



Published in final edited form as:

Int Stat Rev. 2021 December ; 89(3): 605–634. doi:10.1111/insr.12452.

A Review of Spatial Causal Inference Methods for Environmental and Epidemiological Applications

Brian J. Reich¹, Shu Yang¹, Yawen Guan², Andrew B. Giffin¹, Matthew J. Miller¹, Ana Rappold³

¹Department of Statistics, North Carolina State University, Raleigh, NC 27695, USA

²Department of Statistics, University of Nebraska-Lincoln, Lincoln, NE 68583, USA

³US Environmental Protection Agency, Research Triangle Park, NC 27709, USA

Summary

The scientific rigor and computational methods of causal inference have had great impacts on many disciplines but have only recently begun to take hold in spatial applications. Spatial causal inference poses analytic challenges due to complex correlation structures and interference between the treatment at one location and the outcomes at others. In this paper, we review the current literature on spatial causal inference and identify areas of future work. We first discuss methods that exploit spatial structure to account for unmeasured confounding variables. We then discuss causal analysis in the presence of spatial interference including several common assumptions used to reduce the complexity of the interference patterns under consideration. These methods are extended to the spatiotemporal case where we compare and contrast the potential outcomes framework with Granger causality and to geostatistical analyses involving spatial random fields of treatments and responses. The methods are introduced in the context of observational environmental and epidemiological studies and are compared using both a simulation study and analysis of the effect of ambient air pollution on COVID-19 mortality rate. Code to implement many of the methods using the popular Bayesian software `OpenBUGS` is provided.

Keywords

interference; potential outcomes; propensity scores; spatial confounding; spillover

1 Introduction

Large-scale environmental and epidemiological studies often use spatially referenced data to examine the effect of treatments or exposures on a health endpoint. Examples include studying the effect of interventions on the spread of an infectious disease, pesticide application on cancer rates and lead exposure on childhood development. While standard analyses of spatial data simply estimate correlations, the ultimate goal of this research is

to establish causal relationships (e.g. Bind, 2019) to inform decision making. Therefore, developing statistical methods to establish causal relationships when data show spatial and temporal variation is invaluable to environmental science and epidemiology.

A rich literature on the theory and methods for causal inference for independent data has emerged (Bind, 2019; Hernán & Robins, 2020), but progress for spatial applications has been slow due to several analytic challenges. First, randomisation is often infeasible due to logistical or ethical concerns, and so studies rely on observational data. Second, exposure and response variables exhibit spatial correlation complicating statistical modelling and computation. Third, the treatment at one location may influence the outcomes at nearby locations, a phenomenon known as spillover or interference. These features of spatial applications violate the assumptions of standard causal inference methods and require new theory and computational tools.

Despite these challenges, major advances in spatial causal inference have been made in recent years. In this paper, we review the recent progress on spatial causal inference, evaluate and compare current methods, and suggest areas of future work. We first review methods to adjust for missing spatial confounding variables (Hodges & Reich, 2010). Most causal inference methods for observational data rely on an assumption of no missing confounding variables (i.e. unmeasured variables correlated with both the treatment and response). However, if the missing confounding variables have prominent spatial patterns, methods have been developed to mitigate the bias caused by their omission. These methods include case-control matching (e.g. Jarner et al., 2002), neighbourhood adjustments by spatial smoothing (e.g. Schnell & Papadogeorgou, 2020) and propensity-score methods (e.g. Davis et al., 2019). We review these methods and conduct a simulation study to compare their precision for estimating a causal treatment effect in the presence of a missing spatial confounding variable. A subset of the methods are applied to a study of the effect of ambient air pollution on the COVID-19 mortality rate.

A second major challenge in spatial causal inference is interference, where the treatment applied at one location affects the outcomes at other locations. For example, an intervention to reduce the emissions from a power plant would affect the air quality at the power plant, but also locations downwind. Capturing these spillover effects requires new definitions of the estimands of interest and new spatial models for the causal effects. In full generality, allowing the treatment at a site to affect the outcomes at all other sites results in an intractable estimation problem. Therefore, assumptions are required to limit the form and spatial extent of interference. We review several models for spatial interference including partial (e.g. Zigler et al., 2012) and network (e.g. Tchetgen Tchetgen et al., 2017) interference. We also discuss recent methods that combine mechanistic and spatial statistical models to anchor the causal analysis to scientific theory.

We begin reviewing these methods using cross-sectional data at a single time point and then extend these methods to the spatiotemporal data. We discuss adapting spatial methods to the spatiotemporal setting and methods specific to the temporal case such as difference-in-difference (DID) methods (e.g. Delgado & Florax, 2015) that exploit changes over time to estimate causal effects. We also compare and contrast causal methods based on the

potential outcomes framework (Rubin, 1974) with Granger causality (Granger, 1969), which is defined specifically for processes that evolve over time. We also discuss extensions of spatial methods for areal data defined at a finite number of regions (e.g. geopolitical units) to point-referenced (geostatistical) data in which case the treatment and response variables can be modelled as continuous random fields over an uncountable number of spatial locations. This requires new definitions of causal effects, new methods for matching observations for case-control studies and new models for missing spatial confounding variables and spillover effects. The paper concludes with a summary of the current literature and discussion of open problems in this rapidly advancing field.

2 Adjusting for Spatial Confounders

To ensure privacy, public health data are often made available only after aggregation to administrative or geopolitical regions. For areal data of this nature, we adapt the notation that YC_{ij} , AC_{ij} and $\mathbf{X}_{ij} = (X_{ij1}, \dots, X_{ijp})$ are the response, treatment and potential confounding variables (with $X_{ij1} = 1$ for the intercept) for observation $j \in \{1, \dots, nCi\}$ in region $i \in \{1, \dots, M\}$ for a total of $n = \sum_{i=1}^M n_i$ observations. The confounding variables in $\mathbf{X}C_{ij}$ can include both covariates specific to observation j within region i or summaries of the region i common to all nCi observations in the region. In addition to these observed variables, we allow for an unobserved confounding variable UC_i in region i , which is assumed to be a purely spatial term and thus the same for all observations in a region.

Example 1. *As a concrete example, consider an environmental epidemiology study where YC_{ij} is the birth weight of the j -th baby born in zip code i and $A_{ij} = 1$ if the average ambient air pollution concentration in the mother's zip code exceeds a high threshold and $A_{ij} = 0$ otherwise. We may adjust for known confounding variables by including the mother's age and family income in $\mathbf{X}C_{ij}$, and describe the mother's environment by including the median income and measurable environmental factors such the average concentration of other known pollutants in region i in $\mathbf{X}C_{ij}$. In this scenario, the missing spatial confounder variable UC_i might be a second pollutant unknown to the researchers. The second pollutant qualifies as a missing spatial confounder if it has a strong spatial pattern, is associated with low birth weight while holding the treatment fixed, and is correlated with the pollutant of interest, perhaps via a common source. Failing to account for this missing spatial confounder, either because its importance is unknown or data are unavailable, may inadvertently attribute the effects of the unknown pollutant to the pollutant of interest, biasing the estimator.*

In this section, we review spatial models for unknown processes such as $\mathbf{U} = (U_1, \dots, U_N)^T$ (Section 2.1). However, we argue that these standard spatial models are insufficient to remove the effects of spatial confounding, and the remainder of the section focuses on methods that explicitly consider missing spatial confounder variables. We begin with causal inference methods that would apply if \mathbf{U} were observed (Section 2.2). The remainder of the section is dedicated to methods that attempt to control for the missing confounder variable by exploiting its spatial structure.

2.1 Review of Spatial Confounding

Consider the spatial regression model

$$Y_{ij} = A_{ij}\beta + \mathbf{X}_{ij}\boldsymbol{\gamma} + U_i + \varepsilon_{ij}, \quad (1)$$

where β is the treatment effect of interest, $\boldsymbol{\gamma}$ determines the effects of the confounding variables, U_i is the spatial random effect for region i and $\varepsilon_{ij} \stackrel{iid}{\sim} \text{Normal}(0, \tau^2)$. A common approach (Banerjee et al., 2014) for areal data is to model the unobserved spatial effects using a conditionally autoregressive (CAR) model (also known as a Gaussian Markov random field model). The CAR model specifies spatial dependence in terms of the adjacencies between the regions. The full conditional distribution of the random effect for one region given all other random effects is $U_i | U_k, k \neq i \sim \text{Normal}(\rho \bar{U}_i, \sigma^2/m_i)$, where \bar{U}_i is the mean of \mathbf{U} at the m_i regions adjacent to region i , and $\rho \in (0, 1)$ and $\sigma > 0$ are spatial covariance parameters. These full conditional distributions define a multivariate normal distribution (Appendix S1) for \mathbf{U} , which we denote as $\mathbf{U} \sim \text{CAR}(\rho, \sigma)$.

The spatial regression model in Equation (1), where \mathbf{U} is modelled as a spatial process often gives very different estimates of covariate effects than the non-spatial (NS) model that excludes \mathbf{U} , especially when the treatment variable exhibits a strong spatial pattern (Reich et al., 2006; Paciorek, 2010; Hodges & Reich, 2010). However, simply accounting for spatial correlation does not resolve spatial confounding. For example, Appendix S2 describes a scenario where the bias of the posterior-mean estimator for β depends on the strength of dependence between the treatment variable and the unmeasured confounding variable but is the same whether the residuals are assumed to be independent or spatially correlated. The bias of this approach is confirmed in our simulation study (Section 2.8) when data are generated with correlation between \mathbf{U} and the treatment and response variables. This calls for methods that explicitly adjust for missing spatial confounders by blocking the dependence of \mathbf{U} on either the treatment or response variable.

2.2 Potential Outcomes Framework

In this section, we temporarily assume that U_i is observed (and thus treated the same way as \mathbf{X}_{ij}) to facilitate a review of standard NS causal inference methods. We begin with the potential outcomes framework (Rubin, 1974). Assume that the treatment A_{ij} is binary and that each unit has two potential outcomes, $Y_{ij}(0)$ and $Y_{ij}(1)$, which represent the outcomes if the unit j in region i is given treatment $A_{ij} = 0$ or $A_{ij} = 1$, respectively. Our goal is to estimate the average treatment effect (ATE),

$$\delta = \text{E} \left[\frac{1}{n} \sum_{i=1}^N \sum_{j=1}^{n_i} \{Y_{ij}(1) - Y_{ij}(0)\} \right], \quad (2)$$

where the expectation is taken with respect to both \mathbf{X}_{ij} and $\{Y_{ij}(0), Y_{ij}(1)\}$. The fundamental problem is that only one of the two potential outcomes can be observed (Holland, 1986) rendering the other as counterfactual. Therefore, assumptions are required to ensure the ATE can be identified.

This notion of potential outcomes implicitly encodes the Stable Unit Treatment Value Assumption (SUTVA Rubin, 1978).

Assumption 1 (SUTVA). *There is no interference and a single version of treatment.*

Stable Unit Treatment Value Assumption is violated under interference, where YC_{ij} depends not only on AC_{ij} but also on the treatment of other units. For instance, the birth weight of a baby in Example 1 could be influenced by the air pollution concentration both in the mother's zip code (AC_{ij}) but also in other zip codes that the mother frequents. In this case, the potential outcomes are not determined by AC_{ij} alone, and we would need to introduce a different potential outcome for each combination of the treatment variables in the mother's vicinity (Section 3).

An example of multiple versions of treatment might be if birth weight actually depends not only on whether the air pollution exceeds a high threshold but also a second extremely high threshold. In this case, AC_{ij} actually has three levels (low, high and extremely high), and there should be three potential outcomes. An analysis that collapses the two high categories into a single group with $A_{ij} = 1$ would violate SUTVA by having multiple versions of the treatment. Violation of this assumption could be rectified by assuming AC_{ij} has three categories, and thus, each unit has three potential outcomes.

While SUTVA links treatments to potential outcomes, the consistency assumption is needed to further link the potential outcomes to the observations.

