Spatial Covariate Balancing Propensity Scores for Causal Inference with Latent Spatial Confounding

Nathan B. Wikle[1,*] and Corwin M. Zigler[2]

[1]Department of Statistics and Actuarial Science, University of Iowa
[2]Department of Biostatistics, Brown University

June 11, 2026

Abstract

Unmeasured spatial confounding is a major obstacle to causal inference with spatial data, particularly in environmental and public health applications where both exposures and outcomes may be influenced by complex spatial processes. We introduce a spatial covariate balancing propensity score (spatial CBPS) that adapts modern design-based balancing weight methods to settings with latent spatial confounding. The method approximates unobserved spatial confounders using low-rank spatial basis functions and estimates propensity scores through a covariate balancing objective function, prioritizing balance in spatial structure rather than likelihood-based model fit. This yields inverse probability weights with improved finite sample behavior and straightforward diagnostics, and remains compatible with doubly robust estimators. We develop practical inference procedures under spatial dependence using spatial HAC and block bootstrap variance estimators, and demonstrate in simulations that spatial CBPS substantially outperforms likelihood-based spatial propensity score methods across a range of confounding regimes. Finally, we analyze the effect of proximity to concentrated animal feeding operations on nitrate concentrations in Iowa private wells. Although estimated effects are not statistically significant, point estimates differ substantially between likelihood-based spatial propensity score methods and spatial CBPS, suggesting that explicitly balancing spatial structure provides more reliable adjustment for spatial confounding.

1 Introduction

A central challenge in estimating causal effects from spatial data is the threat of *unmeasured spatial confounding*, in which latent, spatially-structured variables affect both an exposure and outcome of interest. Without properly accounting for the unmeasured confounders, standard ignorability assumptions fail and causal identification is generally no longer possible (Gilbert et al. 2024). This problem is especially pronounced in environmental settings where common spatial exposures, such as air and water pollution or extreme heat exposure, arise from complex geophysical processes that are hard to fully observe (Papadogeorgou et al. 2019, Tec et al. 2023).

*Email correspondence: nathan-wikle@uiowa.edu

To address this problem, a growing literature has focused on using spatial proximity as a surrogate for latent spatial confounding (Paciorek 2010). Informally, when the confounding process is assumed to be sufficiently smooth in space, residual spatial structure in the observed exposure or outcome variable may serve as a proxy for the unmeasured confounder, allowing certain estimators to adjust for the corresponding confounding bias. To date, this literature has largely emphasized *model-based approaches* for confounding adjustment, typically in settings with continuous exposure and outcome variables (Thaden & Kneib 2018, Schnell & Papadogeorgou 2020, Dupont et al. 2022, Guan et al. 2022, Marques et al. 2023, Dupont et al. 2025, Rainey & Keller 2025, Datta & Stein 2025, Lamouroux et al. 2025). This emphasis reflects both the long-standing role of the spatial linear mixed model in spatial and environmental statistics (Banerjee et al. 2015, Cressie & Wikle 2011) and close connections to work on fixed effect estimation in settings with collinearity between observed covariates and spatial random effects (Hodges & Reich 2010, Hanks et al. 2015, Khan & Calder 2022, Zimmerman & Ver Hoef 2022).

In contrast, *design-based approaches* for confounding adjustment have received comparatively little attention in the spatial confounding literature (Papadogeorgou et al. 2019). This gap is surprising given the central role of design-based methods, such as propensity score matching and weighting, in modern causal inference (Zubizarreta et al. 2023). These approaches seek to mitigate confounding by constructing pseudo-populations of exposed and unexposed units with similar distributions of confounders prior to any outcome analysis, thereby reducing the reliance on correct outcome model specification and enabling transparent diagnostics of key identifying assumptions, such as covariate balance and treatment-control overlap.

In this paper, we introduce covariate balancing propensity score methods to the spatial confounding setting. We propose a spatially informed covariate balancing propensity score (spatial CBPS) that directly targets unmeasured confounding in settings with spatially indexed data. The approach leverages the natural correspondence between basis function representations of latent spatial processes (Cressie et al. 2022) and the moment conditions defining modern balancing weight methods (Chattopadhyay et al. 2020). When unmeasured confounding arises from a smooth, low-frequency spatial process, enforcing balance on a suitably chosen set of spatial basis functions induces approximate balance in the unobserved confounder itself. This yields a flexible, design-based analogue to common model-based adjustment strategies, while accommodating a wide range of outcome types (binary, continuous, counts) and spatial support, including both point-referenced and areal data.

We show that this perspective leads to weighting estimators of the average treatment effect with substantially improved finite sample performance relative to likelihood-based spatial propensity score models, while remaining fully compatible with popular doubly robust estimators. When combined with low-rank spatial basis representations of the unmeasured confounder, estimating the proposed spatial CBPS is computationally efficient. Furthermore, in settings with spatial dependence, uncertainty can be quantified using either spatial heteroskedasticity and autocovariance (HAC) (Kelejian & Prucha 2007) or spatial block bootstrap (Lahiri 2003) variance estimators. To facilitate ease of use, we provide an accompanying R package, `spbalance`, which implements the spatial CBPS approach for use with several common spatial basis functions.

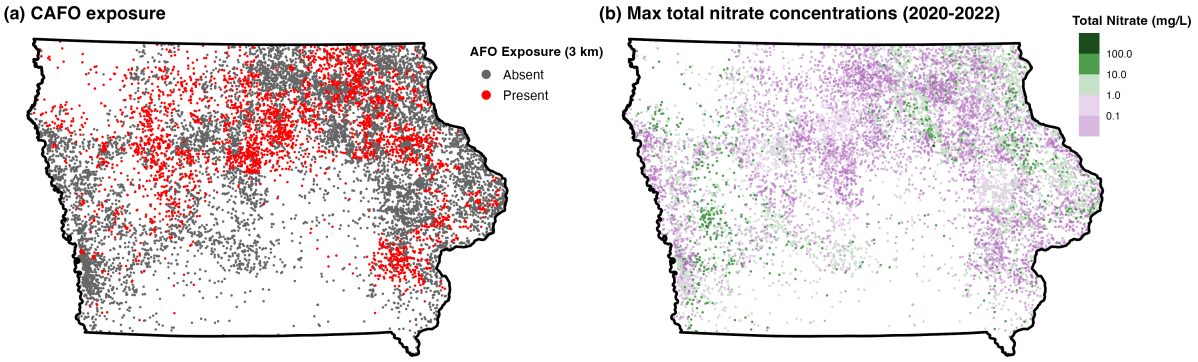

Figure 1: Map of private wells in Iowa that were tested between 2020–2022, showing (a) treatment exposure, with red indicating that the well was in close proximity (within 3 km) of a CAFO, and (b) the max total nitrate concentration recorded at that location. For improved visualization, the outcome data are binned into five categories, but the analysis uses continuous values.

To illustrate the practical value of a design-based approach to spatial confounding adjustment, we study the effect of concentrated animal feeding operations (CAFOs) on groundwater quality in Iowa. Iowa has experienced substantial growth and consolidation in livestock production over the past several decades, with hog production in particular driving the increase in CAFOs (Jones et al. 2019). The increase in CAFOs has been met with growing concern over their environmental and health impacts, and they have been associated with an increased risk of respiratory disease, infection, and waterborne illness (Son et al. 2024) as well as poor air (Burns et al. 2023, Epps et al. 2025) and water quality (Jones et al. 2019, Wheeler et al. 2015). Of particular concern is the risk of groundwater contamination from the application of manure from CAFOs as fertilizer in the surrounding region. In particular, presence or absence of CAFOs has been identified as an important predictor of elevated total nitrate levels in private wells in Iowa (Wheeler et al. 2015). At the same time, elevated nitrate in drinking water is a longstanding public health concern (Ward et al. 2018), and both CAFO siting decisions and groundwater vulnerability are shaped by complex, spatially structured processes that are only partially observed. These features make the setting a natural candidate for spatial confounding: proximity to CAFOs is not randomized across the landscape, and important determinants of both exposure and well water nitrate may vary smoothly over space.

A primary objective of this article is to employ our spatial CBPS estimator to investigate whether CAFO proximity has a detrimental effect on max total nitrate concentrations in groundwater found in private wells in Iowa. Using well water data obtained from the Iowa Private Well Forecasting System (Iowa DNR 2025*b*) and spatial data on active feeding operations from the Iowa DNR (Iowa DNR 2025*a*), we define exposure based on whether a CAFO operates within close proximity (3 km) of a tested well; CAFO exposure and the corresponding outcome of interest are depicted in Figure 1. Our analysis does not find evidence that close proximity to CAFOs has a statistically significant effect on well water nitrate concentration, however, the analysis

highlights that conclusions about environmental health impacts can depend meaningfully on how propensity score methods adjust for spatial confounding.

The remainder of the paper is structured as follows. Section 2 introduces the causal framework and formalizes spatial confounding assumptions. Section 3 develops the spatial CBPS estimator, provides a bias decomposition that clarifies its advantages over likelihood-based spatial propensity score models, and outlines inference under spatial dependence using spatial HAC and block bootstrap variance estimators. Section 4 presents a simulation study comparing spatial CBPS with alternative approaches across multiple spatial confounding regimes. Section 5 applies the proposed methods to estimate the effect of CAFO proximity on nitrate levels in Iowa private wells. Finally, Section 6 concludes with a discussion of practical guidelines, limitations, and extensions of the proposed spatial covariate balancing propensity score.

2 Causal Framework and Spatial Confounding

2.1 Notation, Assumptions, and Estimand of Interest

Let $\mathcal{D} \subset \mathbb{R}^2$ denote the spatial domain of interest. We will assume throughout that the data are point-referenced (i.e., $\mathcal{D}$ is a continuous domain), however it is straightforward to extend our methodology to areal settings. Suppose we observe data at locations $D_n = \{s_1, \ldots, s_n\} \subset \mathcal{D}$, where $n = |D_n|$ may increase with the study size; for notational simplicity we suppress any explicit dependence of D_n on n beyond its cardinality. For each $s \in D_n$, we observe $Z_s = (Y_s, A_s, \mathbf{L}_s)$. The components of Z_s include an observed outcome, Y_s, a binary exposure variable, A_s, and a set of covariates, $\mathbf{L}_s$, which we assume (for now) to be completely observed. In our example, these correspond to well water nitrate concentration, CAFO presence or absence, and observed well and watershed-level covariates (e.g., well depth, soil characteristics, land use summaries, climatological features, etc.), respectively.

