemented the simulation study. Francis K. C. Hui and Mevin B. Hooten led the writing of the manuscript, while all authors contributed to the drafts and gave final approval for publication.

CONFLICT OF INTEREST STATEMENT

All authors have no conflict of interest to declare.

PEER REVIEW

The peer review history for this article is available at https://www.webofscience.com/api/gateway/wos/peer-review/10.1111/2041-210X.14420.

DATA AVAILABILITY STATEMENT

All code used as part of Sections [4](#) and [5](#) are archived at https://doi.org/10.5281/zenodo.13385532 (Hui et al., [2024](#)). The two datasets used in Section [5](#) and Appendix [D](#) are publicly available at https://doi.org/10.5281/zenodo.1452066 (Wilkinson, [2018](#)) and https://data.aad.gov.au/metadata/records/AADC-00099 (Kitchener & Hosie, [2024](#)), respectively.

ORCID

Francis K. C. Hui ![ORCID] https://orcid.org/0000-0003-0765-3533

REFERENCES

Asher, J., Warren, M., Fox, R., Harding, P., Jeffcoate, G., & Jeffcoate, S. (2001). *The millennium atlas of butterflies in Britain and Ireland*. Oxford University Press.

Astarloa, A., Louzao, M., Boyra, G., Martinez, U., Rubio, A., Irigoien, X., Hui, F. K. C., & Chust, G. (2019). Identifying main interactions in marine predator–prey networks of the Bay of Biscay. *ICES Journal of Marine Science*, 76, 2247–2259.

Borcard, D., Gillet, F., & Legendre, P. (2011). *Numerical ecology with R*. Springer.

Dormann, C. F., McPherson, J. M., Araújo, M. B., Bivand, R., Bolliger, J., Carl, G., Davies, R. G., Hirzel, A., Jetz, W., Daniel Kissling, W., Kühn, I., Ohlemüller, R., Peres-Neto, P. R., Reineking, B., Schröder, B., Schurr, F. M., & Wilson, R. (2007). Methods to account for spatial autocorrelation in the analysis of species distributional data: A review. *Ecography*, 30, 609–628.

Dupont, E., Marques, I., & Kneib, T. (2023). Demystifying spatial confounding. *arXiv preprint arXiv:2309.16861*.

Dupont, E., Wood, S. N., & Augustin, N. H. (2022). Spatial+: A novel approach to spatial confounding. *Biometrics*, 78, 1279–1290.

Graham, M. H. (2003). Confronting multicollinearity in ecological multiple regression. *Ecology*, 84, 2809–2815.

Haak, C. R., Hui, F. K. C., Cowles, G. W., & Danylchuk, A. J. (2020). Positive interspecific associations consistent with social information use shape juvenile fish assemblages. *Ecology*, 101, e02920.

Hanks, E. M., Schliep, E. M., Hooten, M. B., & Hoeting, J. A. (2015). Restricted spatial regression in practice: Geostatistical models, confounding, and robustness under model misspecification. *Environmetrics*, 26, 243–254.

Hefley, T. J., Hooten, M. B., Hanks, E. M., Russell, R. E., & Walsh, D. P. (2017). The Bayesian group lasso for confounded spatial data. *Journal of Agricultural, Biological and Environmental Statistics*, 22, 42–59.

Hodges, J. S., & Reich, B. J. (2010). Adding spatially-correlated errors can mess up the fixed effect you love. *The American Statistician*, 64, 325–334.

Hoeting, J. A. (2009). The importance of accounting for spatial and temporal correlation in analyses of ecological data. *Ecological Applications*, 19, 574–577.

Hoeting, J. A., Davis, R. A., Merton, A. A., & Thompson, S. E. (2006). Model selection for geostatistical models. *Ecological Applications*, 16, 87–98.

Hosie, G., Fukuchi, M., & Kawaguchi, S. (2003). Development of the Southern Ocean continuous plankton recorder survey. *Progress in Oceanography*, 58, 263–283.

Hughes, J., & Haran, M. (2013). Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *Journal of the Royal Statistical Society, Series B: Statistical Methodology*, 75, 139–159.

Hui, F. K. C., & Bondell, H. D. (2022). Spatial confounding in generalized estimating equations. *The American Statistician*, 76, 238–247.