Assumption 2 (Consistency). *The observed response is the potential outcome determined by the observed treatment variable, $Y_{ij} = Y_{ij}(A_{ij})$.*

In addition to these assumptions about the treatment and response variables, a standard assumption that permits unbiased estimation of the ATE is that there are no missing confounder variables other than the observed covariates XC_{ij} and the latent spatial confounder UC_i . Following Frangakis & Rubin (1999), we term this assumption as the latent ignorability assumption:

Assumption 3 (Latent ignorability). *The potential outcomes $\{YC_{ij}(0), YC_{ij}(1)\}$ and treatments AC_{ij} are independent given XC_{ij} and UC_i .*

The notion of latent ignorability was proposed by Frangakis & Rubin (1999) for identifying the complier treatment effect from randomised experiments. They assume that the missing outcomes are ignorable given the latent complier status (a complier would have been determined had the subject received the opposite treatment). Yang et al. (2019) formulate a latent ignorability assumption to deal with partially observed confounders. Although this assumption alone does not guarantee identifying the causal estimand of interest, it can help to incorporate subject matter knowledge and formulate plausible assumptions to scrutinise.

Because U is generally a latent (i.e. unknown) variable in the spatial setting, this assumption presumes that there exists some variable U that blocks dependence between the treatment variable and potential outcomes; if U is observed then this is the usual assumption

that there are no unmeasured confounding variables. This assumption implies that the confounding variables $\{XC_{ij}, UC_i\}$ are sufficient to adjust for correlation between the observed treatment and response that is due to non-randomised treatment allocation and not an actual causal effect. This requirement highlights the importance of careful evaluation of the system under study to ensure that all relevant variables are considered in XC_{ij} .

The final assumption deals with the distribution of observed treatment variables, that is, the propensity score. The propensity score is the probability of the treatment assignments, $\text{Prob}\{A_{ij} = 1 | \mathbf{X}_{ij}, U_i, Y_{ij}(0), Y_{ij}(1)\}$. Under Assumption 3, the propensity score becomes

$$e(\mathbf{X}_{ij}, U_i) = \text{Prob}(A_{ij} = 1 | \mathbf{X}_{ij}, U_i). \quad (3)$$

Assumption 4 is the standard positivity assumption on the propensity score:

Assumption 4 (Positivity). *Both $e(XC_{ij}, UC_i)$ and $1 - e(XC_{ij}, UC_i)$ are positive for all XC_{ij} and UC_i .*

This assumption implies that both $A_{ij} = 0$ and $A_{ij} = 1$ are possible under the treatment allocation mechanism, which is necessary to estimate the ATE in Equation (2), which averages over the expected potential outcome under both treatments. When this assumption is violated, Yang & Ding (2018) suggest trimming the sample.

Under Assumption 3, the propensity score is a function of known variables XC_{ij} and UC_i and can thus be estimated without knowledge of unobservable counterfactual responses. However, Assumptions 1–3 are difficult or impossible to verify empirically, and thus a causal inference requires scrutinising the study design and the processes of interest to justify that these assumptions hold. One of the main contributions of causal inference is to state explicitly the assumptions needed for an estimator to have a causal interpretation and thus guide a discussion of a study's results.

Assumptions 1–4 underlie many NS causal estimation procedures such as (augmented) inverse probability weighting (e.g. Rosenbaum & Rubin, 1983a; Robins & Greenland, 1994; Bang & Robins, 2005; Cao et al., 2009) and matching (e.g. Stuart, 2010; Abadie & Imbens, 2016). To fix ideas, we focus on the simplest approach of the linear model in Equation (1), where UC_i is observed and thus not given a spatial model. Spatial analyses often rely on parametric models because the lack of independent replications in a region complicates non-parametric methods. The parametric model in Equation (1) makes the additional assumptions of linearity and normality, but gives valid causal inference under the assumed model and Assumptions 1–4. In other words, the regression coefficient β can be interpreted as the ATE, δ . Therefore, if UC_i is observed and these assumptions hold, then the estimate of β from a standard least squares analysis has a causal interpretation. In the remainder of this section, we discuss methods to deal with unknown \mathbf{U} .

2.3 Case-Control Matching Methods

While most of the methods we discuss control for confounding at the analysis stage, a case-control study controls for confounding at the design stage. In a case-control analysis of

a binary response variable (i.e. $Y_{ij} \in \{0, 1\}$), each case ($Y_{ij} = 1$) is matched with one or more controls ($Y_{ij} = 0$) that are drawn from the same underlying population at risk. When applying this study design, investigators sample controls to resemble cases with respect to all factors that may determine the disease status except for the exposure of interest. As discussed below, this design removes the need to adjust for the matching factors at the analysis stage. Matching variables can be specific to the individual, such as age or education level. Partial control for spatial variation of risk can be achieved by matching on confounding factors that vary spatially such as the region's median income. To adjust for unmeasured spatial confounders, controls can be matched based on their proximity to the cases (Järner et al., 2002). Assuming there is replication within region ($n_{Ci} > 1$) and treatment varies within region ($AC_{ij} \neq AC_{il}$ for some j and l) then matching individuals in the same region is an effective means of adjusting for spatial confounding.

Matched case-control data are most often analysed using conditional logistic regression. Assume each case $Y_{ij} = 1$ is paired with a single control $Y_{kl} = 0$. Under the spatial logistic regression model $\text{logit}\{\text{Prob}(Y_{ij} = 1)\} = A_{ij}\beta + \mathbf{X}_{ij}\boldsymbol{\gamma} + U_i$, the log odds that $Y_{ij} = 1$ given either $Y_{ij} = 1$ or $Y_{kl} = 1$ (but not both) is

$$n_{ij} = (A_{ij} - A_{kl})\beta + (\mathbf{X}_{ij} - \mathbf{X}_{kl})\boldsymbol{\gamma} + U_i - U_k.$$

To account for variability within each pair (strata), a random intercept z_{Cij} is added so the likelihood contribution of the pair is

$$\text{Prob}(Y_{ij} = 1 | Y_{ij} + Y_{kl} = 1) = \exp(n_{ij} + z_{ij}) / \{1 + \exp(n_{ij} + z_{ij})\}.$$

Because the covariates appear in the likelihood only through the difference $\mathbf{X}_{Cij} - \mathbf{X}_{Ckl}$, the effect of covariates used for matching cannot be estimated and these covariates can be removed from the model. Similarly, if cases are paired with observations from the same region (i.e. $i = k$), then the spatial random effects \mathbf{U} do not appear in the likelihood and an NS analysis is sufficient. Thus, while the matched case-control analysis is an excellent means of controlling for confounders, its drawbacks include discarding data and not being able to estimate all covariate effects and spatial variation in risk.

Pairing observations in the same region can also be applied for continuous responses. For a continuous response, there is no natural definition of a case or control, but regressing the difference between the responses in the same region removes spatial confounding. For example, under the linear model in Equation (1), the model for the difference between responses in the same region is

$$Y_{ij} - Y_{il} = (A_{ij} - A_{il})\beta + (\mathbf{X}_{ij} - \mathbf{X}_{il})\boldsymbol{\gamma} + \tilde{\epsilon}_i, \quad (4)$$

where $\tilde{\epsilon}_i$ is independent error. Again, differencing eliminates the latent variable U_{Ci} and thus the differences can be analysed with NS methods. This approach relies on a parametric linear outcome model and matching observations in the same location. He (2018) and Yang (2018) propose alternative approaches that rely on a parametric propensity score model. He

(2018) uses weighting based on a sufficient statistic of the treatments to control cluster-level confounding, while Yang (2018) suggests calibration of treatments within clusters.

2.4 Neighbourhood Adjustments

In Equation (4), modelling the difference between observations in the same region eliminated the unmeasured confounders. In cases without replication and a missing confounder that varies smoothly across space, its effect can be reduced by removing large-scale spatial trends from the response, the treatment or both. Removing large-scale trends isolates local variation in the response, which is arguably less prone to spatial confounding than large-scale variation. In this section, we review several methods that have been proposed for removing large-trends in spatial regression.

2.4.1 Simultaneous autoregressive models—For simplicity, assume there are no replications within each region and temporarily drop the replication subscript by defining $Y_{i1} = Y_i$, $\mathbf{X}_{i1} = \mathbf{X}_i$ and $A_{i1} = A_i$. Rather than specifying the regression on the response, the simultaneous autoregressive (SAR) model first subtracts regional means

$$Y_i - \phi \bar{Y}^i = (A_i - \phi \bar{A}_i)\beta + (\mathbf{X}_i - \phi \bar{\mathbf{X}}^i)\boldsymbol{\gamma} + \varepsilon_i, \quad (5)$$

where \bar{Y}^i , \bar{A}_i and $\bar{\mathbf{X}}^i$ are the means of the response, treatment and covariates at the mCi regions adjacent to region i , ϕ is an unknown parameter and $\varepsilon_i \stackrel{iid}{\sim} \text{Normal}(0, \sigma^2)$. Taking differences reduces the effect of missing confounding variables that are constant across neighbouring regions. In vector form, Equation (5) can be expressed as $\mathbf{Y} = \mathbf{A}\boldsymbol{\beta} + \mathbf{X}\boldsymbol{\gamma} + \boldsymbol{\varepsilon}$ where the spatial covariance of $\boldsymbol{\varepsilon}$ is given in Appendix S1. Wall (2004) compares differences in covariance implied by the SAR and CAR models and finds the models produce similar regression coefficient estimates despite sometimes large differences in covariances between regions.

2.4.2 Neighbourhood adjustment via spatial smoothing—Rather than simply subtracting the mean of neighbouring sites, spatial trends can be removed by joint spatial modelling of the treatment and the missing spatial confounder. Consider the spatial regression model in Equation (1) without replicates. The bias is a result of attributing the effect of the confounder on \mathbf{Y} to the treatment variable when \mathbf{A} and \mathbf{U} are correlated (Appendix S3). Schnell & Papadogeorgou (2020) provide a set of assumptions (given in the supporting information) to identify the unmeasured confounding bias $E(U_{i1}|\mathbf{A}) = B_i(\mathbf{A})$. They model $BC_i(\mathbf{A})$ by specifying a joint distribution for \mathbf{U} and \mathbf{A} that allows each process to have a different range of spatial correlation and permits correlation between \mathbf{U} and \mathbf{A} . The confounding bias is mitigated by fitting a joint model

$$Y_i = A_i\beta - B_i(\mathbf{A}) + \mathbf{X}_i\boldsymbol{\gamma} + \mathbf{e}_{i1}, A_i = \mathbf{X}_i\boldsymbol{\alpha} + \mathbf{e}_{i2}, \quad (6)$$

where the form of $BC_i(\mathbf{A})$ and the spatial covariance of $eCi1$ and $eCi2$ are given in Appendix S3. As noted by Schnell & Papadogeorgou (2020) and was also suggested by Paciorek (2010), if the spatial scale of treatment is larger or about the same as the unmeasured confounder, the confounding bias cannot be mitigated.

2.5 Propensity Score Methods

Propensity scores are used in a wide range of causal inference methods. Assuming a binary treatment variable, the propensity score for observation j in region i is $\text{Prob}(A_{ij} = 1) = e_{ij}$. In a standard analysis, the propensity scores are modelled as a function of the known covariates \mathbf{X}_{Cij} , and the estimated propensity scores are used to alleviate the imbalance of the covariates between treatment groups. Here, we face the additional challenge that the propensity scores may depend on the unobserved spatial process, UC_i .

For example, consider the simple hierarchical model that includes the unobserved spatial process in the propensity score,

$$Y_{ij} = A_{ij}\beta + \mathbf{X}_{ij}\boldsymbol{\gamma} + U_i + \varepsilon_{ij} \quad (7)$$

$$A_{ij} \sim \text{Bernoulli}(e_{ij}) \text{ with } \text{logit}(e_{ij}) = \mathbf{X}_{ij}\boldsymbol{\alpha} + \phi U_i + V_i, \quad (8)$$

where V_i accounts for spatial patterns in treatment allocation not accounted for by the covariates or the missing confounder UC_i . To emphasise the effect of the propensity score on the response model, Equations (7)–(8) can be reparameterised ($UC_i = u_i + \psi v_i$ and $V_i = v_i - \phi u_i - \psi v_i$) as

$$Y_{ij} = A_{ij}\beta + \mathbf{X}_{ij}\boldsymbol{\gamma} + u_i + \psi v_i + \varepsilon_{ij} \quad (9)$$

$$A_{ij} \sim \text{Bernoulli}(e_{ij}) \text{ with } \text{logit}(e_{ij}) = \mathbf{X}_{ij}\boldsymbol{\alpha} + v_i. \quad (10)$$

The shared spatial random effect v_i adjusts for the missing confounder by absorbing signal in the response that can be explained by spatial trends in the treatment allocation. If the spatial trend in the treatment variable is strong and thus $A_{Cij} \approx e_{Cij}$, this method will be unstable, and it will be difficult to estimate the causal effect. The spatial random effects can be assigned priors $\mathbf{u} = (u_1, \dots, u_N)^T \sim \text{CAR}(\rho_u, \sigma_u)$ independent of $\mathbf{v} = (v_1, \dots, v_N)^T \sim \text{CAR}(\rho_v, \sigma_v)$. Fitting this joint model for the treatment and response processes is straightforward using hierarchical Bayesian methods.