We assume that each location s has two potential outcomes (Neyman 1923, Rubin 1974), $Y_s(0)$ and $Y_s(1)$, representing the outcomes that would be observed at location s under exposure levels $A_s = 0$ and $A_s = 1$, respectively. Implicit in this notation is the assumption of no interference, which requires that a location's potential outcomes depend only on its own exposure level and not on the exposures at other spatial locations, and that there are no hidden versions of treatment (i.e., the potential outcome $Y_s(a)$ is unambiguously defined for $a = 0$ and $a = 1$). These conditions are commonly referred to as the Stable Unit Treatment Value Assumption (SUTVA, Rubin 1980):

Assumption 1 (SUTVA)**.** *For each spatial location $s \in \mathcal{D}$, (i) $Y_s(a)$ depends only on local treatment assignment (no interference); and (ii) there are no hidden versions of treatment.*

Whether SUTVA is plausible must be carefully evaluated in spatial settings (Reich et al. 2021) — notably, it rules out the possibility of spatial interference, in which the outcome at a spatial location s is affected by exposures assigned to nearby locations. In environmental applications this assumption may be more or less plausible depending on the physical mechanism of exposure; see Section 6 for a discussion of settings where SUTVA may fail and how it may be relaxed. For now, we assume SUTVA to hold and define causal estimands via contrasts of the

potential outcomes $Y_s(0)$ and $Y_s(1)$. In particular, we focus our attention on the estimation of the average treatment effect (ATE),

$$\tau = E\big[Y_s(1) - Y_s(0)\big]. \tag{1}$$

In our motivating example, the ATE quantifies the average change in groundwater nitrate concentration that would be observed in private wells in Iowa if, counterfactually, each well location in the study region were exposed to a CAFO ($A_s = 1$) versus not exposed ($A_s = 0$), marginalizing over the distribution of locations (and associated covariates) in $\mathcal{D}$.

Because only one potential outcome is ever observed for any unit in D_n, estimating the ATE (or any other functional of the distributions of the potential outcomes) from the observed data requires additional assumptions beyond SUTVA. Specifically, we assume consistency, conditional ignorability, and positivity.

Assumption 2 (Causal Identification)**.** *For each spatial location $s \in \mathcal{D}$, we assume*

(a) Consistency: $Y_s = Y_s(a)$ if $A_s = a$.

(b) Conditional Ignorability: $\{Y_s(0), Y_s(1)\} \perp\!\!\!\perp A_s \,\big|\, \mathbf{L}_s$.

(c) Positivity: There exists $\varepsilon > 0$ such that

$$\varepsilon \;<\; P(A_s = 1 \mid \mathbf{L}_s = \boldsymbol{\ell}) \;<\; 1 - \varepsilon,$$

for all $\boldsymbol{\ell}$ in the support of $\mathbf{L}_s$.

Consistency states that the outcome observed at location s is equal to the potential outcome corresponding to the location's observed exposure level, and is a direct consequence of SUTVA. Conditional ignorability, also commonly referred to as exchangeability or the assumption of no unmeasured confounding, states that treatment assignment is essentially randomized within covariate strata. In our observational setting, this means that $\mathbf{L}_s$ contains all potential confounding variables which impact both the exposure level and outcome at location s, and that conditional on these measured confounders, the exposure variable is essentially randomized in the sense that it is independent of the distribution of potential outcomes. Finally, the positivity assumption states that each location in our study has some positive probability of receiving either exposure level. Assumption 2 allows the ATE to be nonparametrically identified from the observed data, and a variety of methods have been developed to estimate the ATE in standard (i.e., non-spatial) settings (Imbens & Rubin 2015).

2.2 Unmeasured Spatial Confounding

In settings with complex spatially varying exposures the conditional exchangeability assumption can be difficulty to justify. Instead, a more realistic characterization of confounding is given by

$$\{Y_s(0), Y_s(1)\} \perp\!\!\!\perp A_s \,\big|\, \mathbf{L}_s = (\mathbf{X}_s, U_s), \tag{2}$$

where $\mathbf{X}_s$ and U_s denote observed and unobserved covariates, respectively. If Assumption 2(b) *fails* when U_s is omitted, we say that U_s is an unmeasured confounder. When this occurs, it is generally impossible to estimate causal effects without additional assumptions on the structure of U_s (Gilbert et al. 2024).

However, in many settings with spatially-varying exposures it may be reasonable to assume that the unmeasured confounder varies *smoothly* in space. This assumption offers a distinct advantage — if U_s is suitably smooth relative to A_s, the unmeasured confounder may be approximated as a smooth function of space. Following the approach of Gilbert et al. (2024), this idea can be formalized as follows.

Assumption 3 (Spatial Confounding)**.** *For all $s \in \mathcal{D}$, we assume*

(a) Latent spatial confounding: $U_s = f(s)$ for some fixed, measurable function $f : \mathcal{D} \to \mathbb{R}$.

(b) No induced confounding: $\{Y_s(0), Y_s(1)\} \perp\!\!\!\perp A_s \,\big|\, \mathbf{X}_s, U_s, s$.

(c) Spatial positivity: There exist constants $\varepsilon > 0$ such that

$$\varepsilon < P(A_s = 1 \mid \mathbf{X}_s = \mathbf{x}) < 1 - \varepsilon,$$

for all $\mathbf{x}$ in the support of $\mathbf{X}_s$ and all $s \in \mathcal{D}$.

Assumption 3 states that (a) unmeasured confounding arises through a spatially structured latent process that can be represented as a deterministic function of space, (b) conditioning on spatial location s in addition to $\mathbf{L}_s = (\mathbf{X}_s, U_s)$ does not induce unwanted confounding, and (c) treatment assignment exhibits sufficient variation across space, in the sense that the spatially varying conditional probability of treatment is uniformly bounded away from zero and one at all spatial locations in $\mathcal{D}$. In practice, positivity is required to hold at spatial scales over which the latent confounder $f(s)$ may be regarded as approximately constant. This mirrors previous assertions in the spatial confounding literature that the exposure varies at a smaller spatial scale than the unmeasured confounder (Paciorek 2010, Schnell & Papadogeorgou 2020, Gilbert et al. 2024, Khan & Berrett 2024).

Together, Assumptions 1–3 imply causal identifiability in settings with spatial confounding (Gilbert et al. 2024), and suggest a general approach for estimating the ATE: adjust for confounding bias using spatial location, s, as a proxy for U_s.

2.3 Design-Based Inference with Spatial Confounding

While much of the spatial confounding literature has focused on continuous exposures, a smaller body of work has developed estimators for binary treatments that leverage spatial proximity to mitigate unmeasured confounding. Perhaps the most straightforward approach is to specify an outcome model, $\mu_s(a, \mathbf{x}) := E[Y_s \mid A_s = a, \mathbf{X}_s = \mathbf{x}]$, that includes a flexible spatial effect term (e.g., a spatial random effect or penalized spline), with the goal of absorbing residual spatial variation attributable to U_s (Reich et al. 2021). If the outcome model is correctly specified, estimation can proceed by marginalizing over the distribution of confounders (i.e., g-computation, Robins 1986). One drawback to this approach is that it relies heavily on the

outcome model being correctly specified, which may be challenging when the number of confounding variables is large.

In contrast, design-based estimators, such as weighting and matching methods, are often appealing because they separate the "design" stage (constructing comparable exposed and unexposed groups) from the analysis stage, limiting the risk of model misspecification and yielding transparent diagnostics of covariate balance and overlap. Central to this work is the propensity score, $e(\boldsymbol{\ell}) = P(A_s = 1 \mid \mathbf{L}_s = \boldsymbol{\ell})$, defined as the conditional probability of treatment given confounding variables $\mathbf{L}_s$. First introduced by Rosenbaum & Rubin (1983), the propensity score's utility can be best understood via its balancing score property: the conditional ignorability assumption 2(b) holds when conditioning on the propensity score in place of the full set of confounders. In other words, the conditional probability of treatment given confounders contains sufficient information to perform confounding adjustment, a property which is especially useful when the dimension of $\mathbf{L}_s$ grows large.

The propensity score has been adapted to the spatial confounding setting in several ways. Papadogeorgou et al. (2019) define a distance adjusted propensity score matching algorithm to estimate the average treatment effect on the treated (ATT), in which the propensity score is estimated conditional on the observed covariates $\mathbf{X}_s$ and 1:1 caliper matching is used to find treatment-control pairs that jointly prioritizes units with similar propensity score values and spatial proximity to each other. However, this approach suffers from common matching estimator drawbacks: it can be challenging to extend to general causal estimands and may prove inefficient if observations are discarded.

An alternative class of design-based estimators relies on inverse probability weighting (IPW), in which units are reweighted by the inverse of their estimated propensity score to construct a pseudo-population in which treatment assignment is independent of observed confounders. The standard IPW estimator is given as

$$\hat{\tau}_{IPW} = \frac{1}{w_1} \sum_{s \in D_n} \frac{A_s Y_s}{\hat{e}(\mathbf{L}_s)} - \frac{1}{w_0} \sum_{s \in D_n} \frac{(1 - A_s) Y_s}{1 - \hat{e}(\mathbf{L}_s)}, \tag{3}$$

where $\hat{e}_s(\mathbf{L}_s)$ is the estimated propensity score. The weighted sums in (3) estimate $E[Y_s(1)]$ and $E[Y_s(0)]$, respectively, with weights proportional to the inverse probability of treatment. When $w_1 = w_0 := n$, (3) corresponds to the Horvitz–Thompson (HT) estimator, while choosing $w_1 = \sum_s A_s/\hat{e}(\mathbf{L}_s)$ and $w_0 = \sum_s (1 - A_s)/(1 - \hat{e}(\mathbf{L}_s))$ yields the Hájek estimator, reflecting their connection to estimators from the survey sampling literature (Horvitz & Thompson 1952, Hájek 1971). Under correct specification of the propensity score, the IPW estimator is consistent for the ATE (Lunceford & Davidian 2004).

Despite its conceptual simplicity, the performance of the IPW estimator depends on the quality of the propensity score model and can be sensitive to extreme weights in finite samples (Kang & Schafer 2007). To address these limitations, doubly robust estimators, such as the augmented IPW (AIPW) estimator, jointly utilize propensity score and outcome regression models to construct efficient estimators of the ATE that are less susceptible to model misspecification (Robins et al. 1994). For example, the AIPW estimator can be viewed as an IPW estimator with an additional outcome regression adjustment term, $\mu(a, \boldsymbol{\ell})$, that corrects for potential

misspecification of the propensity score model. The AIPW estimator is given by

$$\hat{\tau}_{AIPW} = \frac{1}{n}\sum_{s\in D_n}\left[\frac{A_s Y_s}{\hat{e}(\mathbf{L}_s)} - \frac{A_s - \hat{e}(\mathbf{L}_s)}{\hat{e}(\mathbf{L}_s)}\hat{\mu}(1,\mathbf{L}_s)\right] - \frac{1}{n}\sum_{s\in D_n}\left[\frac{(1-A_s)Y_s}{1-\hat{e}(\mathbf{L}_s)} + \frac{A_s - \hat{e}(\mathbf{L}_s)}{1-\hat{e}(\mathbf{L}_s)}\hat{\mu}(0,\mathbf{L}_s)\right], \quad (4)$$

and is consistent if either the propensity score or outcome regression is correctly specified (Lunceford & Davidian 2004).