Hui, F. K. C., Hill, N. A., & Welsh, A. H. (2022). Assuming independence in spatial latent variable models: Consequences and implications of misspecification. *Biometrics*, 78, 85–99.

Hui, F. K. C., Vu, Q., & Hooten, M. B. (2024). Spatial confounding in joint species distribution models. *Zenodo*. v 0.6 https://doi.org/10.5281/zenodo.13385532

Hui, F. K. C., Warton, D. I., Foster, S. D., & Haak, C. R. (2023). Spatiotemporal joint species distribution modelling: A basis function approach. *Methods in Ecology and Evolution*, 14, 2150–2164.

Khan, K., & Berrett, C. (2023). Re-thinking spatial confounding in spatial linear mixed models. *arXiv preprint arXiv:2301.05743*.

Khan, K., & Calder, C. A. (2022). Restricted spatial regression methods: Implications for inference. *Journal of the American Statistical Association*, 117, 482–494.

Kitchener, J., & Hosie, G. (2024). *Southern ocean continuous plankton recorder zooplankton records, Ver. 12*. Australian Antarctic Data Centre. https://doi.org/10.26179/r7sm-9y85

Kristensen, K., Nielsen, A., Berg, C. W., Skaug, H., & Bell, B. (2016). TMB: Automatic differentiation and Laplace approximation. *Journal of Statistical Software*, 70, 1–21.

Lany, N. K., Zarnetske, P. L., Finley, A. O., & McCullough, D. G. (2020). Complementary strengths of spatially-explicit and multi-species distribution models. *Ecography*, 43, 456–466.

Lembrechts, J. J., Nijs, I., & Lenoir, J. (2019). Incorporating microclimate into species distribution models. *Ecography*, 42, 1267–1279.

Mäkinen, J., Numminen, E., Niittynen, P., Luoto, M., & Vanhatalo, J. (2022). Spatial confounding in Bayesian species distribution modeling. *Ecography*, 2022, e06183.

Marques, I., Kneib, T., & Klein, N. (2022). Mitigating spatial confounding by explicitly correlating Gaussian random fields. *Environmetrics, 33,* e2727.

Marques, I., & Wiemann, P. F. (2023). Bayesian spatial+: A joint model perspective. *arXiv preprint arXiv:2309.05496.*

Mielke, K. P., Claassen, T., Busana, M., Heskes, T., Huijbregts, M. A., Koffijberg, K., & Schipper, A. M. (2020). Disentangling drivers of spatial autocorrelation in species distribution models. *Ecography, 43,* 1741–1751.

Niku, J., Hui, F. K. C., Taskinen, S., & Warton, D. I. (2021). Analyzing environmental-trait interactions in ecological communities with fourth-corner latent variable models. *Environmetrics, 32,* e2683.

Ovaskainen, O., & Abrego, N. (2020). *Joint species distribution modelling: With applications in R.* Cambridge University Press.

Ovaskainen, O., Roy, D. B., Fox, R., & Anderson, B. J. (2016). Uncovering hidden spatial structure in species communities with spatially explicit joint species distribution models. *Methods in Ecology and Evolution, 7,* 428–436.

Paciorek, C. J. (2010). The importance of scale for spatial-confounding bias and precision of spatial regression estimators. *Statistical Science, 25,* 107–125.

Papadogeorgou, G. (2022). Discussion on "spatial+: A novel approach to spatial confounding" by Emiko Dupont, Simon N. Wood, and Nicole H. Augustin. *Biometrics, 78,* 1305–1308.

Pollock, L. J., Tingley, R., Morris, W. K., Golding, N., O'Hara, R. B., Parris, K. M., Vesk, P. A., & McCarthy, M. A. (2014). Understanding co-occurrence by modelling species simultaneously with a joint species distribution model (JSDM). *Methods in Ecology and Evolution, 5,* 397–406.

Prates, M. O., Assunção, R. M., & Rodrigues, E. C. (2019). Alleviating spatial confounding for areal data problems by displacing the geographical centroids. *Bayesian Analysis, 14,* 623–647.