A concern with this model is that some of its many parametric assumptions could be violated, invalidating inference. Another issue is that of so-called ‘feedback’, which in this context refers to information in the response influencing the posterior of the propensity scores (e.g. Zigler et al., 2013; Zigler, 2016; Saarela et al., 2016). Eliminating this feedback can be done by fitting the model in two stages, that is, first fitting the model for the treatment indicators in Equation (10) to obtain an estimate of \mathbf{v} AND then fitting Equation (9) with \mathbf{v} fixed at its first-stage estimate. Other possible remedies include ‘cutting feedback’ in the steps of the MCMC algorithm (Lunn et al., 2009; McCandless et al., 2010) or post-hoc reweighting of the posterior distribution (Saarela et al., 2015; Davis et al., 2019). These methods are discussed below.

Referring to the joint model in Equations (9)–(10), if the propensity score e_{Cij} were known and $\text{logit}(e_{ij})$ were included as a known confounder in \mathbf{X}_{Cij} , then treatment ignorability

would hold given $\mathbf{X}Cij$, and the resulting estimate of β would have a causal interpretation. Of course, the exact propensity is unknown and must be estimated. Let \hat{e}_{ij} be a first-stage propensity score estimator, for example, as estimated by fitting the spatial logistic regression model in Equation (10). The estimated propensity scores can be included in the mean of the response model to account for spatial confounding. The propensity score can be added to the response model as

$$Y_{ij} = A_{ij}\beta + \mathbf{X}_{ij}\tilde{\gamma} + U_i + f(\hat{e}_{ij}) + \varepsilon_{ij}, \quad (11)$$

where f is the logit function or more generally a non-linear function estimated by, say, smoothing splines. Given the inclusion of the propensity score, it can now be assumed that UCi and $ACij$ are conditionally independent. Assuming the model assumptions hold and the propensity score estimate is accurate, then β has a causal interpretation.

Alternatively, the propensity score estimates can be used to define strata, that is,

$$Y_{ij}|\hat{e}_{ij} \in [T_l, T_{l+1}) = S_l + A_{ij}\beta + \mathbf{X}_{ij}\gamma + U_i + \varepsilon_{ij}, \quad (12)$$

where $0 = T_1 < T_2 < \dots < T_{L+1} = 1$ define the propensity score strata, SCl encodes the unmeasured confounder effect for stratum l and UCi and $ACij$ are conditionally independent. Although the strata are defined irrespective of spatial information, the spatial random effect UCi accounts for spatial dependence.

This joint modelling framework can be extended to continuous treatment variables by replacing the the Bernoulli/logistic model for $ACij$ in Equation (10) with a normal model with $E(A_{ij}|\mathbf{X}_{ij}, v_i) = e_{ij} = \mathbf{X}_{ij}\alpha + v_i$ and $\text{Var}(A_{ij}|\mathbf{X}_{ij}, v_i) = \sigma_a^2$. This method could be fit as a joint model or in two stages where first a Gaussian spatial model for $ACij$ is fit and estimates of e_{ij} are used as generalised propensity scores (Hirano & Imbens, 2004) in the response model as in Equation (11) or (12). Generally, this model-based framework can be adapted to more complex settings as long as a model with reasonable fidelity to the data generating process can be determined and justified.

As an alternative to model-based causal adjustment, Davis et al. (2019) use imputation of potential outcomes and propensity score weighting. They first estimate propensity scores \hat{e}_{ij} using a spatial regression such as Equation (10). Then, in a second stage, they fit the response model in Equation (1), which excludes the propensity score. Rather than use the estimate of β from this analysis, they post-process the model output to remove confounding bias. They estimate the causal effect using concepts from augmented inverse probability weighting (Rosenbaum & Rubin, 1983; Robins et al., 1994; Bang & Robins, 2005; Cao et al., 2009)

$$\begin{aligned} \delta &= \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^{n_i} \delta_{ij} \delta_{ij} = \frac{1}{\hat{e}_{ij}} \{ A_{ij} Y_{ij} - (A_{ij} - \hat{e}_{ij}) \tilde{Y}_{ij1} \} \\ &\quad - \frac{1}{1 - \hat{e}_{ij}} \{ (1 - A_{ij}) Y_{ij} - (\hat{e}_{ij} - A_{ij}) \tilde{Y}_{ij0} \}, \end{aligned} \quad (13)$$

where $\tilde{Y}_{ija} = a\hat{\beta} + \mathbf{X}_{ij}\hat{\gamma} + \hat{U}_i$ is the estimated mean response setting $AC_{ij} = a$ for $a \in \{0, 1\}$. Davis et al. (2019) suggest using bootstrap sampling (which accounts for uncertainty at all stages) or a closed form large-sample variance estimator to quantify uncertainty in δ . Alternatively, in a Bayesian analysis, samples from the posterior distribution of δ can be made by computing δ for each posterior sample of $\{\beta, \gamma, \mathbf{U}\}$.

2.6 Instrumental Variables

An instrumental variable (IV) Z_i is widely used to deal with unmeasured confounding. A valid IV must (i) be associated with the treatment AC_i , (ii) not be related to the unmeasured confounder UC_i and (iii) only affect the outcome through AC_i . Figure 1 illustrates the dependence structure of the random variables. As an example, suppose AC_i and UC_i are the region's concentrations of air pollutants 1 and 2, respectively, and Y_i is the region's asthma rate. Further, assume that Pollutant 1 is the treatment of interest and is produced by both traffic and power plants, while Pollutant 2 is unmeasured and produced only by power plants. Assuming Pollutant 2 has a health impact, it is a confounding variable because it is correlated with Pollutant 1 via their shared source. A potential IV to resolve this confounding is the region's traffic density, Z_i . It could be argued that this is a valid IV because (i) it is a source of pollutant 1 and thus Z_i and AC_i are strongly correlated, (ii) it is not a source of Pollutant 2 and thus Z_i and UC_i are uncorrelated, and (iii) traffic density is unrelated to asthma rate other than via air quality.

The classic causal analysis with IVs is a two-stage least squares regression. The treatment is first regressed onto the IV and then the fitted values from this first-stage regression are used as the treatment variable in the response model. That is, if the first-stage regression gives $\hat{A}_i = \hat{\alpha}_0 + Z_i\hat{\alpha}_1 + \mathbf{X}_i\hat{\alpha}_2$, then the second stage model replaces AC_i with $Z_i\hat{\alpha}_1$, i.e., $Y_i = \hat{\alpha}_1 Z_i\beta + \mathbf{X}_i\gamma + \varepsilon_i$. This confines the treatment variable to the span of the IV, and thus to a space orthogonal to the missing confounding variable. If a valid IV can be identified then this provides a simpler means of estimating average treatment effect instead of adjusting for missing confounders than propensity scores.

Some caution has to be exercised when interpreting causal estimates based on IVs. In the observational setting, as in traffic instrument example, the investigators do not have the ability to enforce treatment (PM) based on treatment assignment (traffic). Although traffic is a major source of variation in PM, other sources can play a role, which leads to differences between intended and observed treatments among units and potentially to the heterogeneity of responses (power plants, wildfires, etc). In randomised treatment-control examples, this equates to the lack of full compliance between treatment assignment and the intake of drug. The implication is that the ATE is estimated only among those whose PM variation is explained by variation in the IV, referred to as the local average treatment effect or complier average treatment effect. Imbens & Angrist (1994) provide the criteria under which the local average treatment effect/complier average treatment effect represents the ATE.

Spatial consideration can be made in both stages of the model. Consider a continuous treatment variable and the joint model

$$Y_{ij} = \alpha_i Z_{ij} \beta + \mathbf{X}_{ij} \boldsymbol{\gamma} + U_i + \epsilon_{1ij} \quad (14)$$

$$A_{ij} = \alpha_0 + Z_{ij} \alpha_1 + \mathbf{X}_{ij} \boldsymbol{\alpha} + \phi U_i + V_i + \epsilon_{2ij}, \quad (15)$$

where $\mathbf{U} \sim \text{CAR}(\rho_U, \sigma_U)$, $\mathbf{V} \sim \text{CAR}(\rho_V, \sigma_V)$, $\epsilon_{1ij} \stackrel{iid}{\sim} \text{Normal}(0, \tau_1^2)$ and $\epsilon_{2ij} \stackrel{iid}{\sim} \text{Normal}(0, \tau_2^2)$. In Equation (14), AC_{ij} in the response model in Equation (1) is replaced by $ZC_{ij}aC1$ in the IV regression. Spatial random effects are included in both stages of the model to provide more efficient estimators of the regression coefficients and valid uncertainty quantification. This model closely resembles the joint propensity score model in Equations (7)–(8) except that only the signal in AC_{ij} that can be explained by the IV enters the response model.

The two models in Equations (14)–(15) can be fit simultaneously, although feedback effects must be considered as in the propensity score methods of Section 2.5. Alternatively, the method can be fit in two stages. The first stage is a spatial regression of AC_i onto ZC_i in Equation (15) and $\mathbf{X}C_i$ gives an estimate of $aC1$. In the second stage spatial regression of the response, $Z\hat{a}_i$ is used as the treatment variable. An important difference between the classical and this spatial IV approach is that in the spatial version the fitted values will not be strictly orthogonal to the errors UC_i . A potential remedy is the use of restricted spatial regression (Reich et al., 2006; Hodges & Reich, 2010; Hughes & Haran, 2013; Hanks et al., 2015), although these methods should be used with caution in light of the recent work of Khan & Calder (2020).

2.7 Structural Equation Modelling

Thaden & Kneib (2018) propose to adjust for spatial confounding using structural equation modelling (SEM). They introduce binary indicator variables for each spatial location in both the models for the treatment and response variables. Therefore, although motivated using structural equation modelling, they arrive at a similar model to the joint model in Equations (9)–(10). They argue that independent priors for the random effects (uC_i and vC_i in Equations 9–10) more effectively resolve spatial confounding than spatial priors. Treating the random effects as independent requires replication within region, which is not always available. However, when there is sufficient replication within regions, independent priors are preferable to spatial models because they are less constrained and thus more completely block spatial confounding.

2.8 Simulation Study

In this section, we conduct a simulation study to compare methods for adjusting for an unmeasured confounding variable. We examine how the methods compare with different levels of spatial correlation in the treatment and confounding variable, and robustness to model misspecification.