IPW and AIPW estimators have been used in the spatial confounding setting with varying success. Perhaps the simplest approach is to model the propensity score with a spatial generalized linear mixed model (GLMM) or generalized additive model (GAM), e.g.,

$$A_s \mid \mathbf{X}_s = \mathbf{x} \sim \text{Bernoulli}\big(e(\mathbf{x}, s)\big), \quad l\big(e(\mathbf{x}, s)\big) = \mathbf{x}^\intercal \boldsymbol{\beta} + f(s), \quad (5)$$

where l is a suitable link function (e.g., logit) and $f(s)$ is modeled as a spatial random effect or penalized spline. For example, Davis et al. (2019) and Reich et al. (2021) consider AIPW estimators where the propensity score model includes a spatial random effect that is assigned an intrinsic conditional autoregressive (ICAR) prior. Ideally, the spatial propensity score model results in a consistent estimator of the true probability of treatment (conditional on observed *and* unobserved confounders), and the advantages of the IPW and AIPW estimators are retained. Davis et al. (2019) show that ignoring latent dependence in the propensity score model can lead to substantial bias in the AIPW estimates, particularly when outcome models are misspecified, while Cohn (2024) shows that the valuable semiparametric properties of AIPW estimators (namely, double robustness and semiparametric efficiency) are retained when the observed spatial data are weakly dependent.

However, while IPW and AIPW estimators offer attractive theoretical guarantees under correct specification of the propensity score model, their practical performance in spatial settings remains sensitive to how spatial structure is modeled. In particular, when unmeasured confounding is driven by latent spatial processes, likelihood-based models for the propensity score may struggle to adequately capture low-frequency spatial variation without strong modeling assumptions. Moreover, even when flexible machine learning methods are used to estimate nuisance functions, weighting estimators based on estimated propensity scores can suffer from instability and poor finite-sample performance if overlap is limited at relevant spatial scales. For example, Tec et al. (2023) consider a spatial confounding process in which an IPW estimator with a spatial GLMM propensity score fails drastically. Similarly, Reich et al. (2021) find spatial AIPW estimators of the ATE to offer only a moderate reduction in bias compared to alternatives, such as a hierarchical Bayesian model of the full data generating process.

These considerations motivate an alternative design-based approach that places covariate balance, rather than propensity score model fit, at the center of confounding adjustment. Covariate balancing methods seek to directly enforce balance in observed confounders through a choice of treatment-control weights, bypassing the need to correctly specify a treatment assignment model (Chattopadhyay et al. 2020, Ben-Michael et al. 2021). In the presence of spatial confounding, this perspective suggests that obtaining balance of suitably chosen spatial features — such as spatial basis representations of U_s — may provide a more direct and transparent

route to mitigating bias from latent spatial structure than likelihood-based propensity score modeling. In the next section, we build on this idea by introducing a spatial covariate balancing propensity score estimator designed to target unmeasured spatial confounding while retaining the interpretability and diagnostic advantages of design-based inference.

3 Spatial Covariate Balancing Propensity Scores

The success of the IPW methods in Section 2 can be viewed through the lens of *covariate balance*. Informally, if the distribution of confounding variables is the same for treatment and control units, then the average difference in outcomes between the two groups can be attributed to the treatment. When the propensity score is known, Assumption 2 guarantees that the inverse propensity weights satisfy a *population balance property*:

$$E\left[\frac{A_s}{e(\mathbf{L}_s)}g(\mathbf{L}_s)\right] = E\big[g(\mathbf{L}_s)\big] \quad \text{and} \quad E\left[\frac{1-A_s}{1-e(\mathbf{L}_s)}g(\mathbf{L}_s)\right] = E\big[g(\mathbf{L}_s)\big], \quad \text{for all bounded } g(\cdot). \tag{6}$$

In fact, the inverse propensity score weights are the unique weighting function that satisfies (6) for all bounded functions g of the confounders (Ben-Michael et al. 2021). When the propensity score is estimated as in (5), the law of large numbers suggests that sample-based analogues to (6) should hold for large n. However, in practice, covariate balance may be quite poor when the propensity score is misspecified or the sample size is small; furthermore, the corresponding weights may be highly variable and produce unstable estimators of the ATE (Kang & Schafer 2007, Chattopadhyay et al. 2020).

This perspective suggests an alternative approach to inference: rather than estimating a propensity score model via the usual Bernoulli likelihood in (5), we instead estimate the propensity score under the requirement that certain prespecified covariate balance conditions are satisfied. With origins in survey sampling (Folsom 1991, Deville & Särndal 1992), these balancing weight estimators have received considerable interest in causal inference in non-spatial settings (Hainmueller 2012, Imai & Ratkovic 2014, Zubizarreta 2015, Zhao 2019, Wong & Chan 2017, Tan 2019, Hazlett 2020, Hirshberg & Wager 2021, Bruns-Smith et al. 2025). In the spatial confounding setting, this approach leads naturally to a *spatial covariate balancing propensity score (CBPS)* which directly targets covariate balance in both the measured and unmeasured confounders.

3.1 Spatial CBPS via Basis Function Representations of U_s

We motivate the spatial CBPS estimator as follows. Under Assumption 3, unmeasured confounding arises through a latent spatial process, $U_s \equiv f(s)$, that varies smoothly over the spatial domain $\mathcal{D}$. While f is unknown, we assume it belongs to a class of smooth functions, $\mathcal{F}$, that can be well approximated by finite-dimensional *spatial basis functions* (Cressie et al. 2022). Consequently, enforcing balance on a sufficiently rich set of basis functions spanning a subspace of $\mathcal{F}$ yields approximate balance in the latent confounder itself.

To operationalize this idea, we approximate f using a basis function expansion of the form,

$$f(s) \approx \sum_{k=1}^{K} \alpha_k \, b_k(s), \tag{7}$$

where $\{b_k(\cdot) : k = 1, 2, \ldots, K\}$ denotes a finite basis and $\{\alpha_k : k = 1, 2, \ldots, K\}$ are the corresponding coefficients. Each basis function $b_k(\cdot)$ generates a particular spatially-varying feature describing some component of the underlying variation in f, and f is approximated via a linear combination of these features. The basis functions may be defined by truncations of a countable basis, or they may simply be functions thought to be important for representing spatial variability. The specification of spatial basis functions is a central focus in spatial and environmental statistics, and a variety of functions have been proposed, including splines (Wahba 1990), wavelets (Cressie & Johannesson 2008), Wendland functions (Nychka et al. 2015), empirical orthogonal functions (Cressie & Wikle 2011), and finite elements (Lindgren et al. 2011); see Cressie et al. (2022) for a comprehensive review. In practice, the choice of basis function specification and rank K is often dictated by (i) the spatial support of the data, (ii) prior beliefs regarding the smoothness of the spatial surface, and (iii) computational considerations.

Given the representation in (7), the task of balancing the unobserved confounder U_s can be recast as that of balancing the *observable* spatial basis functions. In particular, enforcing balance on the individual basis function elements ensures balance of any linear combination of those basis elements, i.e.,

$$E\left[\frac{A_s}{e(\mathbf{X}_s, s)} b_k(s)\right] = E\big[b_k(s)\big], \quad \text{for all } k = 1, \ldots, K \tag{8}$$

implies

$$E\left[\frac{A_s}{e(\mathbf{X}_s, s)} f(s)\right] = E\big[f(s)\big], \quad \text{for all } f \in \text{span}\big\{b_k(\cdot)\big\}. \tag{9}$$

Thus, if $U_s \in \text{span}\{b_k(\cdot) : k = 1, \ldots, K\}$, then empirically balancing the spatial basis functions implies empirical balance of the latent confounder itself. More generally, when U_s is only approximately represented by (7), balancing the basis functions yields approximate balance in U_s, with any remaining imbalance governed by the approximation error. In this sense, the choice of basis and its dimension K controls a bias–variance tradeoff: richer bases reduce residual confounding at the expense of increased variability in the resulting weights. This perspective parallels sieve and kernel-based approaches to covariate balancing in non-spatial settings (Zhao 2019, Hazlett 2020, Wong & Chan 2017), but here the balanced function class is explicitly chosen to reflect spatial structure.

The spatial CBPS estimator can now be defined using a penalized covariate balancing loss function (Zhao 2019) that places explicit balance constraints on both the observed covariates, $\mathbf{X}_s$, and the spatial basis functions, $\mathbf{b}_s = (b_1(s), \ldots, b_K(s))^\intercal$. Define the propensity score model, $e_s(\boldsymbol{\theta}) \equiv e(\tilde{\mathbf{x}}_s; \boldsymbol{\theta})$, as

$$l\big(e_s(\boldsymbol{\theta})\big) = \mathbf{x}_s^\intercal \boldsymbol{\beta} + \mathbf{b}_s^\intercal \boldsymbol{\alpha}, \tag{10}$$

where $\tilde{\mathbf{x}}_s = (\mathbf{x}_s^\intercal, \mathbf{b}_s^\intercal)^\intercal$ collects the observed covariates (or possibly functions of the covariates)

and spatial basis functions evaluated at location s. The parameter vector $\boldsymbol{\theta} = (\boldsymbol{\beta}^\intercal, \boldsymbol{\alpha}^\intercal)^\intercal$ corresponds to the coefficients associated with these components, and $l(\cdot)$ denotes the logit link function. Rather than estimating $\boldsymbol{\theta}$ using the typical Bernoulli likelihood (5), we instead utilize a penalized M-estimator,

$$\hat{\boldsymbol{\theta}}_\lambda = \arg\min_{\boldsymbol{\theta}\in\boldsymbol{\Theta}} \left\{ \frac{1}{n} \sum_{s\in D_n} \rho\big(Z_s; \boldsymbol{\theta}\big) + \lambda J\big(\boldsymbol{\theta}\big) \right\}, \tag{11}$$

where

$$\rho\big(Z_s; \boldsymbol{\theta}\big) = -\left(\log \frac{e_s(\boldsymbol{\theta})}{1 - e_s(\boldsymbol{\theta})} - \frac{1}{e_s(\boldsymbol{\theta})} \right) A_s - \left(\log \frac{1 - e_s(\boldsymbol{\theta})}{e_s(\boldsymbol{\theta})} - \frac{1}{1 - e_s(\boldsymbol{\theta})} \right) (1 - A_s) \tag{12}$$

is a covariate balancing tailored loss function first introduced by Zhao (2019) and $J(\boldsymbol{\theta})$ is a regularization function that penalizes model complexity, with tuning parameter λ. For the remainder of the article, we will consider quadratic penalty terms, $J(\boldsymbol{\theta}) = \boldsymbol{\theta}^\intercal \mathbf{S} \boldsymbol{\theta}/2$, where $\mathbf{S}$ is a positive semi-definite penalty matrix.