Schliep, E. M., Lany, N. K., Zarnetske, P. L., Schaeffer, R. N., Orians, C. M., Orwig, D. A., & Preisser, E. L. (2018). Joint species distribution modelling for spatio-temporal occurrence and ordinal abundance data. *Global Ecology and Biogeography, 27,* 142–155.

Schulz, T., Vanhatalo, J., & Saastamoinen, M. (2020). Long-term demographic surveys reveal a consistent relationship between average occupancy and abundance within local populations of a butterfly metapopulation. *Ecography, 43,* 306–317.

Shirota, S., Gelfand, A. E., & Banerjee, S. (2019). Spatial joint species distribution modeling using Dirichlet processes. *Statistica Sinica, 29,* 1127–1154.

Stoklosa, J., Blakey, R. V., & Hui, F. K. C. (2022). An overview of modern applications of negative binomial modelling in ecology and biodiversity. *Diversity, 14,* 320.

Thaden, H., & Kneib, T. (2018). Structural equation models for dealing with spatial confounding. *The American Statistician, 72,* 239–252.

Thorson, J. T., Adams, G., & Holsman, K. (2019). Spatio-temporal models of intermediate complexity for ecosystem assessments: A new tool for spatial fisheries management. *Fish and Fisheries, 20,* 1083–1099.

Thorson, J. T., Ianelli, J. N., Larsen, E. A., Ries, L., Scheuerell, M. D., Szuwalski, C., & Zipkin, E. F. (2016). Joint dynamic species distribution models: A tool for community ordination and spatio-temporal monitoring. *Global Ecology and Biogeography, 25,* 1144–1158.

Thorson, J. T., Scheuerell, M. D., Shelton, A. O., See, K. E., Skaug, H. J., & Kristensen, K. (2015). Spatial factor analysis: A new tool for estimating joint species distributions and correlations in species range. *Methods in Ecology and Evolution, 6,* 627–637.

Tikhonov, G., Duan, L., Abrego, N., Newell, G., White, M., Dunson, D., & Ovaskainen, O. (2020). Computationally efficient joint species distribution modeling of big spatial data. *Ecology, 101,* e02929.

Tzeng, S., & Huang, H.-C. (2018). Resolution adaptive fixed rank kriging. *Technometrics, 60,* 198–208.

van der Veen, B., Hui, F. K. C., Hovstad, K. A., & O'Hara, R. B. (2023). Concurrent ordination: Simultaneous unconstrained and constrained latent variable modelling. *Methods in Ecology and Evolution, 14,* 683–695.

Van Ee, J. J., Ivan, J. S., & Hooten, M. B. (2022). Community confounding in joint species distribution models. *Scientific Reports, 12,* 1–14.

Velleman, P. F., & Welsch, R. E. (1981). Efficient computing of regression diagnostics. *The American Statistician, 35,* 234–242.

Warton, D. I., Blanchet, F. G., O'Hara, R. B., Ovaskainen, O., Taskinen, S., Walker, S. C., & Hui, F. K. C. (2015). So many variables: Joint modeling in community ecology. *Trends in Ecology & Evolution, 30,* 766–779.

Wilkinson, D. (2018). Doi90/jsdm_inference: Publication. *Zenodo.* v 0.1.6 https://doi.org/10.5281/zenodo.1452066

Yang, H.-D., Chiou, Y.-H., & Chen, C.-S. (2023). Estimation and selection for spatial confounding regression models. *Communications in Statistics - Theory and Methods, 52,* 939–955.

Yuan, Y., Bachl, F. E., Lindgren, F., Borchers, D. L., Illian, J. B., Buckland, S. T., Rue, H., & Gerrodette, T. (2017). Point process models for spatio-temporal distance sampling data from a large-scale survey of blue whales. *Annals of Applied Statistics, 11,* 2270–2297.

Zimmerman, D. L., & Ver Hoef, J. M. (2022). On deconfounding spatial confounding in linear models. *The American Statistician, 76,* 159–167.

SUPPORTING INFORMATION

Additional supporting information can be found online in the Supporting Information section at the end of this article.

Appendix A. Discussion on finite sample correlations.

Appendix B. Details and additional results of simulation study.

Appendix C. Additional results for application to Great Britain butterfly data.

Appendix D. Case study-zooplankton records in the SOCPR survey.

Appendix E. Impacts on model-based unconstrained ordination.