2.8.1 Data generation—We simulate data with a missing spatial confounder variable from a general form that permits performance evaluation under both correctly and incorrectly specified spatial models. The general data-generating model is

$$Y_i | A_i \overset{indep}{\sim} \text{Normal}(A_i \beta + U_i, 1) \text{ and } A_i \overset{indep}{\sim} \text{Bernoulli}[\text{expit}\{g(V_i, \phi U_i)\}], \quad (16)$$

where the spatial terms are drawn from the model $\mathbf{U} \sim \text{CAR}(\rho_U, 2)$, $\mathbf{V} \sim \text{CAR}(\rho_V, 2)$ and the transformation function g is given below. The correlation structure is determined by three parameters: $\rho_C U$ and $\rho_C V$ control the range of spatial dependence and ϕ controls the strength of spatial confounding. For simplicity, we exclude known confounders \mathbf{X}_i to isolate the effects of spatial confounding. The first four scenarios have $g(V_i, \phi U_i) = V_i + \phi U_i$ and vary $\rho_C U, \rho_C V \in \{0.90, 0.99\}$ to study the performance of the joint model when it is correctly specified. Setting the CAR dependence parameter to 0.99 gives strong spatial dependence with correlation 0.54 between adjacent regions in the center of the grid, while the value 0.90 gives moderate correlation of 0.35 between adjacent regions in the center of the grid. The final two scenarios have $\rho_U = \rho_V = 0.99$ and either nonlinear or nonstationary g . The non-linear case has $g(V_i, \phi U_i) = V_i + \phi \{U_i I(U_i > 0) - 0.63\}$ ('non-linear'). The nonstationary case has $g(V_i, \phi U_i) = V_i + \phi U_i c_i$, where c_i increases linearly from zero to one across the columns of the grid ('non-stationary'). These scenarios are included to investigate the performance of the joint model when it is misspecified. A stimulation study with more complex data-generating mechanism using the observed covariates from the data analysis in Section 2.9 is presented in Appendix S4.

We generated 100 data sets on a 30×30 square grid of regions with rook neighbours and $\beta = \varphi = 0.5$. For each data set, we fit the following models.

- NS: NS least squares, $Y_i \overset{indep}{\sim} \text{Normal}(\gamma + A_i \beta, \tau^2)$
- NS + P: NS least squares with a spline function of the propensity score, $Y_i \overset{indep}{\sim} \text{Normal}\{\gamma + A_i \beta + f(\hat{e}_i), \tau^2\}$
- S: Spatial CAR regression without confounder adjustment, $Y_i \overset{indep}{\sim} \text{Normal}(\gamma + A_i \beta + U_i, \tau^2)$
- S + P: Spatial CAR regression with a spline function of the spatial propensity score, $Y_i \overset{indep}{\sim} \text{Normal}(\gamma + A_i \beta + U_i + f(\hat{e}_i), \tau^2)$
- S + AIPW: Spatial CAR regression with post-hoc IDW debiasing step, that is, model S with post-processing as in Equation (13)
- Joint: Joint model in Equations (9)–(10)
- Cut: Joint model with feedback cut as in McCandless et al. (2010)

In these models, \hat{e}_i is computed using the spatial logistic regression in Equation (10) and f is a B-spline basis expansion with 5 degrees of freedom. In the model-fitting stage, the spatial processes \mathbf{U} and \mathbf{V} are assumed to be unknown and given priors $\mathbf{U} \sim \text{CAR}(\rho_U, \sigma_U)$ and $\mathbf{V} \sim \text{CAR}(\rho_V, \sigma_V)$. For all models, the hyperpriors are $\rho_U, \rho_V \sim \text{Uniform}(0, 1)$, all mean parameters have $\text{Normal}(0, 10)$ priors and all variances have $\text{InvGamma}(0.5, 0.005)$ priors. All of these methods are fit in `OpenBUGS`, and the code is available at <https://github.com/reich-group/SpatialCausalReview/>.

Figure 2 plots the causal effect estimates across data sets for each scenario and statistical method. As expected, the NS method without causal adjustment is biased and has low coverage in all cases. The spatial model without causal adjustment (S) provides only a small improvement. The NS model with spatial propensity score (NS + P) substantially reduces bias although its coverage remains below the nominal level. The spatial model with causal post-processing (AIPW) and the joint model that cuts feedback (Cut) have large bias and low coverage in the cases we considered.

In this simulation, the most effective methods are the spatial model with propensity score adjustment (S + P) and the full joint model (Joint). This is not surprising in the first four scenarios because the joint model was used to generate the data. In these cases, the joint model appears to have less bias than the two-stage spatial propensity score model, but both methods are similar. These models are misspecified in the final two scenarios but still outperform the other methods. Surely more extreme scenarios where these methods fail to deliver reliable inference can be devised, but these results suggest some robustness to model assumptions.

The strength of the spatial correlation in the treatment allocation process appears to be more predictive of reliable performance than model misspecification. In scenarios (b) and (d) with $\rho_U = 0.9$, all of the methods are biased and have low coverage. In these cases, the spatial model of the treatment allocation process has low predictive power, and thus, all subsequent causal adjustments are ineffective. In these cases, the unmeasured confounder cannot be explained by known covariates or spatial patterns, and there is simply no structure that can be exploited to remove its effect.

2.9. Effect of PM_{C2.5} Exposure on COVID-19 Mortality

To illustrate the spatial confounder adjustment methods, we reanalyse the data provided by Wu et al. (2020). The response YC_i for county i is the number of COVID-19 related deaths through 12 May 2020. The treatment variable AC_i is the long-term (2000–2016) average fine particulate matter (PM_{C2.5}) concentration. These variables are plotted in Figure 3, and both show strong spatial trends. The known confounder variables in \mathbf{XC}_i include $p = 15$ measures of the county’s demographic, socio-economic and climate conditions (refer to table 2 of Wu et al. (2020) for a complete list). Some covariates (number of hospital beds, body mass index and smoking rate) have a high proportion of missing values. Rather than removing the counties with missing value, which would complicate the spatial adjacency structure, we remove the covariates with missing value. Removing these covariates does not greatly affect the effect estimates (as discussed further).

Because the data set is large and the treatment is continuous, we consider only the non-spatial (‘NS’) and spatial (‘S’) models and these models with a two-stage propensity score adjustment (‘NS + P’ and ‘S + P’). The response model is $Y_i \sim \text{Poisson}(N_i \lambda_i)$, where N_i is the county’s population and λ_i is the mortality rate. Wu et al. (2020) use a quasi-Poisson model with state-level random effects; we use county-level random effects and allow these random effects to account for overdispersion. Specifically, the mortality rate is modelled as

$$\log \lambda_i = A_i \beta + \mathbf{X}_i \boldsymbol{\gamma} + U_i + f(\hat{e}_i) \quad (17)$$

where $\mathbf{U} \sim \text{CAR}(\rho_u, \sigma_u)$, \hat{e}_i is the estimated generalised propensity score (Hirano & Imbens, 2004), and f is a B-spline basis with 5 degrees of freedom. The generalised propensity score is the fitted negative log-likelihood (ignoring constants) $\hat{e}_i = (A_i - \mathbf{X}_i \hat{\alpha} - \hat{V}_i)^2$, where $\hat{\alpha}$ and \hat{V}_i are the posterior means from the model $A_i = \mathbf{X}_i \boldsymbol{\alpha} + V_i + \varepsilon_i$ and $\mathbf{V} \sim \text{CAR}(\rho_v, \sigma_v)$ and $\varepsilon_i \stackrel{iid}{\sim} \text{Normal}(0, \sigma_\varepsilon^2)$. The priors are the same as in Section 2.8. The NS models set $\rho_u = 0$ (the county-level random effect remain in the model to account for overdispersion) and the methods without a propensity score set $f(\hat{e}_i) = 0$.

The posterior distributions of β under these four models are plotted in Figure 4. The spatial models give smaller posterior mean and larger posterior variance than the NS models. Including the generalised propensity score leads to a slightly higher effect estimate for both the spatial and NS analyses. The results are generally similar to those in Wu et al. (2020) who found an 8% increase in COVID-19 related mortality for a unit increase in long-term average PMC2.5. Therefore, this analysis does not detect a missing spatial confounder that dramatically affects the causal effect estimate.

3 Methods for Spatial Interference/Spillover

Interference (also called spillover) occurs when the treatment received by one unit can affect the outcomes of other units. The ubiquitous no interference assumption in Section 2.2 was first discussed in Cox (1958), where it was referred to as ‘no interaction between units’ (Hernán & Robins, 2020). In the subsequent literature, it is often simply referenced as part of SUTVA. Despite a variety of data and treatments exhibiting interference, methods that account for interference have only recently begun to proliferate in the statistics literature, in part because interference significantly complicates the potential outcomes approach and requires additional assumptions about the form of the interference.

In this section, we review the challenges associated with accounting for interference and the current literature on this topic. In Section 3.1, we give a general formulation of potential outcomes in the presence of interference and define several quantities of interest under this framework. The remainder of the section discusses different assumptions about the nature of interference and subsequent estimation methods.

3.1 Potential Outcomes Framework

In the potential outcomes framework in Section 2.2 with binary treatment and no interference, there are two potential outcomes defined for each unit: $YC_{ij}(0)$ and $YC_{ij}(1)$. Allowing for general treatment interference entails considering 2^n potential outcomes, each corresponding to a different combination of treatments received by all units. As a result, the estimands under interference are more complicated because they require considering treatment that could be applied to multiple units. Therefore, defining the potential outcomes and estimands requires additional notation. We distinguish between the treatment applied to unit (i, j) in the observed data set, AC_{ij} , and a hypothetical treatment that could be applied

to unit (i, j) , denoted a_{Cij} . To describe potential outcomes under interference, we denote the treatments that could be applied to all n units as $\mathbf{a} = \{a_{ij}; i = 1, \dots, N; j = 1, \dots, n_i\}$, and the collection of the $n - 1$ treatments excluding a_{Cij} as \mathbf{a}_{C-ij} . The potential outcome for each unit is then written as $Y_{Cij}(a_{Cij}, \mathbf{a}_{C-ij})$, where the first term is the treatment received by unit (i, j) and the second term are the treatments received by other units.

The average treatment effect in Equation (2) is insufficient in the presence of interference as it depends only on the treatment assigned to unit (i, j) . Rather, several treatment effects are needed to provide a comprehensive summary. Halloran and Struchiner (1991,1995) and Hudgens & Halloran (2008) describe four key estimands assuming binary treatments. The direct effect (DE) is

$$DE_{ij}(\mathbf{a}_{-ij}) = E\{Y_{ij}(1, \mathbf{a}_{-ij}) - Y_{ij}(0, \mathbf{a}_{-ij})\}. \tag{18}$$

The DE compares the difference potential between outcomes for unit (i, j) with treatments $A_{ij} = 1$ versus $A_{ij} = 0$ and holding all other treatments fixed at \mathbf{a}_{C-ij} . Unlike Equation (2), there is not a single DE, as Equation (18) may be different for each unit and for all 2^{n-1} combinations of \mathbf{a}_{C-ij} . While the direct effect isolates the local treatment effect, the indirect effect (IE) measures the contribution of other treatments,

$$IE_{ij}(\mathbf{a}_{-ij}, \mathbf{a}'_{-ij}) = E\{Y_{ij}(0, \mathbf{a}_{-ij}) - Y_{ij}(0, \mathbf{a}'_{-ij})\}. \tag{19}$$

The IE is also called the spillover effect because it compares the difference between potential outcomes for two combinations of treatments for the other units, \mathbf{a}_{C-ij} and \mathbf{a}'_{-ij} , to an untreated unit with $a_{ij} = 0$ to quantify how much of the other treatment effects spill over to observation (i, j) . The DE and IE can be combined using either the total (TE) or overall effects (OE):

$$\begin{aligned} TE_{ij}(\mathbf{a}_{-ij}, \mathbf{a}'_{-ij}) &= DE_{ij}(\mathbf{a}_{-ij}) + IE_{ij}(\mathbf{a}_{-ij}, \mathbf{a}'_{-ij}) = E\{Y_{ij}(1, \mathbf{a}_{-ij}) - Y_{ij}(0, \mathbf{a}'_{-ij})\} \\ OE_{ij}(\mathbf{a}, \mathbf{a}') &= E\{Y_{ij}(a_{ij}, \mathbf{a}_{-ij}) - Y_{ij}(a'_{ij}, \mathbf{a}'_{-ij})\}. \end{aligned}$$

These effects are similar, except that the total effect always compares $a_{ij} = 1$ versus $a_{ij} = 0$, whereas the overall effect allows the local treatment to be the same for \mathbf{a} and \mathbf{a}' .