The penalized loss in (11) has several desirable properties. First, given the logit link function in (10) and the quadratic penalty, (11) is strictly convex and second-order differentiable in $\boldsymbol{\theta}$. Consequently, solving for $\hat{\boldsymbol{\theta}}_\lambda$ is computationally efficient: it is straightforward to implement a Fisher scoring algorithm for computing $\hat{\boldsymbol{\theta}}_\lambda$ for any fixed λ (see Section 1 of the Supplementary Material for details). Perhaps more importantly, the estimated inverse propensity weights can be shown to empirically balance the J individual elements of $\tilde{\mathbf{X}}_s = (\mathbf{X}_s^\intercal, \mathbf{b}_s^\intercal)^\intercal$. To see this, note that

$$\frac{\partial}{\partial \theta_j} \left[\frac{1}{n} \sum_{s\in D_n} \rho\big(Z_s, \boldsymbol{\theta}\big) \right] = \frac{1}{n} \sum_{s\in D_n} \left(\frac{A_s}{e_s(\boldsymbol{\theta})} - \frac{1 - A_s}{1 - e_s(\boldsymbol{\theta})} \right) \tilde{X}_{sj}, \quad j = 1, \ldots, J. \tag{13}$$

As a result, when $\lambda = 0$, the spatial CBPS estimate $e_s(\hat{\boldsymbol{\theta}})$ provides inverse propensity weights such that the gradient of the loss function is exactly 0 for all features in $\tilde{\mathbf{X}}$. In other words, the weighted empirical distribution of observed features, $\tilde{X}_{sj}$, is identical across exposure levels:

$$\frac{1}{n} \sum_{s\in D_n} \frac{A_s}{e_s(\hat{\boldsymbol{\theta}})} \tilde{X}_{sj} = \frac{1}{n} \sum_{s\in D_n} \frac{1 - A_s}{1 - e_s(\hat{\boldsymbol{\theta}})} \tilde{X}_{sj}, \quad \text{for all } j = 1, \ldots, J. \tag{14}$$

Since $\{b_k(\cdot) : k = 1, \ldots, K\}$ are included in $\tilde{\mathbf{X}}$, (14) implies that the weighted average of the unmeasured spatial confounder approximated by (7), or indeed *any* linear combination of $b_k(\cdot)$, is the same across exposure levels. In stark contrast to the (penalized) likelihood-based approach, estimating $\hat{e}_s$ with (11) exchanges the priority of estimating the spatial surface to that of ensuring that features of the spatial surface are balanced between treatment groups, which more directly corresponds to the priority of confounding adjustment.

3.2 Bias Decomposition

To further understand the advantage of the spatial CBPS over the likelihood-based alternative, consider the following bias decomposition of the IPW estimator. Under Assumption 3, suppose that $\mu_s(a, \boldsymbol{\ell}) = E\big[Y_s \,|A_s = a, \mathbf{L}_s = \boldsymbol{\ell}\big]$ has an additive structure, such that $\mu_s(0, \boldsymbol{\ell}) = g(\mathbf{x}_s) + f(s)$ for

some square integrable functions $g(\cdot)$ and $f(\cdot)$. Furthermore, assume that the treatment effect is homogeneous, i.e., $\mu_s(1,\ell) = \tau + \mu_s(0,\ell)$, and that $Y_s = \mu_s(A_s, \mathbf{L}_s) + \epsilon_s$, where $E[\epsilon_s | A_s, \mathbf{L}_s] = 0$. Finally, suppose that $\mu_s(0,\ell)$ admits a series expansion, $\mu_s(0,\ell) = \sum_{j=1}^{\infty} \beta_j \phi_j(\mathbf{x}_s) + \sum_{k=1}^{\infty} \alpha_k b_k(s)$.

Under this set-up, it is straightforward to show that the IPW estimator (3) has the following decomposition (Zhao 2019),

$$\hat{\tau}_{IPW} - \tau = \underbrace{\left[\sum_{A_s=1} w_s \mu(0, \mathbf{L}_s) - \sum_{A_s=0} w_s \mu(0, \mathbf{L}_s) \right]}_{\text{imbalance}} + \underbrace{\left[\sum_{A_s=1} w_s \epsilon_s - \sum_{A_s=0} w_s \epsilon_s \right]}_{\text{noise}}, \tag{15}$$

where $w_s \equiv w(A_s, \mathbf{X}_s)$ are the normalized (i.e., Hájek) IPW weights from (3). Note that the second term in (15) has mean 0, so the bias of $\hat{\tau}_{IPW}$ is completely determined by the imbalance of $\mu(0, \mathbf{L}_s)$. Due to the series expansion of $\mu(0, \mathbf{L}_s)$, the imbalance term can be equivalently expressed as

$$\text{imbalance} = \sum_{j=1}^{\infty} \beta_j \left[\sum_{A_s=1} w_s \phi_j(\mathbf{X}_s) - \sum_{A_s=0} w_s \phi_j(\mathbf{X}_s) \right] + \sum_{k=1}^{\infty} \alpha_k \left[\sum_{A_s=1} w_s b_k(s) - \sum_{A_s=0} w_s b_k(s) \right]. \tag{16}$$

In other words, the bias of the IPW estimator is governed by how well the inverse propensity weights balance the functions $\{\phi_j(\cdot)\}_{j\geq 1}$ and $\{b_k(\cdot)\}_{k\geq 1}$ that span the baseline outcome regression $\mu(0, \mathbf{L}_s)$. This observation clarifies the advantage of the spatial CBPS over likelihood-based propensity score models: rather than prioritizing accurate prediction of treatment assignment or precise modeling of the spatial surface — objectives that are only indirectly related to the performance of the IPW estimator — spatial CBPS directly targets the balance conditions that determine bias.

In practice, the series expansion of $\mu(0, \mathbf{L}_s)$ must be truncated, and any remaining bias is driven by imbalance in the truncation remainder. Consequently, the effectiveness of the spatial CBPS approach depends on how well the chosen feature set $\tilde{\mathbf{X}}_s$ captures the dominant components of variation in $\mu(0, \mathbf{L}_s)$. From this perspective, the specification of $\tilde{\mathbf{X}}_s$ is naturally guided by features that are predictive of the outcome, rather than by parsimony or goodness-of-fit considerations for the propensity score model itself.

3.3 Inference Under Spatial Dependence

Because spatially indexed outcomes and exposures are typically correlated over space, standard variance estimators that assume independent observations can substantially understate uncertainty for IPW and AIPW estimators. Assuming i.i.d. observations and suitable regularity conditions on the propensity score model, Hirano et al. (2003) show that the IPW estimator with a likelihood-based sieve propensity score model is consistent and asymptotically normal, with asymptotic linear expansion

$$\hat{\tau}_{IPW} = \tau_0 + \frac{1}{n} \sum_{s \in D_n} \left[\psi(Z_s; \tau_0, e_0) + \alpha(Z_s) \right] + o_p(n^{-1/2}), \tag{17}$$

where τ_0 and e_0 denote the true ATE and propensity score, $\psi(\cdot)$ is the IPW estimating equation,

$$\psi(Z;\tau,e) = \frac{AY}{e(X)} - \frac{(1-A)Y}{1-e(X)} - \tau, \tag{18}$$

and $\alpha(\cdot)$ is a mean-zero correction term reflecting first-order uncertainty due to replacing the true propensity score e_0 with its estimate $\hat{e}$. The expansion of $\hat{\tau}_{IPW}$ in (17) implies that its asymptotic variance, σ^2_{IPW}, is given by the variance of the corresponding mean-zero influence function,

$$\sigma^2_{IPW} = E\big[\big(\psi(Z;\tau_0,e_0) + \alpha(Z)\big)^2\big]. \tag{19}$$

Hirano et al. (2003) derive an explicit form for (19) and show that $\hat{\tau}_{IPW}$ is semiparametrically efficient, provided the regularity conditions for their sieve propensity score estimator are met.

When the propensity score is estimated with the CBPS loss, these asymptotic results still hold (Zhao 2019), albeit with a slightly modified correction term in (17). Using the spatial CBPS with the linear-logistic propensity score model from (10) and assuming appropriate rate conditions on the basis dimension K and tuning parameter λ (outlined in Section 2 of the Supplementary Material), the correction term is given as

$$\alpha(Z) = -\mathbf{M}^\intercal\, \mathbf{A}^{-1}\, s_{\boldsymbol{\theta}}(Z;\boldsymbol{\theta}_0), \tag{20}$$

where $s_{\boldsymbol{\theta}}(Z;\boldsymbol{\theta})$ is the score of the spatial CBPS objective function in (11),

$$s_{\boldsymbol{\theta}}(Z;\boldsymbol{\theta}) = \partial_{\boldsymbol{\theta}}\, \rho(Z;\boldsymbol{\theta}), \tag{21}$$

$\mathbf{A}$ is the corresponding Hessian,

$$\mathbf{A} = E\big[\partial_{\boldsymbol{\theta}}\, s_{\boldsymbol{\theta}}(Z;\boldsymbol{\theta}_0)\big], \tag{22}$$

and $\mathbf{M}$ is the derivative of the IPW estimating equation with respect to nuisance parameters $\boldsymbol{\theta}$,

$$\mathbf{M} = E\big[\partial_{\boldsymbol{\theta}}\, \psi(Z;\tau_0,e_{\boldsymbol{\theta}})\big]_{\boldsymbol{\theta}=\boldsymbol{\theta}_0}. \tag{23}$$

Note that $\partial_{\boldsymbol{\theta}} := \frac{\partial}{\partial \boldsymbol{\theta}^\intercal}$ and $\boldsymbol{\theta}_0$ denotes the true parameter. The exact expressions for $s_{\boldsymbol{\theta}}(Z;\boldsymbol{\theta})$, $\mathbf{A}$, and $\mathbf{M}$ for the spatial CBPS estimator are given in the Section 2 of the Supplementary Material.

The asymptotic arguments of Hirano et al. (2003) suggest a simple approach to variance estimation in the i.i.d. setting: because the influence function contributions are uncorrelated across units, the asymptotic variance reduces to the variance of a single observation, which may be consistently estimated using a plug-in estimator of (19) based on empirical analogues of (22) and (23) (Lunceford & Davidian 2004). In spatial settings, however, the independence assumption underlying (19) is typically violated. Even after conditioning on observed covariates and spatial basis functions, nearby observations may remain correlated due to residual spatial structure in outcomes, exposures, or covariates. As a result, the asymptotic variance of $\hat{\tau}_{IPW}$ is no longer characterized by the variance of a single influence function contribution, but instead by a long-run variance that aggregates spatial covariances across locations. Formally, if $\hat{\tau}_{IPW}$ admits an asymptotically linear representation of the form (17) under weak spatial dependence

(Cohn 2024), then its asymptotic variance can be written as

$$\tilde{\sigma}^2_{IPW} = \lim_{n\to\infty} \frac{1}{n} \text{Var}\left(\sum_{s\in D_n} \left[\psi(Z_s; \tau_0, e_0) + \alpha(Z_s) \right] \right), \tag{24}$$

which reduces to (19) only in the absence of spatial dependence.