If these effects can be estimated, then the user can interrogate the fitted model by selecting any scenarios defined by \mathbf{a} and \mathbf{a}' . For example, in the context of Example 1, the DE might be computed by fixing the air pollution status of all other units \mathbf{a}_{C-ij} at their current value to determine the effect of a local action that changes the air pollution concentration in the mother's zip code but does not affect other zip codes. For the IE, we might fix all the treatment variables at their observed values except set the air pollution variable for the zip codes neighbouring a mother's zip code to one in \mathbf{a}_{C-ij} versus zero in \mathbf{a}'_{-ij} to determine the impact of changing the air pollution in zip codes where the mother spends some time outdoors. The sum of these two effects is the total effect of changing the air pollution status of all zip codes in the mother's home range (her zip code and those the mother frequents). This total effect equals the overall effect of setting $\mathbf{a} = \mathbf{1}$ for the mother's home range, $\mathbf{a}' =$

$\mathbf{0}$ for the mother's home range and both \mathbf{a} and \mathbf{a}' equal to the current value for all other zip codes.

While measures such as $DEC_{ij}(\mathbf{a}_{C-ij})$ are useful for understanding the implications of individual actions on local outcomes, assessing the overall impact of the treatment requires averaging over units and potential actions. Rather than weight all potential actions equally, they can be assigned probabilities, $\text{Prob}(\mathbf{a} = \tilde{\mathbf{a}}) = \psi(\tilde{\mathbf{a}})$. The probability mass function ψ is called the treatment policy. For example, the policy-averaged expected counterfactual outcome under treatment $a_{Cij} = a$ for unit (i, j) is

$$\bar{Y}^{ij}(a, \psi) = \sum_{\mathbf{a}_{-ij}} E\{Y_{ij}(a, \tilde{\mathbf{a}}_{-ij})\} \text{Prob}(\mathbf{a}_{-ij} = \tilde{\mathbf{a}}_{-ij} | a_{ij} = a) \quad (20)$$

where the sum is over all 2^{n-1} possible values of \mathbf{a}_{C-ij} and $\text{Prob}(\mathbf{a}_{-ij} | a_{ij} = a)$ is determined by the policy, ψ . The policy-averaged direct effect for unit (i, j) is then $\bar{Y}^{ij}(1, \psi) - \bar{Y}^{ij}(0, \psi)$, and the spatial average DE is

$$DE(\psi) = \frac{1}{n} \sum_{i=1}^N \sum_{j=1}^{n_i} \bar{Y}^{ij}(1, \psi) - \bar{Y}^{ij}(0, \psi). \quad (21)$$

Policy-averaged IE, TE and OE have similar forms.

In the context of the environmental epidemiology study described in Example 1, a simple policy is to assume that the a_{Cij} are independent over units with $\text{Prob}(a_{ij} = 1) = p$ and compute Equation (21) for several values of p to understand the DE. A policy more tailored to anticipating short-term effects of interventions in a given region is to assume that the a_{Cij} are independent over units with $\text{Prob}(a_{ij} = 1) = p_a$ if the current value of the treatment in unit (i, j) is $a_{Cij} = a$. Under this policy, a zip code currently below the threshold is converted to exceed the threshold with probability p_{C0} , and a zip code currently above the threshold is converted to below the threshold with probability $1 - p_{C1}$. The policy-averaged DE, IE and TE can be approximated via Monte Carlo simulation for a range of p_{C0} and p_{C1} to evaluate the overall effects of a campaign to reduce air pollution.

While these summaries are well defined for any potential outcome model, estimation is virtually impossible without simplifying assumptions. In the remainder of this section, we discuss several methods that exploit the spatial structure of the units to simplify the interference pattern. These methods are summarised in Figure 5.

3.2 Partial Interference

Partial interference, a term coined in Sobel (2006), or clustered interference, assumes that the units can be partitioned into groups so that interference can occur only between observations in the same group. In Example 1, partial interference might be evoked if it is reasonable to partition the zip codes into cities, and that birth weight is dependent only on the air pollution concentration in the mother's city, and not air pollution in other cities. A further parametric assumption might be that the potential outcome is a function only of

the air pollution concentration in the mother's zip code and the proportion of her city's zip codes that exceed the threshold excluding zipcode i , denoted by \tilde{a}_i . A linear model with these assumptions is

$$Y_{ij}(a_{ij}, \mathbf{a}_{-ij}) = a_{ij}\beta_1 + \tilde{a}_{ij}\beta_2 + \mathbf{X}_{ij}\boldsymbol{\gamma} + \varepsilon_{ij}, \quad (22)$$

where β_1 and β_2 entail the DE and IE, respectively. This parametric model and assumptions analogous to Assumptions, 1, 3 and 4 that \mathbf{A} is independent of all potential outcomes given the n vectors \mathbf{X}_{ij} and that $\phi(\mathbf{a}) > 0$ for all \mathbf{a} endows the parametric model

$$Y_{ij} = A_{ij}\beta_1 + \tilde{A}_{ij}\beta_2 + \mathbf{X}_{ij}\boldsymbol{\gamma} + \varepsilon_{ij} \quad (23)$$

with a causal interpretation. Of course, this model relies on strong assumptions that are difficult to verify, and thus a more flexible approach may be desirable.

There is an extensive literature that explores and expands on NS partial interference (Halloran & Struchiner, 1991; 1995; Halloran, 2012; Tchetgen Tchetgen & VanderWeele, 2012; VanderWeele et al., 2014; Liu et al., 2016; Barkley et al., 2017; Baird et al., 2018; Papadogeorgou et al., 2019). Zigler et al. (2012) assume partial interference in a spatial analysis of the health effects of environmental regulations, with clusters of sites defined by their attainment status. Perez-Heydrich et al. (2014) and Zigler & Papadogeorgou (2021) assume partial interference for groups defined by spatial proximity. Zigler & Papadogeorgou (2021) deal with additional complications that arise when the spatial resolutions of the treatment and response differ.

3.3 Spatial Network Interference

With the rise of social network data, there is a fast-growing literature on network-based interference, where observations can interfere with each other along connected edges. These methods can be applied to areal spatial data by viewing the regions as the network's nodes and defining spatial adjacency by the network's edges (e.g. Verbitsky-Savitz & Raudenbush, 2012). For example, as in the CAR model defined in Section 2.1, regions i and k can be defined as sharing an edge if they share a common border. A simple example of a model to study spatial network interference for Example 1 is Equation (23) with \tilde{A}_{ij} redefined as the mean treatment variable across the m_{ij} neighbours of region i .

More generally, Forastiere et al. (2016) propose a model that allows for interference between an observation and its immediate neighbours, creating a local interference neighbourhood around each observation. Treatment effects are estimated by conditioning on propensity scores for the direct and indirect treatment effects. Aronow et al. (2017) consider network data in a similar vein but loosen the restrictions on interference by defining an exposure mapping function. Tchetgen Tchetgen et al. (2017) examine arbitrary network interference subject only to a local Markov property that observations are conditionally independent after taking into account the nodes between them. This gives both a reasonable constraint for estimation and also allows for treatment effects to propagate through the network. Under a non-parametric structural equation model, Ogburn et al. (2020) clarify assumptions required to estimate spillover effects based on a single realisation of the network and propose

a targeted maximum likelihood estimator allowing dual dependence due to contagion and homophily (i.e. latent similarities). In a further generalisation of the spatial network interference assumption, Giffin et al. (2020) use the distance between units themselves, rather than a network approximation, to develop a generalised propensity score method to balance the spillover effect, \bar{A}_i .

3.4 Process-Based Spillover Models

Partial and network interference make assumptions that are conducive to a statistical analysis, such as the simple spillover effect in Equation (23), but are likely crude representations of reality. Mechanistic methods that encode scientific understanding of the physical processes of interest offer increased fidelity to the true interference structure. Mechanistic models are indispensable in environmental attribution studies. For example, climate models play a central role in the Intergovernmental Panel on Climate Change's conclusion that human activities likely caused the majority of the observed increase in global mean surface temperature from 1951 to 2010 (Bindoff et al., 2013). As reviewed by Hegerl & Zwiers (2011), unlike purely statistical models that are limited to scenarios observed in the data, mechanistic models can be run under counterfactual scenarios that have not, or could not, be observed. This provides a key link to the potential outcomes framework in Section 3.1.

While mechanistic models can be used to estimate direct effects, they are more critical in the presence of interference because they can rule out many of the massive number of potential spillover paths, greatly reducing the complexity of the problem. Despite these strengths, mechanistic models are only approximations and thus need to be calibrated and validated using observed data. Most relevant for our purposes is the recent work that combines mechanistic modelling with spatial statistical methods to estimate causal effects. For example, Larsen et al. (2020) fit a Bayesian geostatistical model to observed air pollution concentrations and mechanistic model output under scenarios with and without wildland fires to map the total causal effect of wildland fires on fine particulate matter concentration and the resulting health burden. Rather than post-processing model runs, Forastiere et al. (2020) build a statistical model based on a dispersion model to track air pollution from power plants in a causal analysis of health effects, and Cross et al. (2019) embed an epidemiological model for disease spread in a hierarchical Bayesian model to estimate spillover effects. These examples that highlight the important roles of mechanistic models not only likely provide more accurate estimates of causal effects but also ensure the results are tethered to scientific theory.

4 Spatiotemporal Methods

Data collected over space and time are more informative about causal relationships than cross-sectional data, because they afford the opportunity to observe variables coevolve. This reduces the potential for spurious associations. For example, if a treatment is applied in the course of the study, comparing a site's responses before and after the treatment can control for missing spatial confounding variables assuming they and their effects are time-invariant.

This narrows the search for potential confounding variables to those with a similar pattern as the treatments over both space and time.

To describe spatiotemporal methods, we adopt new notation to accommodate the temporal dimension. For simplicity, we assume areal spatial units, discrete time steps, and that each region $i \in \{1, \dots, N\}$ has a single observation at each time step $t \in \{1, \dots, T\}$. We denote the response, treatment, known and unknown confounding variables as YC_{it} , AC_{it} , \mathbf{XC}_{it} and UC_{it} , respectively. The potential outcomes framework and assumptions in Section 2.2 apply with the time step t replacing the replication number j . Similarly, many of the spatial methods in Section 2 such as matching (Section 2.3), neighbourhood adjustments (Section 2.4), propensity score methods (Section 2.5) and the IV approach (Section 2.6) apply for spatiotemporal data by viewing time as a third spatial dimension, with a different degree of correlation in this third dimension.

4.1 Testing for Missing Spatial Confounders

Janes et al. (2007) propose a method to test for unmeasured spatial confounders using spatiotemporal data. Letting \bar{A}_t denote the average of AC_{it} at time t , their approach can be adapted to our setting via the model

$$Y_{it} = \eta_1 \bar{A}_t + \eta_2 (A_{it} - \bar{A}_t) + \mathbf{X}_{it} \boldsymbol{\gamma} + \varepsilon_{it} \quad (24)$$

where \mathbf{XC}_{it} includes smooth functions of t to account for missing temporally varying confounders. In this model, η_1 and η_2 measure global and local effects of treatment, respectively, and they argue that if the estimated values of η_1 and η_2 are equal and non-zero then this represents an average causal effect of AC_{it} on YC_{it} , and that a large difference between the estimated η_1 and η_2 suggests there may be a missing spatial confounder.

4.2 Difference in Difference Methods

Difference-in-difference estimators (Ashenfelter & Card, 1985) aim to quantify the treatment effect on the increase in the mean response over time. For simplicity, we assume a binary treatment variable and two time steps ($T=2$). If the treatment at the both time steps is $a_{C1i} = a_{C2i} = a$, the increase in counterfactuals at site i is $\delta_i(a) = Y_{i2}(a) - Y_{i1}(a)$. Therefore, $\delta_i(0)$ is the increase over time in the absence of treatment, and $\delta_i(1) - \delta_i(0)$ is the increase that can be attributed to treatment. The DID average treatment effect is then

$$\delta^{DID} = \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N \{ \delta_i(1) - \delta_i(0) \} \right], \quad (25)$$

which is analogous to Equation (2) except that the outcomes are changes over time. Assume the potential outcomes follow the model $Y_{it}(a) = \beta_1 a + \beta_2 t + \beta_3 t a + \mathbf{X}_{it} \boldsymbol{\gamma} + U_{it} + \varepsilon_{it}$. Under Assumptions 1–4, the observed outcome model follows the induced linear model

$$Y_{it} = \beta_1 A_{it} + \beta_2 t + \beta_3 t A_{it} + \mathbf{X}_{it} \boldsymbol{\gamma} + U_{it} + \varepsilon_{it}. \quad (26)$$

Moreover, $\beta_3 = \delta^{DID}$ has a causal interpretation.

To render Assumptions 1–3 plausible, it is important to include information on a rich enough set of time-varying confounders in $\mathbf{X}C_{it}$ that affect both AC_{it} and YC_{it} . In the spatiotemporal settings, the time-varying confounders $\mathbf{X}C_{it}$ include the observed information on the past treatments and outcomes.

Delgado & Florax (2015) extend the spatial DIDs by assuming Markov interference, where treatment effects only impact neighbours. This gives the model

$$Y_{it} = \beta_1 A_{it} + \beta_2 t + \beta_3 t A_{it} + \beta_4 \bar{A}_{it} + \beta_5 t \bar{A}_{it} + \mathbf{X}_{it} \boldsymbol{\gamma} + U_{it} + \varepsilon_{it} \quad (27)$$

where \bar{A}_{it} is the mean of AC_{it} over the mCi neighbours of region i at time step t . The neighbourhood coefficients, β_4 and β_5 , can be viewed either as indirect spillover effects or added terms to adjust for local confounders to give more precise estimates of the direct causal effect, β_3 .

Matched wake analysis combines the DID approach with a spatiotemporal analogue to coarsened exact matching (Schutte & Donnay, 2014). It was developed in the political science literature for studying responses to whether insurgent violence in Iraq causes civilians to help the US military. In this scenario, insurgent violence leading to civilian casualties is the ‘treatment’ and violence not resulting in casualties is the ‘control’. The response is the act of turning in salvaged unexploded ordinance to the US military, so that it will not be used in an improvised explosive device. The data are divided into sliding spatiotemporal windows called ‘wakes’ and matched. Then, a DID approach is applied to the matched sample by counting the number of explosives turned in before and after events. A drawback to this method is that in some cases, the sliding windows may overlap, which will violate SUTVA.

4.3 Granger Causality

Granger causality is a fundamentally different concept from the potential outcomes framework. It is defined by temporal relationships and not potential outcomes. In a time series analysis with response YC_t , treatment AC_t , and all other relevant variables at time t , $\mathbf{X}C_t$, the treatment is said to Granger cause the response if $\text{Var}(Y_t | \mathbf{H}_t) > \text{Var}(Y_t | \mathbf{H}_t, A_1, \dots, A_{t-1})$, where the history up to time t is $\mathbf{H}_t = \{Y_1, \dots, Y_{t-1}, \mathbf{X}_1, \dots, \mathbf{X}_{t-1}\}$. In other words, Granger causality implies that given the history of all other variables, knowledge of past treatments reduces predictive uncertainty. If a linear lag L time series model is assumed, $Y_t = \sum_{i=1}^L (A_{t-i} \beta_i + \mathbf{X}_{t-i} \boldsymbol{\gamma}_i + Y_{t-i} \rho_i) + \varepsilon_t$, then the treatment is said to Granger cause the response if $\beta_i \neq 0$ for any $i \in \{1, \dots, L\}$.

Because this notion of causality is inherently defined for temporal data, extending these methods to the spatiotemporal case is straightforward. The simplest model is the linear no-interference model

$$Y_{it} = \sum_{l=1}^L \{A_{it} - l\beta_l + \mathbf{X}_{it-l}\gamma_l + Y_{it-l}\rho_l\} + U_{it} + \varepsilon_{it}, \quad (28)$$

where U_{it} is correlated over space (e.g. following a CAR or SAR distribution) but independent over time. It is also straightforward to include spillover effects by including spatial averages as covariates, that is, under a Markov interference assumption the mean of $ACit - 1$ over region is mCi neighbours could be added as a covariate.

Granger causality and Rubin causality based on potential outcomes are fundamentally different. Granger causality is defined in terms of predictive uncertainty, as might be useful to a passive observer of the system trying to maximise predictive power. In contrast, Rubin causality is defined in terms of the effects of an active intervention, as might be performed by a scientist conducting a controlled experiment. Despite their different definitions and objectives, these two approaches share similarities. White & Lu (2010) show that Granger causality is equivalent to Rubin causality for times series data with no missing confounders and valid parametric assumptions. For example, the model in Equation (28) could be motivated by Granger causality or Rubin causality with Assumptions 1–4 and further assumptions (normality, linearity, etc) on the form of the potential outcomes model. For further discussion of the similarities and differences between types of causality, refer to Holland (1986) or Eichler (2012).

5 Methods for Point-Referenced Data

Point-referenced, or geostatistical, data are not measurements of a region, but rather taken at a specific point (latitude/longitude). Let $\mathbf{s}_i \in \mathcal{R}^2$ be the spatial location corresponding to observation $i \in \{1, \dots, n\}$. The spatial regression model becomes

$$Y_i = A_i\beta + \mathbf{X}_i\gamma + U(\mathbf{s}_i) + \varepsilon_i \quad (29)$$

where the unknown confounder $U(\mathbf{s})$ is a continuous spatial processes and $\varepsilon_i \stackrel{iid}{\sim} \text{Normal}(0, \tau^2)$. This notation allows for replications at sites if, say $\mathbf{s}_i = \mathbf{s}_j$ in which case observations i and j share the spatial term $U(\mathbf{s}_i) = U(\mathbf{s}_j)$. The covariate vector \mathbf{X}_i can include spatial covariates such as the elevation at sCi and NS covariates such as the time of day the measurement was taken.

Unlike an areal data analysis as in Section 2 where the number of potential sampling locations is finite, a geostatistical analysis must consider an uncountable number of potential sampling locations $\mathbf{s} \in \mathcal{D} \subset \mathcal{R}^2$. We use the bold to denote a process over the entire spatial domain, for example, $\mathbf{U} = \{U(\mathbf{s}) : \mathbf{s} \in \mathcal{D}\}$. An unknown spatial process such as \mathbf{U} is typically assumed to be a continuous function of \mathbf{s} over \mathcal{D} and modelled as a Gaussian process with mean zero and isotropic covariance function (i.e. a covariance that depends only on the distance between locations). Although other covariance functions are available (Banerjee et al., 2014), the simplest choice is the exponential covariance function $\text{Cov}\{U(\mathbf{s}_i), U(\mathbf{s}_j)\} = \sigma^2 \exp(-d_{ij}/\rho)$, where d_{ij} is the distance between sCi and sCj . We denote this Gaussian process model as $U \sim \text{GP}(\rho, \sigma)$.

5.1 Potential Outcomes Framework

In the most general form, the potential outcomes for observation i depend on the entire spatial field of potential treatments, $\mathbf{a} = \{a(\mathbf{s}) : \mathbf{s} \in \mathcal{D}\}$. Therefore, we define the potential outcome for observation i as $YC_i(\mathbf{a})$. In the context of Example 1, $a(\mathbf{s})$ might be the air pollution concentration at spatial location \mathbf{s} , as opposed to the average concentration in a zip code. In this geostatistical setting, a mother's exposure to air pollution would integrate the concentration $a(\mathbf{s})$ along the path the mother travels. This could be estimated by a backpack the mother wears that continuously measures her local air pollution concentration. Therefore, changing $a(\mathbf{s})$ for any \mathbf{s} in the spatial domain could affect her potential outcome.

The potential outcomes framework simplifies dramatically under the no interference assumption. With a binary treatment, the two potential outcomes for unit i are $YC_i(0)$ if $a(\mathbf{s}_i) = 0$ and $YC_i(1)$ if $a(\mathbf{s}_i) = 1$. In this simple case, the potential outcomes concepts, definitions and assumptions introduced in Section 2.2 directly apply to the geostatistical setting. Many of the methods developed to adjust for missing spatial confounders described for areal data can also be applied. For example, all of the propensity score methods in Section 2.5 and IVs methods in Section 2.6 can be adapted for geostatistical data by replacing the CAR model for the missing spatial confounder with a Gaussian process model. Similarly, the adjustments based on spatial smoothing described in Section 2.4.2 can be extended to the geostatistical case as in Keller & Szpiro (2020) and Dupont et al. (2020) using splines and Guan et al. (2020) using spectral methods. Many of the other methods introduced for areal data can also be modified for geostatistical applications, as described in the remainder of this section.

5.2 Matching Methods

The matching methods described in Section 2.3 that pair observations from the same region can be applied for geostatistical data with replications at spatial locations. Distance adjusted propensity score matching (Papadogeorgou et al., 2018) can be used when there are not replications. This method alters propensity score matching (Rosenbaum & Rubin, 1983a) by using a standardised distance that combines the propensity score difference and geographic distance. The logic is that if unmeasured spatial confounders exist, then observations that are close together will have confounders that are the most alike. Similar to the neighbourhood adjustment methods, this method balances treatment and control by including geographic distances as a proxy for the unmeasured confounders in the matching process. The difference for a pair with $A_i = 1$ and $A_j = 0$ is defined as

$$D_{ij} = w|\hat{e}_i - \hat{e}_j| + (1 - w)d_{ij}/m \quad (30)$$

where \hat{e}_i and \hat{e}_j are estimated propensity scores, m is the maximum distance between pairs of locations in the study domain and $w \in [0, 1]$ is a weight. The authors propose an algorithm to select pairs with small DC_{ij} .

5.3 Regression Discontinuity

Regression discontinuity designs are generally used when treatment assignment is determined by whether the covariate value for a unit exceeds a threshold (Imbens &

Lemieux, 2008; Bor et al., 2014; Keele & Titiunik, 2015), for example, students are admitted to a college if and only if their SAT score exceeds a threshold. These cases provide a natural experiment if it can be assumed that units slightly above and slightly below the threshold are similar in every way except the treatment assignment, and thus, the difference between these groups can be attributed to the causal effect of the treatment. Natural experiments of this form often arise in environmental and epidemiological studies, where the variable being thresholded to determine treatment is the spatial location. In the context of Example 1, the treatment might be whether a state is subject to an air pollution regulation, and the objective is to determine if this affects health outcomes. Figure 6 shows a hypothetical example where treatment is applied to locations in the region $\mathcal{A} \subset \mathcal{D}$. If it can be assumed that all other factors are balanced across the border of \mathcal{A} , then comparing observations on either side of the border provides information about the causal effect of treatment. Under this assumption, the causal effect can be estimated by simply fitting the geostatistical model in Equation (29) with $A_j = 1$ if $\mathbf{s}_j \in \mathcal{A}$ and $A_j = 0$ otherwise.

5.4 Neighbourhood Adjustments

5.4.1 Stochastic partial differential equation modelling—Section 2.4 introduces the SAR model that defines the regression of the response onto the treatment after subtracting the means across neighbouring regions. The motivation for building a model on the differences is to remove the effects of spatially smooth confounding variables. The stochastic partial differential equation models of Lindgren et al. (2011) can be viewed as an extension of this idea to the continuous (geostatistical) spatial domain. In the stochastic partial differential equation framework, models are specified on the partial derivatives of the response surface, which is a generalisation of the SAR model that can be applied to differentiable functions such as \mathbf{U} . Lindgren et al. (2011) show that this approach can be used to approximate Gaussian processes with the Matérn covariance function and develop approximations that resemble the SAR covariance model.

5.5 Spillover/Interference Methods

Defining interference for geostatistical applications requires returning to the general potential outcomes formulation in Section 5.1, where the potential outcome for observation i depends on the entire field of treatments, \mathbf{a} , and is denoted as $Y_i(\mathbf{a})$. Relating the spatial field \mathbf{a} with the scalar potential outcome requires assumptions about the form of interference. A general form of the interference is

$$Y_i(\mathbf{a}) = a(\mathbf{s}_i)\beta_1 + \bar{a}(\mathbf{s}_i)\beta_2 + \mathbf{X}_i\gamma + U_i + \varepsilon_i \quad (31)$$

$$\bar{a}(\mathbf{s}_i) = \int_{\mathcal{D} \setminus \mathbf{s}_i} w(\mathbf{s}_i, \mathbf{s})a(\mathbf{s})d\mathbf{s}, \quad (32)$$

where w is a weighting function that determines the spillover effect $\bar{a}(\mathbf{s}_i)$ and β_1 and β_2 control the direct and indirect effects, respectively. Given this potential outcome model, the four causal effects (direct, indirect, total and overall) can be defined and interpreted as in Section 3.1 with \mathbf{a}_{-i} defined as the surface \mathbf{a} excluding $a(\mathbf{s}_i)$, or perhaps excluding \mathbf{a} for all sites within a small radius of \mathbf{s}_i .