To circumvent this issue, note that under an increasing domain asymptotic regime (Jenish & Prucha 2009) and with an appropriate characterization of weak dependence (Cohn 2024), regions of space that are sufficiently far apart from one another may be considered approximately independent. This motivates a spatial analogue to the plug-in variance estimator from the independent data setting. In particular, we utilize a spatial heteroskedasticity and autocovariance (HAC) variance estimator (Conley 1999, Kelejian & Prucha 2007),

$$\hat{\sigma}^2_{\text{HAC}} = \frac{1}{n} \sum_{s\in D_n} \sum_{s'\in D_n} K\left(\frac{\|s - s'\|_2}{h} \right) \hat{\phi}_s \hat{\phi}_{s'}, \tag{25}$$

where $\hat{\phi}_s = \psi(Z_s; \hat{\tau}, \hat{e}) + \hat{\alpha}(Z_s)$ is the empirical influence function evaluated at location s and $K(\cdot)$ denotes a bounded, symmetric kernel function with compact support, such as the Bartlett or uniform kernel, with bandwidth parameter h. The spatial HAC estimator can be viewed as a spatially-smoothed plug-in variance estimator, with covariance contributions favoring observations that are close in spatial proximity. The bandwidth parameter, h, should be large enough that observations located outside the bandwidth window are considered effectively independent. In practice, we suggest selecting h using an empirical semivariogram estimated from the influence function residuals, $\hat{\xi}_s = \hat{\phi}_s - \bar{\hat{\phi}}$, as a data-driven heuristic to determine an optimal range for which residual spatial correlation is no longer significantly different than zero (Lambert et al. 2008). Alternatively, a spatial block bootstrap variance estimator may be used, in which bootstrap resamples are constructed by stitching together spatial blocks drawn from the observed sampling region (Lahiri 2003). This approach preserves local spatial dependence while approximating independence across sufficiently separated blocks. In practice, selecting the block size can be difficult; we use the nonparametric plug-in (NPI) method of Nordman et al. (2007).

Although the block bootstrap estimator may offer some benefits in finite samples, in practice it is much more computationally intensive than the HAC estimator. Empirically, we find the two estimators to behave similarly, each achieving nominal coverage in simulated data settings with moderate sample sizes (Section 4). Consequently, we recommend the HAC variance estimator be used in most settings, with the block bootstrap serving as a possible robustness check.

The arguments above extend directly to AIPW and other doubly robust estimators that admit an asymptotically linear representation. Under weak spatial dependence and suitable rate conditions on the nuisance functions, the AIPW estimator is asymptotically normal with a long-run variance determined by the spatial covariance structure of its influence function (Cohn 2024). Valid inference may therefore be obtained by replacing the IPW influence function in the spatial HAC estimator with the corresponding AIPW influence function. Additional details on the asymptotic regime, regularity conditions, and implementation of the variance estimators

are provided in the Supplementary Material.

3.4 Practical Implementation: Spatial Basis Specification and Regularization

We provide a user-friendly R package, `spbalance`, that implements the spatial CBPS and inference procedures outlined in Sections 3.1 and 3.3. The package includes three implementations of the penalized spatial CBPS estimator in (11):

1. **Thin-plate regression splines (TPRS)**. The latent spatial confounder, $f(s)$, is represented using thin-plate regression splines (Wood 2003), i.e., low-rank, isotropic basis functions that avoids knot placement and controls smoothness via a quadratic penalty that regularizes rapidly varying spatial components while leaving low-frequency, smoothly varying components relatively unpenalized.

2. **Multiresolution Wendland basis (MRW)**. The latent spatial confounder, $f(s)$, is modeled as an additive sum of spatial functions, i.e., $f(s) = \sum_j f_j(s)$, where each $f_j(s)$ is represented in turn as a linear combination of compact, two-dimensional Wendland basis functions centered at knots organized on regular grids of increasing resolution; consequently, each component function $f_j(\cdot)$ corresponds to a different scale of spatial confounding. At each resolution, coefficients are regularized using a sparse quadratic penalty analogous to a Markov random field prior on the knot lattice, enforcing local smoothness while progressively shrinking fine-scale spatial variation (Nychka et al. 2015).

3. **Random Fourier features (RFF)**. The latent spatial confounder, $f(s)$, is modeled using random Fourier features (Rahimi & Recht 2007), yielding a finite-dimensional approximation to a reproducing kernel Hilbert space (RKHS) associated with a shift-invariant kernel (e.g., exponential or squared exponential). This approach captures spatial confounding through global, nonlocal basis functions operating in the frequency domain. Regularization is imposed via a ridge penalty on the Fourier coefficients, which controls smoothness and stabilizes the resulting CBPS weights.

In addition, the spatial CBPS implementation allows for the specification of an arbitrary basis function and penalty matrix, so that practitioners can tailor the basis function representation to specific applications of interest. Furthermore, the HAC and block bootstrap variance estimates from Section 3.3 can be found for either IPW or AIPW estimates of the ATE.

The specification of the tuning parameter λ is an important determinant in the success of the spatial CBPS method. In theory, the asymptotic results for the IPW estimator requires the propensity score model to be undersmoothed, reflecting similar findings from the semiparametric regression literature (Rice 1986). In practice, data driven tuning parameter selection is preferred. Optimal tuning parameter selection for balancing weights remains an open problem (Zhao 2019), and suggested approaches include cross-validation as well as user-specified tolerances which restrict the allowable total coefficient of variation in the estimated weights (Ben-Michael et al. 2021). The Supplementary Material (Section 3) outlines several approaches for choosing λ that we have implemented in `spbalance`. We find that Zhao (2019)'s suggestion of choosing the largest λ such that the coefficient of variation of its associated weights is greater

than 0.9 times the largest coefficient of variation in the regularization path, has performed well empirically.

4 Simulation Study

We conduct a simulation study to compare the performance of spatial CBPS IPW and AIPW estimators against several alternatives in settings with unmeasured spatial confounding. We evaluate methods using the mean squared error (MSE) of the ATE estimates and, when appropriate, compare covariate balance of observed and latent confounders for the different weighting estimators.

4.1 Study Design

We replicate the data generating procedure outlined in Tec et al. (2023), in which (i) two spatially-varying covariates, $\boldsymbol{X}_s$, are generated from a Gaussian process (GP) with a squared exponential covariance function on a dense 100×100 grid, (ii) an unobserved spatial function, $f(s)$, is defined (via a convolution process) as a nonlocal function of the observed covariates $\boldsymbol{X}_s$; (iii) the binary exposure, A_s, is generated with the probability of treatment defined via a linear-logit model with additive terms involving $\boldsymbol{X}_s$ and $f(s)$, and (iv) the outcome, Y_s, is defined as the linear combination of $f(s)$, $\boldsymbol{X}_s$, A_s, and a spatially dependent error term. Notably, it is assumed that $f(s)$ is unobserved; thus, $U_s = f(s)$ is an unmeasured confounder.

We ultimately consider four data generating procedures, with only minor changes to the general procedure outlined above. First, the convolution generating $f(s)$ has two forms, one in which $f(s)$ is a linear transformation of $\boldsymbol{X}_s$ (the *linear* set-up) and another where $\boldsymbol{X}_s$ is passed through a nonlinear transformation before convolution (the *nonlinear* set-up). Second, we consider the case where data on the entire 100×100 grid is observed (the *dense* set-up) as well as the case where only 1000 randomly selected observations are retained from the grid (the *sparse* set-up). Consequently, the four data generating procedures include settings with linear-dense, linear-sparse, nonlinear-dense, and nonlinear-sparse observations. The complete details of each data generating procedure (along with figures showing representative simulated data) are outlined in Section 4 of the Supplementary Material.

For each set of simulated data, we compare the performance of six ATE estimators: three IPW estimators and three outcome regression/AIPW estimators. The IPW estimators differ in how the propensity score is estimated, including (i) a logistic regression model that only includes the observed covariates, $\boldsymbol{X}_s$, in the linear predictor — we denote this as `glm(no sp))`; (ii) a generalized additive model in which the propensity score is modeled as in (5), with $f(s)$ modeled using thin plate regression splines (TPRS) — we refer to this as `gam`; and (iii) the spatial CBPS estimator with the latent confounder represented with TPRS — we refer to this as `bal`. The outcome and AIPW estimators include (iv) an outcome regression model with fixed effects consisting of A_s and $\mathbf{X}_s$ and a TPRS spatial effect term (`outcome`); (v) an AIPW estimator using the previously described GAM propensity score model from (ii) and the outcome model from (iv) — we refer to this as `gam(DR)`; and finally (vi) an AIPW estimator with spatial CPBS

from (iii) and outcome regression from (iv), which we call `bal(DR)`. Normalized (i.e., Hajek) weights were used for all weighting estimators.

The GAM models were fitted using the `mgcv` package from R (Wood 2000) with spline basis dimension $K = 500$, while the spatial CBPS models were fitted as described in Section 3 with the same basis dimension ($K = 500$). Additional estimators were considered, including propensity score and outcome regression models that included spatial random effects, however the results were not sufficiently different from the spline estimators to warrant inclusion.