The form of spillover in Equation (32) encompasses many common interference assumptions. For example, partial/cluster interference can be implemented by fixing $w(\mathbf{s}_j, \mathbf{s}) = 0$ if sites $\mathbf{s} \in \mathcal{C}_i$ and \mathbf{s} are in different groups. A structure resembling Markov/network interference assumes that $w(\mathbf{s}_j, \mathbf{s}) = 1/(\pi r^2)$ if \mathbf{s} is within radius r of $\mathbf{s} \in \mathcal{C}_i$ and $w(\mathbf{s}_j, \mathbf{s}) = 0$ otherwise. This reduces the spillover measure $\bar{a}(\mathbf{s}_i)$ to the average treatment within radius r of $\mathbf{s} \in \mathcal{C}_i$. If strict bounds on the range of interference cannot be assumed, then the weight function could be a decreasing function of the distance from $\mathbf{s} \in \mathcal{C}_i$, such as the Gaussian kernel function with $w(\mathbf{s}_i, \mathbf{s}) = \exp\{-0.5(\|\mathbf{s} - \mathbf{s}_i\|/\phi)^2\}/\sqrt{2\pi\phi^2}$.

Even after reducing the complexity of the model by selecting a simple form for the weighting function, computing the spatial integral in Equation (32) is often impossible because the treatments are only observed at a finite number of locations. One remedy is to use spatial interpolation (Kriging) to impute the treatments onto a fine grid of locations covering the spatial domain and then approximate the integrals as sums over the grid points. In this case, uncertainty about the estimated spillover variables should be accounted for using Bayesian or multiple imputation methods.

Given a form of interference and the assumption of no missing confounders, estimation of the direct and indirect effects can proceed with the usual spatial linear model. One approach to accounting for missing spatial confounders is to include spatial propensity score models for both the direct treatment AC_i and the spillover effect \bar{A}_i (Giffin et al., 2020). The propensity score for AC_i can be estimated as in the areal case with say a spatial logistic regression to give $\hat{e}(\mathbf{s}_i)$.

6 Summary and Future Work

The field of spatial causal inference has seen impressive advances in recent years. There are now methods to address the fundamental problems including accounting for missing spatial confounding variables and modelling spatial interference. However, there are many opportunities for future work that we discuss below, including orthogonalisation of confounders and treatment, combining data types, relaxing model assumptions, going beyond mean estimation and using causal estimates for decision making.

Our discussion of spatial confounding began with the observation that fixed effects (e.g. treatment) estimates can be quite different between spatial and NS regressions because the spatial covariates and spatial random effects compete for the same spatial signal. In a spatial casual analysis, the treatment variables, covariates and random effects may all have spatial patterns. One way to resolve this conflict is to restrict the spatial random effects to be orthogonal to the observed treatment variables (Reich et al., 2006; Hughes & Haran, 2013; Hanks et al., 2015; Page et al., 2017; Prates et al., 2019). However, Khan & Calder (2020) showed that this can lead to poor performance for treatment estimates. A motivation for the orthogonal regression approach is that it is easier to interpret a regression model if the signal is attributed to known quantities (e.g. Plumlee & Joseph, 2018). While this may be appealing in some settings, it is contrary to the conservative causal-inference approach that

the treatment effect is what remains after adjusting for confounding variable. Resolving these two approaches is an area of future work.

We have discussed methods for areal data (Section 2) and point-referenced/geostatistical data (Section 5) separately, but many analyses require utilising both types of data. For example, treatments may be defined at point locations (e.g. air pollution concentration) while the response variable is defined regionally (e.g. hospital admission rate by zip code). In spatial statistics, this is referred to as the change of support problem (Gotway & Young, 2002; Gelfand et al., 2010). One approach to combining data with different supports is to conceptualise the areal data as an aggregation of a continuous latent process and then specify geostatistical models such as those presented in Section 5 on the latent process. Extending these methods to the causal inference would require carefully specifying the causal estimand and devising computationally efficient methods for estimation. Zigler & Papadogeorgou (2021) may provide a template for this work.

Change of support issues also arise when the treatment is a point source, such as an oil spill, power plant or wildland fire. The effect of point source treatment variables can be direct, but their most prominent causal effects will likely be the spillover effects (Section 3) felt by nearby locations. The spillover effects can be modelled as a function of the distance from the response location to the point source or mechanistically using a mathematical dispersion model (Section 3.4). These methods can also be extended to the spatiotemporal setting using spillover effects that decay in space and time (e.g. Kim et al., 2018; 2019). Inferential methods that rely on modelling the treatment variables (e.g. propensity scores) could apply a spatial point pattern analysis (Baddeley et al., 2015), such as an inhomogeneous Poisson process model, to estimate the treatment intensity. It may also be possible to leverage work on informative sampling (Diggle et al., 2010; Pati et al., 2011) that uses a joint model for the sampling locations and the responses to reduce the effects of systematic bias in the sampling design.

Most of the methods discussed in this review rely on strong parametric assumptions such as linearity and normality. Parametric methods dominate spatial statistics because in the canonical problem with one observation at each spatial location there is insufficient data to relax these assumptions. In contrast, most causal inference methods aim to be robust to model misspecification. There is a body of work on non-parametric spatial methods (Gelfand et al., 2010; Reich & Fuentes, 2015) that might be used to relax the parametric assumptions in spatial causal inference, but these ideas have yet to be applied in this context.

We focused only on the average treatment effect, and future work is to extend spatial causal inference to other types of treatment effects. For example, extreme events are often the most impactful in environmental studies, and thus, it would be of great interest to extend causal inference ideas to spatial quantile regression (e.g. Reich et al., 2011; Reich, 2012; Lum & Gelfand, 2012) or extreme value analysis (e.g. Davison & Huser, 2019). Another simplification made throughout the review is that the confounder and treatment effects are the same throughout the spatial domain. A more general approach is a locally adaptive model with spatially varying coefficients (Gelfand et al., 2003), which would be a spatial application of conditional treatment effects (Wu et al., 2020).

Ultimately, causal effect estimates can be used to influence decision making. An area of future work is to use these estimates to derive individualised/localised treatment rules. This is complicated in the spatial case by interference between regions that require considering simultaneously assigning the treatments to all regions to achieve optimality. Laber et al. (2018) and Guan et al. (2020) propose a policy-search method for optimal treatment allocation for spatiotemporal problems, but a general theory awaits development.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

ACKNOWLEDGEMENTS

This work was partially supported by the National Institutes of Health (R01ES031651-01, R01ES027892-01) and King Abdullah University of Science and Technology (3800.2). The research described in this article has been reviewed by the Center for Public Health and Environmental Assessment, US Environmental Protection Agency and approved for publication. Approval does not signify that the contents necessarily reflect the views and the policies of the Agency, nor does mention of trade names of commercial products constitute endorsement or recommendation for use. The authors declare that they have no conflict of interest.

References

- Abadie A & Imbens GW 2016. Matching on the estimated propensity score. *Econometrica*, 84, 781–807.
- Aronow PM, Samii C et al. 2017. Estimating average causal effects under general interference, with application to a social network experiment. *Ann. Appl. Stat.*, 11(4), 1912–1947.
- Ashenfelter O & Card D 1985. Using the longitudinal structure of earnings to estimate the effect of training programs. *Rev. Econ. Stat.*, 67(4), 648–660.
- Baddeley A, Rubak E & Turner R 2015. *Spatial Point Patterns: Methodology and Applications With R*. Boca Raton, Florida: Chapman and Hall/CRC.
- Baird S, Bohren JA, McIntosh C & Özler B 2018. Optimal design of experiments in the presence of interference. *Rev. Econ. Stat.*, 100(5), 844–860.
- Banerjee S, Carlin BP & Gelfand AE 2014. *Hierarchical Modeling and Analysis for Spatial Data*. Boca Raton, Florida: Chapman and Hall/CRC.
- Bang H & Robins JM 2005. Doubly robust estimation in missing data and causal inference models. *Biometrics*, 61, 962–973. [PubMed: 16401269]
- Barkley BG, Hudgens MG, Clemens JD, Ali M & Emch ME 2017. Causal inference from observational studies with clustered interference. arXiv preprint arXiv:1711.04834.
- Bind M-A 2019. Causal modeling in environmental health. *Ann. Rev. Public Health*, 40, 23–43. [PubMed: 30633715]
- Bindoff NL, Stott PA, AchutaRao KM, Allen MR, Gillett N, Gutzler D, Hansingo K, Hegerl G, Hu Y, Jain S & Mokhov II 2013. *Detection and Attribution of Climate Change: From Global to Regional*. New York, NY, USA: Cambridge University Press.
- Bor J, Moscoe E, Mutevedzi P, Newell M-L & Bärnighausen T 2014. Regression discontinuity designs in epidemiology: causal inference without randomized trials. *Epidemiology (Cambridge, Mass.)*, 25(5), 729. [PubMed: 25061922]
- Cao W, Tsiatis AA & Davidian M 2009. Improving efficiency and robustness of the doubly robust estimator for a population mean with incomplete data. *Biometrika*, 96, 723–734. [PubMed: 20161511]
- Cox DR 1958. *Planning of Experiments*. New York, NY, USA: Wiley.
- Cross PC, Prosser DJ, Ramey AM, Hanks EM & Pepin KM 2019. Confronting models with data: the challenges of estimating disease spillover. *Phil. Trans. R. Soc. B*, 374(1782), 20180435. [PubMed: 31401965]

- Davis ML, Neelon B, Nietert PJ, Hunt KJ, Burgette LF, Lawson AB & Egede LE 2019. Addressing geographic confounding through spatial propensity scores: a study of racial disparities in diabetes. *Stat. Methods Med. Res.*, 28(3), 734–748. [PubMed: 29145767]
- Davison AC & Huser R 2019. *Spatial Extremes*. Boca Raton, Florida: CRC Press.
- Delgado MS & Florax RJGM 2015. Difference-in-differences techniques for spatial data: local autocorrelation and spatial interaction. *Econ. Lett.*, 137, 123–126.
- Diggle PJ, Menezes R & Su T 2010. Geostatistical inference under preferential sampling. *J. R. Stat. Soc.: Ser. C (Appl. Stat.)*, 59(2), 191–232.
- Dupont E, Wood SN & Augustin N 2020. Spatial+: a novel approach to spatial confounding. arXiv preprint arXiv:2009.09420.
- Eichler M 2012. Causal inference in time series analysis. In *Causality: Statistical Perspectives and Applications*, 1st edn., Ed. Carlo Berzuini LB pp. 326–354. Chichester, UK: Wiley Online Library.
- Forastiere L, Airoidi EM & Mealli F 2016. Identification and estimation of treatment and interference effects in observational studies on networks. arXiv preprint arXiv:1609.06245.
- Forastiere L, Mealli F & Zigler C 2020. Bipartite interference and air pollution transport: estimating health effects of power plant interventions. Submitted.
- Frangakis CE & Rubin DB 1999. Addressing complications of intention-to-treat analysis in the combined presence of all-or-none treatment-noncompliance and subsequent missing outcomes. *Biometrika*, 86(2), 365–379.
- Gelfand AE, Diggle P, Guttorp P & Fuentes M 2010. *Handbook of Spatial Statistics*. Boca Raton, Florida: CRC Press.
- Gelfand AE, Kim H-J, Sirmans CF & Banejee S 2003. Spatial modeling with spatially varying coefficient processes. *J. Am. Stat. Assoc.*, 98(462), 387–396.
- Giffin A, Reich BJ, Yang S & Rappold AG 2020. Generalized propensity score approach to causal inference with spatial interference. arXiv preprint arXiv:2007.00106.
- Gotway CA & Young LJ 2002. Combining incompatible spatial data. *J. Am. Stat. Assoc.*, 97(458), 632–648.
- Granger CW 1969. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: J. Economet. Soc.*, 37, 424–438.
- Guan Y, Page GL, Reich BJ, Ventrucchi M & Yang S 2020. A spectral adjustment for spatial confounding. arXiv preprint arXiv:2012.11767.
- Guan Q, Reich BJ & Laber EB 2020. A spatiotemporal recommendation engine for malaria control. arXiv preprint arXiv:2003.05084.
- Halloran ME 2012. The minicommunity design to assess indirect effects of vaccination. *Epidemiologic Methods*, 1(1), 83–105. [PubMed: 23599908]
- Halloran ME & Struchiner CJ 1991. Study designs for dependent happenings. *Epidemiology*, 2, 331–338. [PubMed: 1742381]
- Halloran ME & Struchiner CJ 1995. Causal inference in infectious diseases. *Epidemiology*, 6, 142–151. [PubMed: 7742400]
- Hanks EM, Schliep EM, Hooten MB & Hoeting JA 2015. Restricted spatial regression in practice: geostatistical models, confounding, and robustness under model misspecification. *Environmetrics*, 26(4), 243–254.
- He Z 2018. Inverse conditional probability weighting with clustered data in causal inference. arXiv preprint arXiv:1808.01647.
- Hegerl G & Zwiers F 2011. Use of models in detection and attribution of climate change. *Wiley Interdiscipl. Rev.: Climate Change*, 2(4), 570–591.
- Hernán MA & Robins JM 2020. *Causal Inference: What if*. Chapman & Hall/CRC: Boca Raton.
- Hirano K & Imbens GW 2004. The propensity score with continuous treatments. *Appl. Bayesian Model. Causal Inference Incomplete-Data Perspect.*, 22, 73–84.
- Hodges JS & Reich BJ 2010. Adding spatially-correlated errors can mess up the fixed effect you love. *Am. Stat.*, 64(4), 325–334.
- Holland PW 1986. Statistics and causal inference. *J. Am. Stat. Assoc.*, 81, 945–960.