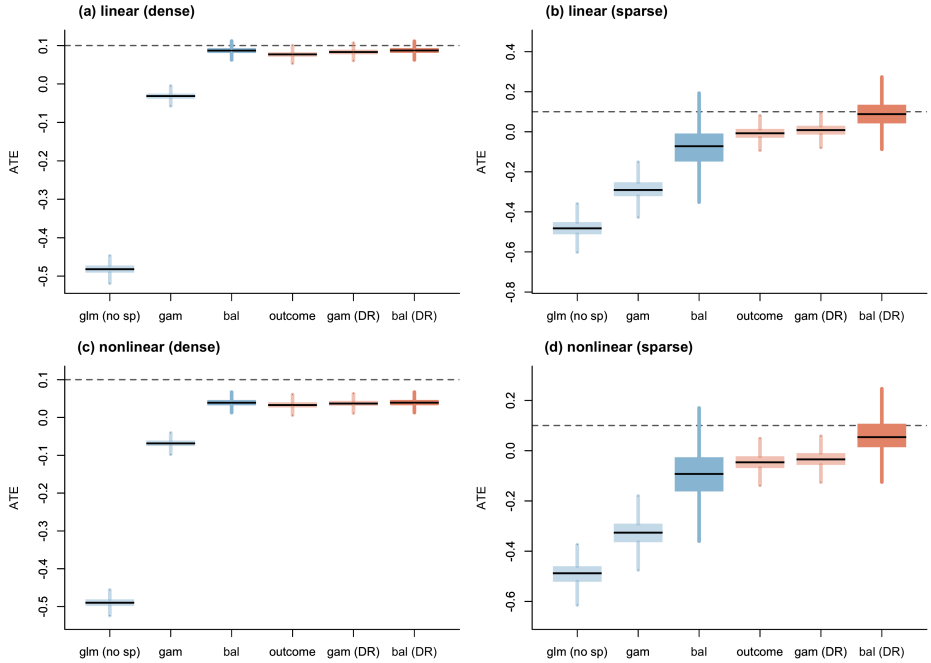

Figure 2: The performance of IPW (shown in blue) and outcome regression/AIPW estimators (shown in orange) of the ATE (dashed-line representing the true value of 0.1) in four settings: (a) observations generated on a 100×100 (dense) grid using the linear simulation set-up; (b) observations generated via 1000 randomly selected points (sparse) using the linear set-up; (c) dense grid with nonlinear set-up; and (d) sparse points with nonlinear set-up. Estimators that utilize spatial CBPS are indicated with a darker shading. The box plots summarize ATE estimates from 1000 simulations. See the main text for a description of the six estimators.

Table 1: Mean squared error for the six estimators and four data generating regimes; the lowest MSE for each set-up is shown in bold.

	Estimator					
Set-Up	`glm(no sp)`	`gam`	`bal`	`outcome`	`gam(DR)`	`bal(DR)`
linear, dense	0.3382	0.0174	0.0003	0.0006	0.0004	**0.0002**
linear, sparse	0.3400	0.1517	0.0416	0.0127	0.0102	**0.0053**
nonlinear, dense	0.3482	0.0285	**0.0038**	0.0046	0.0040	**0.0038**
nonlinear, sparse	0.3504	0.1876	0.0471	0.0226	0.0196	**0.0065**

4.2 Empirical Results

The performance of the six estimators on 1000 simulated data sets is summarized in Figure 2, with corresponding MSE shown in Table 1. The IPW estimator with naive (i.e., non-spatial) propensity score model is universally the worst of the six, exhibiting strong bias in every data generating and sampling scenario. Similarly, the IPW estimator that utilizes the likelihood-based spatial propensity score (`gam`) has considerably worse MSE compared to the other spatial confounding adjustment procedures. In contrast, the spatial CBPS IPW estimator (`bal`) is competitive with the outcome regression and AIPW estimators in the dense sampling regimes, with only slightly worse MSE than the best estimator (e.g., 0.0003 vs. 0.0002 in the linear and dense set-up). The spatial CBPS estimator is not as competitive relative to the outcome regression and AIPW estimators in the sparse sampling regimes, although it continues to outperform the likelihood-based IPW estimator. Similarly, the spatial CBPS AIPW estimator exhibits the lowest MSE in every data generating scenario, with considerable advantages in the sparse sampling regimes (Figures 2(b) and (d)), where there is not enough data for the outcome regression model to adequately capture the variability due to the latent spatial confounder.

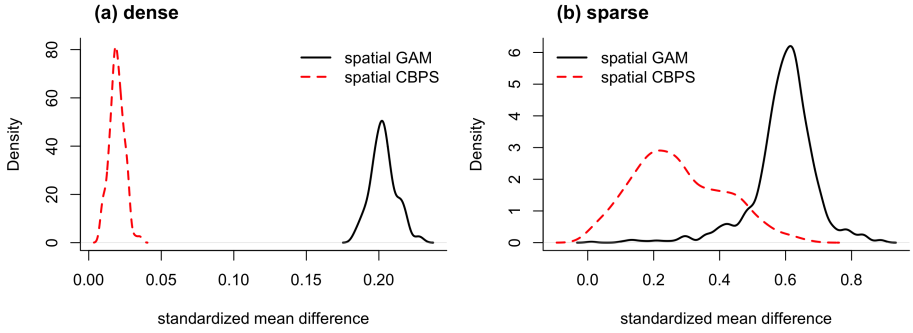

Figure 3: Distributions of weighted standardized mean differences (SMD) of the latent spatial confounder, U_s, when using IPW weights with a spatial CBPS vs. spatial GAM propensity score for the (a) dense and (b) sparse sampling regimes. Distributions are SMD across 100 simulation replicates.

To further understand the performance benefits of the spatial CBPS estimators compared to their likelihood-based counterparts, we compare the weighted standardized mean difference (SMD, see Austin 2009, for details) of the true unmeasured spatial confounder, U_s, across 100 simulations. In practice, we would never be able to assess covariate balance diagnostics with respect to the unobserved U_s, but in our simulation setting this comparison is illustrative of the advantages offered by the spatial CBPS estimator. The sample distribution of the weighted SMD is shown in Figure 3 for the linear data generating process with (a) dense and (b) sparse sampling regimes. In both cases, the spatial CBPS weights result in lower weighted SMDs on average; the advantage is especially pronounced in the dense data setting. As a rule of thumb, practitioners are often satisfied if weighted SMDs are below 0.1 (Austin 2009). This is satisfied by the spatial CBPS weights every time in the dense data setting, while the spatial GAM weights never reach this threshold. Both methods exhibit larger variance in the sparse data setting, although the spatial CBPS weights result in SMDs below a 0.2 cutoff 33 percent of the time, compared with 8 percent for the spatial GAM method.

5 Estimating the Effect of CAFO Proximity on Well Water Quality in Iowa

Beginning in the 1990s, Iowa has experienced a rapid increase in large-scale intensive livestock production, commonly referred to as animal feeding operations or concentrated animal feeding operations (AFOs/CAFOs). The increase in CAFOs has been primarily driven by hog production — in 2002, the state produced 14 million hogs (Herriges et al. 2005); by 2023, the total was 25 million (USDA 2024). In recent years a growing literature has highlighted the negative impacts of CAFOs on the environment and human health, with animal waste and emissions originating from CAFOs being linked with poor air, water, and soil quality and increased risk of respiratory disease, infection, and waterborne disease (Son et al. 2024).

The impact of CAFOs on human health can be especially pronounced when animal waste contaminates water sources. Livestock manure has been identified as a primary driver of increased nitrate levels in Iowa surface waters (Jones et al. 2019). Perhaps more concerning is the possibility that animal waste can also affect groundwater quality. Injection of liquid swine manure into soils has become more common with the increase in CAFOs, and proximity to CAFOs has been shown to be an important predictor of nitrate concentrations in private wells in Iowa (Wheeler et al. 2015). These results suggest that CAFO proximity may be a public health risk due to possible impacts on drinking water quality, as elevated nitrate concentrations in drinking water are associated with several adverse health outcomes, including infant methemoglobinemia, colorectal cancer, thyroid disease, and neural tube defects (Ward et al. 2018). Consequently, estimating the average causal effect of CAFO proximity on total nitrate levels in private well water in Iowa remains a pressing challenge, as the complex spatial confounding processes dictating where CAFOs operate and groundwater quality can be hard to fully quantify.

We collected Iowa private well water data from the Iowa Private Well Forecasting System (Iowa DNR 2025*b*). The data consist of measurements of max total nitrate (mg/L) taken from 14893 private wells, tested between 2020–2022 in Iowa. Figure 1(b) shows the spatial distribution of the tested wells as well as their corresponding outcomes. Note that increased risk of adverse health outcomes has been associated with nitrate levels of 5.0 mg/L or higher (Ward et al. 2018).

CAFO data were obtained from the Iowa DNR (Iowa DNR 2025*a*) They include all active feeding operations with 300 animal units (AUs) or greater, which represents a lower bound for regulatory action on CAFOs. The data are from 2023 (earlier years were unavailable), however, farm locations are quite stable from year to year, suggesting that the data are likely a reasonable proxy for CAFO exposure from 2020–2022. We define treatment A_s to be 1 if there exists a CAFO operating within 3 km of a private well, and 0 otherwise; approximately 28% of the tested wells are classified as exposed. Figure 1(a) displays the CAFO exposures, which are most concentrated in the central and northern portions of the state.

In addition to the exposure and outcome data, potential confounders were compiled from various sources. These include elevation, whether the well is located in regions with carbonate bedrock or within 1000 feet of a sinkhole, climate variables such as 30-year annual average

precipitation, max and min temperature, and maximum vapor pressure deficit, demographic variables such as census block-level population density, proportion white, and proportion with high school education, well depth, and land cover use summaries, such as proportion developed land and proportion cropland within 100 m and 5 km of the well. A complete list of all confounders and data sources can be found in the Supplementary Material.

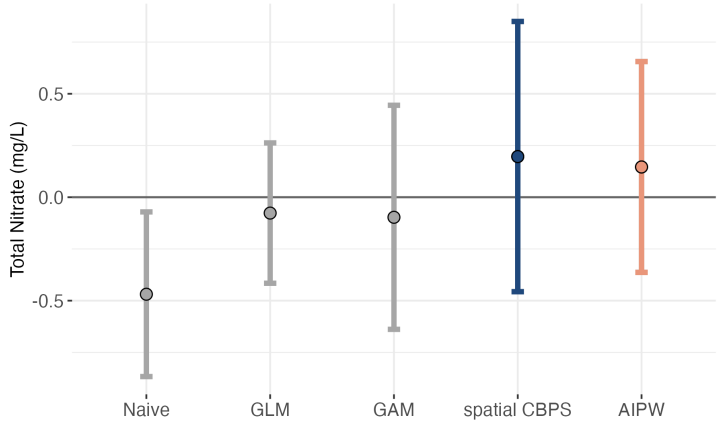

Figure 4: Effect estimates and 95% CIs for CAFO proximity (within 3 km) on max total nitrate concentration (mg/L) of Iowa private well water during 2020–2022, using the naive difference-in-means estimator, the non-spatial GLM, spatial GAM, and spatial CBPS IPW (shown in blue) estimators, and the spatial CBPS AIPW estimator (shown in orange).

5.1 Results

We estimated the effect of CAFO proximity on private well water total nitrate level using five methods. First, a naive difference in mean outcome between exposed and unexposed wells is included as a baseline comparison (Figure 4). We see that the naively estimated average effect of CAFO exposure on max nitrate is -0.47, with 95% confidence interval $(-0.87, -0.07)$. In other words, without any confounding adjustment the data suggest that CAFO proximity actually *reduces* total nitrate levels by -0.47 mg/L. This result is counterintuitive, as there is no reason to believe that CAFOs would result in cleaner water, serving as a strong indication that confounding adjustment is needed.

Consequently, we also consider three IPW estimators: one where the propensity score is estimated using a logistic regression with only the observed data as predictors (hereafter, 'GLM'), another that uses the spatial GAM model given in (5), with spatial effect modeled with TPRS (denoted as 'GAM') and with a ridge penalty on the first-order fixed effect terms, and finally a spatial CBPS IPW estimator with TPRS spatial basis functions to model any latent spatial confounding and a ridge penalty for the observed covariates. Finally, we include a spatial CBPS AIPW estimator, where the regression model contains all first-order covariates as fixed effects plus an additive spatial effect term, which is modeled using TPRS.

The results from the analysis are shown in Figure 4. Compared to the naive difference-in-means estimator, all confounding adjustment methods result in effect estimates that are closer to zero; in fact, none of the confounding adjustment estimates are significantly different than zero (using a 95% significance level). Qualitatively, we see very little difference between the

GLM and spatial GAM point estimates ($\hat{\tau}_{GLM} = -0.08$ vs $\hat{\tau}_{GAM} = -0.10$), suggesting that the addition of the spatial effect to the propensity score did little to adjust any unmeasured spatial confounding. In contrast, the spatial CBPS estimators (IPW, shown in blue; and AIPW, shown in orange) result in point estimates greater than zero ($\hat{\tau}_{IPW} = 0.20$ and $\hat{\tau}_{AIPW} = 0.15$, respectively). Consequently, although we do not have statistically significant evidence to suggest that the presence of a CAFO within 3 km of a well results in an increase in total nitrate, we have some indication that adjusting for unmeasured spatial confounding is an important consideration when estimating causal effects in this setting. Sensitivity checks on these results, including covariate balance and overlap diagnostics, are included in the Supplementary Material.

6 Discussion

Unobserved confounding is a persistent threat when analyzing observational environmental data, however, in some settings it may be reasonable to use spatial proximity to adjust for possible latent spatial confounding. While the vast majority of such methods focus on regression adjustment strategies, we argue that design-based perspectives for confounding adjustment remain underutilized in the spatial confounding setting. We demonstrate that modern balancing weight estimators may be particularly useful in such settings, and we propose spatial CBPS to construct weights that directly target balance of a spatial basis function representation of the latent confounder. Low rank basis function representations can be used to facilitate fast computation for point-referenced and areal spatial data, and we use a computationally efficient spatial HAC variance estimator for inference. An accompanying R package (`spbalance`) implements spatial CBPS for use with IPW and AIPW estimators, and includes options for HAC and block bootstrap variance estimation. The package includes three common basis function representations of spatial processes (thin plate regression splines, multiresolution Wendland bases, and random Fourier features) and is designed to flexibly accommodate user-specified inputs.

From a practical standpoint, the spatial CBPS approach is most appropriate in settings where unmeasured confounding is plausibly dominated by smooth, low-frequency spatial processes. In such cases, the choice of spatial basis and degree of regularization plays a role analogous to bandwidth selection in nonparametric regression: too little smoothing may result in residual confounding, while overly rich representations can induce unstable weights. Importantly, these trade-offs can be assessed empirically through balance diagnostics on the estimated spatial features themselves and through sensitivity of estimates across basis specifications. In contrast, spatial CBPS is likely to perform poorly in settings where confounding varies at very fine spatial scales, as specifying appropriate basis function representations of such processes is challenging (Nychka et al. 2015).

The analysis of Iowa groundwater quality illustrates several of these considerations in practice. In this setting, both CAFO location and groundwater nitrate concentrations are shaped by complex, spatially structured processes that are difficult to fully observe, making unmeasured spatial confounding a primary concern. The naive and likelihood-based spatial propensity score estimators produced effect estimates that were qualitatively similar despite

incorporating spatial structure, suggesting limited confounding adjustment at the relevant spatial scales. In contrast, the spatial CBPS estimators yielded substantively different point estimates, reflecting improved balance in spatial features that might plausibly approximate the latent confounding process. Although uncertainty remains substantial and the estimated effects are not significantly different from zero, the analysis highlights how explicitly targeting spatial balance can meaningfully alter causal conclusions in environmental health applications where standard approaches may inadequately address spatial confounding.

At the same time, this application also underscores that the benefits of spatial CBPS arise from specific modeling and computational choices, and that alternative formulations of balancing weights may be attractive in other settings. For example, the bias decomposition in Section 3.2 assumes a homogeneous treatment effect, which may be hard to justify in some settings. Balancing weights for ATE estimation with heterogeneous treatments can be obtained through a slight modification of (11), in which two sets of balancing weights are separately estimated — one set of weights for the treatment units, such that the weighted average of treatment unit covariates equals the empirical average of covariates from the entire sample, and another set of weights that does the same thing for control (Tan 2019). In theory, this approach can be implemented in our spatial CBPS set-up, although we expect a bias-variance tradeoff due to the increased variability in the weights. This remains an avenue of future work.

The spatial CBPS objective in (11) was chosen, in part, due to computational convenience: the regularized loss function is strictly convex and allows us to easily implement a Fisher scoring algorithm that is fast even in settings with moderate-to-large spatial data sets. However, the CBPS objective has received some criticism (Wong & Chan 2017, Ben-Michael et al. 2021), as it still requires construction of the weights through a (logistic) propensity score model that may prove arbitrarily restrictive in practice. An alternative perspective for balancing weights instead specifies an optimization problem that directly estimates weights $\mathbf{w} = (w_1, \ldots, w_n)^\intercal$, subject to constraints on weight variability (Zubizarreta 2015, Wong & Chan 2017). The utility of this method for spatial data remains unexplored, although implementation of kernel-based weight estimators in non-spatial settings (Wong & Chan 2017, Hazlett 2020) indicates that it may be a promising direction for future research.

Finally, we note that the proposed spatial CBPS method is at present limited to settings with binary exposures at a single time-point. Extensions of balancing weight estimators to settings with time-varying (Imai & Ratkovic 2015) and continuous treatments (Fong et al. 2018) have been considered in non-spatial settings, and the results of this paper suggest that adapting them to the spatial confounding setting may offer benefits over alternative approaches to confounding adjustment (Gilbert et al. 2024, Dupont et al. 2025). Furthermore, in many spatial settings we may expect treatment effects to "spillover" across space, violating the no interference component of SUTVA (Reich et al. 2021). Estimating causal effects in such settings is challenging (Wikle & Zigler 2024), and to date, existing approaches have ignored the possibility of spatial confounding. Adapting the spatial CBPS procedure to settings with spatial interference remains an inviting challenge.

References

Austin, P. C. (2009), 'Balance diagnostics for comparing the distribution of baseline covariates between treatment groups in propensity-score matched samples', *Statistics in Medicine* **28**(25), 3083–3107.

Banerjee, S., Carlin, B. P. & Gelfand, A. E. (2015), *Hierarchical Modeling and Analysis for Spatial Data*, Monographs on Statistics and Applied Probability, 2nd edn, Chapman and Hall/CRC, Boca Raton, FL.

Ben-Michael, E., Feller, A., Hirshberg, D. A. & Zubizarreta, J. R. (2021), 'The balancing act in causal inference'. *arXiv*:2110.14831 [stat.ME].

Bruns-Smith, D., Dukes, O., Feller, A. & Ogburn, E. L. (2025), 'Augmented balancing weights as linear regression', *Journal of the Royal Statistical Society Series B: Statistical Methodology* p. qkaf019.

Burns, A. M., Chandler, G., Dunham, K. J. & Carlton, A. G. (2023), 'Data gap: Air quality networks miss air pollution from concentrated animal feeding operations', *Environmental Science & Technology* **57**(49), 20718–20725.

Chattopadhyay, A., Hase, C. H. & Zubizarreta, J. R. (2020), 'Balancing vs modeling approaches to weighting in practice', *Statistics in Medicine* **39**(24), 3227–3254.

Cohn, E. R. (2024), 'Causal Inference for Spatially Dependent Data', PhD thesis, Harvard University. Chapter 3 in PhD dissertation: *Statistical Methods for the Study of Effect Modification and Spatial Causal Inference: Theory and Applications*.

Conley, T. (1999), 'Gmm estimation with cross sectional dependence', *Journal of Econometrics* **92**(1), 1–45.

Cressie, N. & Johannesson, G. (2008), 'Fixed Rank Kriging for Very Large Spatial Data Sets', *Journal of the Royal Statistical Society Series B: Statistical Methodology* **70**(1), 209–226.

Cressie, N., Sainsbury-Dale, M. & Zammit-Mangion, A. (2022), 'Basis-function models in spatial statistics', *Annual Review of Statistics and Its Application* **9**(Volume 9, 2022), 373–400.

Cressie, N. & Wikle, C. K. (2011), *Statistics for Spatio-Temporal Data*, John Wiley & Sons, Hoboken, NJ.

Datta, A. & Stein, M. L. (2025), 'Consistent infill estimability of the regression slope between gaussian random fields under spatial confounding'. *arXiv*:2506.09267 [math.ST].

Davis, M. L., Neelon, B., Nietert, P. J., Hunt, K. J., Burgette, L. F., Lawson, A. B. & Egede, L. E. (2019), 'Addressing geographic confounding through spatial propensity scores: a study of racial disparities in diabetes', *Statistical Methods in Medical Research* **28**(3), 734–748.

Deville, J.-C. & Särndal, C.-E. (1992), 'Calibration estimators in survey sampling', *Journal of the American Statistical Association* **87**(418), 376–382.

Dupont, E., Marques, I. & Kneib, T. (2025), 'Demystifying spatial confounding'. *arXiv*:2309.16861 [stat.ME].

Dupont, E., Wood, S. N. & Augustin, N. H. (2022), 'Spatial+: A novel approach to spatial confounding', *Biometrics* **78**(4), 1279–1290.

Epps, A., Dressel, I. M., Guo, X., Odanibe, M., Fields, K. P., Carlton, A. M. G., Sun, K. & Pusede, S. E. (2025), 'Satellite observations of atmospheric ammonia inequalities associated with industrialized swine facilities in eastern north carolina', *Environmental Science & Technology* **59**(5), 2651–2664.

Folsom, R. (1991), Exponential and logistic weight adjustments for sampling and nonresponse error reduction, *in* 'Proceedings of the Social Statistics Section', American Statistical Association, pp. 197–202.

Fong, C., Hazlett, C. & Imai, K. (2018), 'Covariate balancing propensity score for a continuous treatment: Application to the efficacy of political advertisements', *The Annals of Applied Statistics* **12**(1), 156 – 177.

Gilbert, B., Datta, A., Casey, J. A. & Ogburn, E. L. (2024), 'A causal inference framework for spatial confounding'. *arXiv*:2112.14946 [stat.ME].

Guan, Y., Page, G. L., Reich, B. J., Ventrucci, M. & Yang, S. (2022), 'Spectral adjustment for spatial confounding', *Biometrika* **110**(3), 699–719.

Hainmueller, J. (2012), 'Entropy balancing for causal effects: A multivariate reweighting method to produce balanced samples in observational studies', *Political Analysis* **20**(1), 25–46.

Hájek, J. (1971), Comment on "an essay on the logical foundations of survey sampling, part one", *in* V. P. Godambe & D. A. Sprott, eds, 'The Foundations of Survey Sampling', Holt, Rinehart and Winston, p. 236.

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**(4), 243–254.

Hazlett, C. (2020), 'Kernel balancing: A flexible non-parametric weighting procedure for estimating causal effects', *Statistica Sinica* **30**, 1155–1189.

Herriges, J. A., Secchi, S. & Babcock, B. A. (2005), 'Living with hogs in iowa: The impact of livestock facilities on rural residential property values', *Land Economics* **81**(4), 530–545.

Hirano, K., Imbens, G. W. & Ridder, G. (2003), 'Efficient estimation of average treatment effects using the estimated propensity score', *Econometrica* **71**(4), 1161–1189.

Hirshberg, D. A. & Wager, S. (2021), 'Augmented minimax linear estimation', *The Annals of Statistics* **49**(6), 3206 – 3227.

Hodges, J. S. & Reich, B. J. (2010), 'Adding spatially-correlated errors can mess up the fixed effect you love', *The American Statistician* **64**(4), 325–334.

Horvitz, D. G. & Thompson, D. J. (1952), 'A generalization of sampling without replacement from a finite universe', *Journal of the American Statistical Association* **47**(260), 663–685.

Imai, K. & Ratkovic, M. (2014), 'Covariate balancing propensity score', *Journal of the Royal Statistical Society Series B: Statistical Methodology* **76**(1), 243–263.

Imai, K. & Ratkovic, M. (2015), 'Robust estimation of inverse probability weights for marginal structural models', *Journal of the American Statistical Association* **110**(511), 1013–1023.

Imbens, G. W. & Rubin, D. B. (2015), *Causal Inference in Statistics, Social, and Biomedical Sciences*, Cambridge University Press, Cambridge, UK.

Iowa DNR (2025*a*), 'Animal feeding operations (afo) data', https://programs.iowadnr.gov/animalfeedingoperations/. Geospatial dataset and regulatory records. Data portal: https://geodata.iowa.gov/documents/abfbd972640d4e87b6c48dc669775767.

Iowa DNR (2025*b*), 'Private well tracking system (pwts) database', https://programs.iowadnr.gov/pwts/. Statewide database of private well testing records.

Jenish, N. & Prucha, I. R. (2009), 'Central limit theorems and uniform laws of large numbers for arrays of random fields', *Journal of Econometrics* **150**(1), 86–98.

Jones, C. S., Drake, C. W., Hruby, C. E., Schilling, K. E. & Wolter, C. F. (2019), 'Livestock manure driving stream nitrate', *Ambio* **48**(10), 1143–1153.
URL: *https://doi.org/10.1007/s13280-018-1137-5*

Kang, J. D. Y. & Schafer, J. L. (2007), 'Demystifying double robustness: A comparison of alternative strategies for estimating a population mean from incomplete data', *Statistical Science* **22**(4), 523–539.

Kelejian, H. H. & Prucha, I. R. (2007), 'Hac estimation in a spatial framework', *Journal of Econometrics* **140**(1), 131–154.

Khan, K. & Berrett, C. (2024), 'Re-thinking spatial confounding in spatial linear mixed models'. *arXiv*:2301.05743 [stat.ME].

Khan, K. & Calder, C. A. (2022), 'Restricted spatial regression methods: Implications for inference', *Journal of the American Statistical Association* **117**(537), 482–494.

Lahiri, S. N. (2003), *Resampling Methods for Dependent Data*, Springer Series in Statistics, Springer, New York.

Lambert, D. M., Florax, R. J. & Cho, S.-H. (2008), Bandwidth selection for spatial hac and other robust covariance estimators, Working Paper 08-10, Department of Agricultural Economics, Purdue University.
URL: *https://ageconsearch.umn.edu/record/45964/files/08-10.pdf*

Lamouroux, J., Geffroy, A., Leblond, S., Meyer, C. & Albert, I. (2025), 'Addressing spatial confounding in geostatistical regression models: An r-inla approach', *Methods in Ecology and Evolution* **16**(9), 2082–2097.

Lindgren, F., Rue, H. & Lindström, J. (2011), 'An explicit link between gaussian fields and gaussian markov random fields: the stochastic partial differential equation approach', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **73**(4), 423–498.

Lunceford, J. K. & Davidian, M. (2004), 'Stratification and weighting via the propensity score in estimation of causal treatment effects: a comparative study', *Statistics in Medicine* **23**(19).

Marques, I., Kneib, T. & Klein, N. (2023), 'Mitigating spatial confounding by explicitly correlating gaussian random fields', *Environmetrics* **34**(4), e2801.

Neyman, J. (1923), 'On the application of probability theory to agricultural experiments: Essay on principles', *Statistical Science* **5**(4), 465–480. Originally published in 1923 in Polish; republished in 1990.

Nordman, D. J., Lahiri, S. N. & Fridley, B. L. (2007), 'Optimal block size for variance estimation by a spatial block bootstrap method', *Sankhyā: The Indian Journal of Statistics, Series A* **69**(3), 468–493.

Nychka, D., , Soutir, B., , Dorit, H., , Finn, L., & Sain, S. (2015), 'A Multiresolution Gaussian Process Model for the Analysis of Large Spatial Datasets', *Journal of Computational and Graphical Statistics* **24**(2), 579–599.

Paciorek, C. J. (2010), 'The importance of scale for spatial-confounding bias and precision of spatial regression estimators', *Statistical Science* **25**(1), 107–125.

Papadogeorgou, G., Mealli, F. & Zigler, C. M. (2019), 'Causal inference with interfering units for cluster and population level treatment allocation programs', *Biometrics* **75**(3), 778–787.

Rahimi, A. & Recht, B. (2007), Random features for large-scale kernel machines, *in* 'Advances in Neural Information Processing Systems', Vol. 20, pp. 1177–1184.

Rainey, M. J. & Keller, K. P. (2025), 'Semiparametric approaches for mitigating spatial confounding in large environmental epidemiology cohort studies', *Environmetrics* **36**(6), e70028.

Reich, B. J., Yang, S., Guan, Y., Giffin, A. B., Miller, M. J. & Rappold, A. (2021), 'A review of spatial causal inference methods for environmental and epidemiological applications', *International Statistical Review* **89**(3), 605–634.

Rice, J. (1986), 'Convergence rates for partially splined models', *Statistics Probability Letters* **4**(4), 203–208.

Robins, J. (1986), 'A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect', *Mathematical Modelling* **7**(9), 1393–1512.

Robins, J. M., Rotnitzky, A. & Zhao, L. P. (1994), 'Estimation of regression coefficients when some regressors are not always observed', *Journal of the American Statistical Association* **89**(427), 846–866.

Rosenbaum, P. R. & Rubin, D. B. (1983), 'The central role of the propensity score in observational studies for causal effects', *Biometrika* **70**(1), 41–55.

Rubin, D. B. (1974), 'Estimating causal effects of treatments in randomized and nonrandomized studies', *Journal of Educational Psychology* **66**, 688–701.

Rubin, D. B. (1980), 'Randomization analysis of experimental data: The Fisher randomization test comment', *Journal of the American Statistical Association* **75**(371), 591–593.

Schnell, P. M. & Papadogeorgou, G. (2020), 'Mitigating unobserved spatial confounding when estimating the effect of supermarket access on cardiovascular disease deaths', *The Annals of Applied Statistics* **14**(4), 2069 – 2095.

Son, J.-Y., Heo, S., Byun, G., Foo, D., Song, Y., Lewis, B. M., Stewart, R., Choi, H. M. & Bell, M. L. (2024), 'A systematic review of animal feeding operations including concentrated animal feeding operations (cafos) for exposure, health outcomes, and environmental justice', *Environmental Research* **259**, 119550.

Tan, Z. (2019), 'Regularized calibrated estimation of propensity scores with model misspecification and high-dimensional data', *Biometrika* **107**(1), 137–158.

Tec, M., Scott, J. G. & Zigler, C. (2023), Weather2vec: Representation learning for causal inference with non-local confounding in air pollution and climate studies, *in* 'Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI-23)', AAAI Press, Washington, DC, USA, pp. 14504–14513.

Thaden, H. & Kneib, T. (2018), 'Structural equation models for dealing with spatial confounding', *The American Statistician* **72**(3), 239–252.

USDA (2024), Iowa ag news – hogs & pigs (ia-hogs-12-24), Technical Report IA-Hogs-12-24, National Agricultural Statistics Service, Des Moines, IA.
URL: *https://data.nass.usda.gov/Statistics*$_b y_S tate/Iowa/Publications/Livestock_R eport$/2024/$IA-Hogs-12-24.pdf$

Wahba, G. (1990), *Spline Models for Observational Data*, Society for Industrial and Applied Mathematics.

Ward, M. H., Jones, R. R., Brender, J. D., De Kok, T. M., Weyer, P. J., Nolan, B. T., Villanueva, C. M. & Van Breda, S. G. (2018), 'Drinking water nitrate and human health: An updated review', *International Journal of Environmental Research and Public Health* **15**(7).

Wheeler, D. C., Nolan, B. T., Flory, A. R., DellaValle, C. T. & Ward, M. H. (2015), 'Modeling groundwater nitrate concentrations in private wells in iowa', *Science of The Total Environment* **536**, 481–488.

Wikle, N. B. & Zigler, C. M. (2024), 'Causal health impacts of power plant emission controls under modeled and uncertain physical process interference', *The Annals of Applied Statistics* **18**(4), 2753–2774.

Wong, R. K. W. & Chan, K. C. G. (2017), 'Kernel-based covariate functional balancing for observational studies', *Biometrika* **105**(1), 199–213.

Wood, S. (2000), *mgcv: Mixed GAM Computation Vehicle with Automatic Smoothness Estimation*. R package version 1.9-4.
URL: *https://cran.r-project.org/package=mgcv*

Wood, S. N. (2003), 'Thin plate regression splines', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **65**(1), 95–114.

Zhao, Q. (2019), 'Covariate balancing propensity score by tailored loss functions', *The Annals of Statistics* **47**(2), 965–993.

Zimmerman, D. L. & Ver Hoef, J. M. (2022), 'On deconfounding spatial confounding in linear models', *The American Statistician* **76**(2), 159–167.

Zubizarreta, J. R. (2015), 'Stable weights that balance covariates for estimation with incomplete outcome data', *Journal of the American Statistical Association* **110**(511), 910–922.

Zubizarreta, J. R., Stuart, E. A., Small, D. S. & Rosenbaum, P. R., eds (2023), *Handbook of Matching and Weighting Adjustments for Causal Inference*, CRC Press, United States.