- Hudgens MG & Halloran ME 2008. Toward causal inference with interference. *J. Am. Stat. Assoc.*, 103, 832–842. [PubMed: 19081744]
- Hughes J & Haran M 2013. Dimension reduction and alleviation of confounding for spatial generalized linear mixed models. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)*, 75(1), 139–159.
- Imbens GW & Angrist JD 1994. Identification and estimation of local average treatment effects. *Econometrica*, 62(2), 467–475.
- Imbens GW & Lemieux T 2008. Regression discontinuity designs: a guide to practice. *J. Economet.*, 142(2), 615–635.
- Janes H, Dominici F & Zeger SL 2007. Trends in air pollution and mortality: an approach to the assessment of unmeasured confounding. *Epidemiology*, 18, 416–423. [PubMed: 17568215]
- Jarner MF, Diggle P & Chetwynd AG 2002. Estimation of spatial variation in risk using matched case-control data. *Biometrical J.: J. Math. Methods Biosci.*, 44(8), 936–945.
- Keele LJ & Titiunik R 2015. Geographic boundaries as regression discontinuities. *Political Anal.*, 23(1), 127–155.
- Keller JP & Szpiro AA 2020. Selecting a scale for spatial confounding adjustment. *J. R. Stat. Soc.: Ser. A*, 183, 1121–1143.
- Khan K & Calder CA 2020. Restricted spatial regression methods: implications for inference. *J. Am. Stat. Assoc.*, 1–13.
- Kim M, Pains D & Jurdak R 2018. Causal inference in disease spread across a heterogeneous social system. arXiv preprint arXiv:1801.08133.
- Kim M, Pains D & Jurdak R 2019. Modeling stochastic processes in disease spread across a heterogeneous social system. *Proc. Nat. Acad. Sci.*, 116(2), 401–406. [PubMed: 30587583]
- Laber EB, Meyer NJ, Reich BJ, Pacifici K, Collazo JA & Drake JM 2018. Optimal treatment allocations in space and time for on-line control of an emerging infectious disease. *J. R. Stat. Soc.: Ser. C (Appl. Stat.)*, 67(4), 743–789.
- Larsen A, Yang S, Reich BJ & Rappold AG 2020. A spatial causal analysis of wildland fire-contributed PM_{2.5} using numerical model output. arXiv preprint arXiv:2003.06037.
- Lindgren F, Rue H & Lindström J 2011. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *J. R. Stat. Soc.: Ser. B (Stat. Methodol.)*, 73(4), 423–498.
- Liu L, Hudgens MG & Becker-Dreps S 2016. On inverse probability-weighted estimators in the presence of interference. *Biometrika*, 103(4), 829–842. [PubMed: 29422692]
- Lum K & Gelfand AE 2012. Spatial quantile multiple regression using the asymmetric laplace process. *Bayesian Anal.*, 7(2), 235–258.
- Lunn D, Best N, Spiegelhalter D, Graham G & Neuenschwander B 2009. Combining MCMC with ‘sequential’ PKPD modelling. *J. Pharmacokinetics Pharmacodyn.*, 36(1), 19.
- McCandless LC, Douglas IJ, Evans SJ & Smeeth L 2010. Cutting feedback in Bayesian regression adjustment for the propensity score. *Int. J. Biostat.*, 6(2), 1–22.
- Ogburn EL, Sofrygin O, Diaz I & van der Laan MJ 2020. Causal inference for social network data. arXiv preprint arXiv:1705.08527.
- Paciorek CJ 2010. The importance of scale for spatial-confounding bias and precision of spatial regression estimators. *Stat. Sci.*, 25(1), 107–125. [PubMed: 21528104]
- Page GL, Liu Y, He Z & Sun D 2017. Estimation and prediction in the presence of spatial confounding for spatial linear models. *Scandinavian J. Stat.*, 44(3), 780–797.
- Papadogeorgou G, Choirat C & Zigler CM 2018. Adjusting for unmeasured spatial confounding with distance adjusted propensity score matching. *Biostatistics*, 20(2), 256–272.
- Papadogeorgou G, Mealli F & Zigler CM 2019. Causal inference with interfering units for cluster and population level treatment allocation programs. *Biometrics*, 75(3), 778–787. [PubMed: 30859545]
- Pati D, Reich BJ & Dunson DB 2011. Bayesian geostatistical modelling with informative sampling locations. *Biometrika*, 98(1), 35–48. [PubMed: 23956461]
- Perez-Heydrich C, Hudgens MG, Halloran ME, Clemens JD, Ali M & Emch ME 2014. Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70, 731–741. [PubMed: 24845800]

- Plumlee M & Joseph VR 2018. Orthogonal Gaussian process models. *Stat. Sin*, 28, 601–619.
- Prates MO, Assunção RM & Rodrigues EC 2019. Alleviating spatial confounding for areal data problems by displacing the geographical centroids. *Bayesian Anal*, 14(2), 623–647.
- Reich BJ 2012. Spatiotemporal quantile regression for detecting distributional changes in environmental processes. *J. R. Stat. Soc.: Ser. C (Appl. Stat.)*, 61(4), 535–553.
- Reich BJ & Fuentes M 2015. Spatial Bayesian nonparametric methods. In *Nonparametric Bayesian Inference in Biostatistics*. New York: Springer, pp. 347–357.
- Reich BJ, Fuentes M & Dunson DB 2011. Bayesian spatial quantile regression. *J. Am. Stat. Assoc.*, 106(493), 6–20. [PubMed: 23459794]
- Reich BJ, Hodges JS & Zadnik V 2006. Effects of residual smoothing on the posterior of the fixed effects in disease-mapping models. *Biometrics*, 62(4), 1197–1206. [PubMed: 17156295]
- Robins JM & Greenland S 1994. Adjusting for differential rates of prophylaxis therapy for PCP in high-versus low-dose AZT treatment arms in an AIDS randomized trial. *J. Am. Stat. Assoc.*, 89, 737–749.
- Robins JM, Rotnitzky A & Zhao LP 1994. Estimation of regression coefficients when some regressors are not always observed. *J. Am. Stat. Assoc.*, 89, 846–866.
- Rosenbaum PR & Rubin DB 1983. Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome. *J. R. Stat. Soc. Ser. B*, 45, 212–218.
- Rosenbaum PR & Rubin DB 1983a. The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70, 41–55.
- Rubin DB 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *J Educational Psychol*, 66, 688–701.
- Rubin DB 1978. Bayesian inference for causal effects: the role of randomization. *Ann. Statist*, 6, 34–58.
- Saarela O, Belzile LR & Stephens DA 2016. A Bayesian view of doubly robust causal inference. *Biometrika*, 103(3), 667–681.
- Saarela O, Stephens DA, Moodie EEM & Klein MB 2015. On Bayesian estimation of marginal structural models. *Biometrics*, 71(2), 279–288. [PubMed: 25677103]
- Schnell P & Papadogeorgou G 2020. Mitigating unobserved spatial confounding when estimating the effect of supermarket access on cardiovascular disease deaths. *Ann. Appl. Stat*, 14, 2069–2095.
- Schutte S & Donnay K 2014. Matched wake analysis: finding causal relationships in spatiotemporal event data. *Polit. Geography*, 41, 1–10.
- Sobel ME 2006. What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *J. Am. Stat. Assoc.*, 101, 1398–1407.
- Stuart EA 2010. Matching methods for causal inference: a review and a look forward. *Stat. Sci*, 25, 1–21. [PubMed: 20871802]
- Tchetgen Tchetgen EJ, Fulcher I & Shpitser I 2017. Auto-g-computation of causal effects on a network. arXiv preprint arXiv:1709.01577.
- Tchetgen Tchetgen EJ & VanderWeele TJ 2012. On causal inference in the presence of interference. *Stat. Methods Med. Res.*, 21(1), 55–75. [PubMed: 21068053]
- Thaden H & Kneib T 2018. Structural equation models for dealing with spatial confounding. *Am. Statistician*, 72(3), 239–252.
- VanderWeele TJ, Tchetgen Tchetgen EJ & Halloran ME 2014. Interference and sensitivity analysis. *Stat. Sci.: A Rev. J. Inst. Math. Stat*, 29(4), 687.
- Verbitsky-Savitz N & Raudenbush SW 2012. Causal inference under interference in spatial settings: a case study evaluating community policing program in Chicago. *Epidemiologic Methods*, 1(1), 107–130.
- Wall MM 2004. A close look at the spatial structure implied by the CAR and SAR models. *J. Stat. Plann. Inference*, 121(2), 311–324.
- White H & Lu X 2010. Granger causality and dynamic structural systems. *J. Financial Economet*, 8(2), 193–243.
- Wu X, Nethery RC, Sabath BM, Braun D & Dominici F 2020. Exposure to air pollution and COVID-19 mortality in the United States. medRxiv.

- Wu L, Yang S, Reich BJ & Rappold AG 2020. Estimating spatially varying health effects in app-based citizen science research. arxiv.org/pdf/2005.12017v2.pdf.
- Yang S 2018. Propensity score weighting for causal inference with clustered data. *J. Causal Inference*, 6, 1–19. 10.1515/jci-2017-0027
- Yang S & Ding P 2018. Asymptotic inference of causal effects with observational studies trimmed by the estimated propensity scores. *Biometrika*, 105, 487–493.
- Yang S, Wang L & Ding P 2019. Causal inference with confounders missing not at random. *Biometrika*, 106, 875–888.
- Zigler CM 2016. The central role of Bayes' theorem for joint estimation of causal effects and propensity scores. *Am. Stat*, 70(1), 47–54. [PubMed: 27482121]
- Zigler CM, Dominici F & Wang Y 2012. Estimating causal effects of air quality regulations using principal stratification for spatially correlated multivariate intermediate outcomes. *Biostatistics*, 13, 289–302. [PubMed: 22267524]
- Zigler CM & Papadogeorgou G 2021. Bipartite causal inference with interference. *Stat. Sci*, 36, 109. [PubMed: 33867656]
- Zigler CM, Watts K, Yeh RW, Wang Y, Coull BA & Dominici F 2013. Model feedback in Bayesian propensity score estimation. *Biometrics*, 69(1), 263–273. [PubMed: 23379793]




Figure 1. A directed acyclic graph (DAG) represents the dependence of the random variables. Z is the instrumental variable, A is the treatment, Y is the outcome, X is the observed confounder and U is the unobserved confounder.

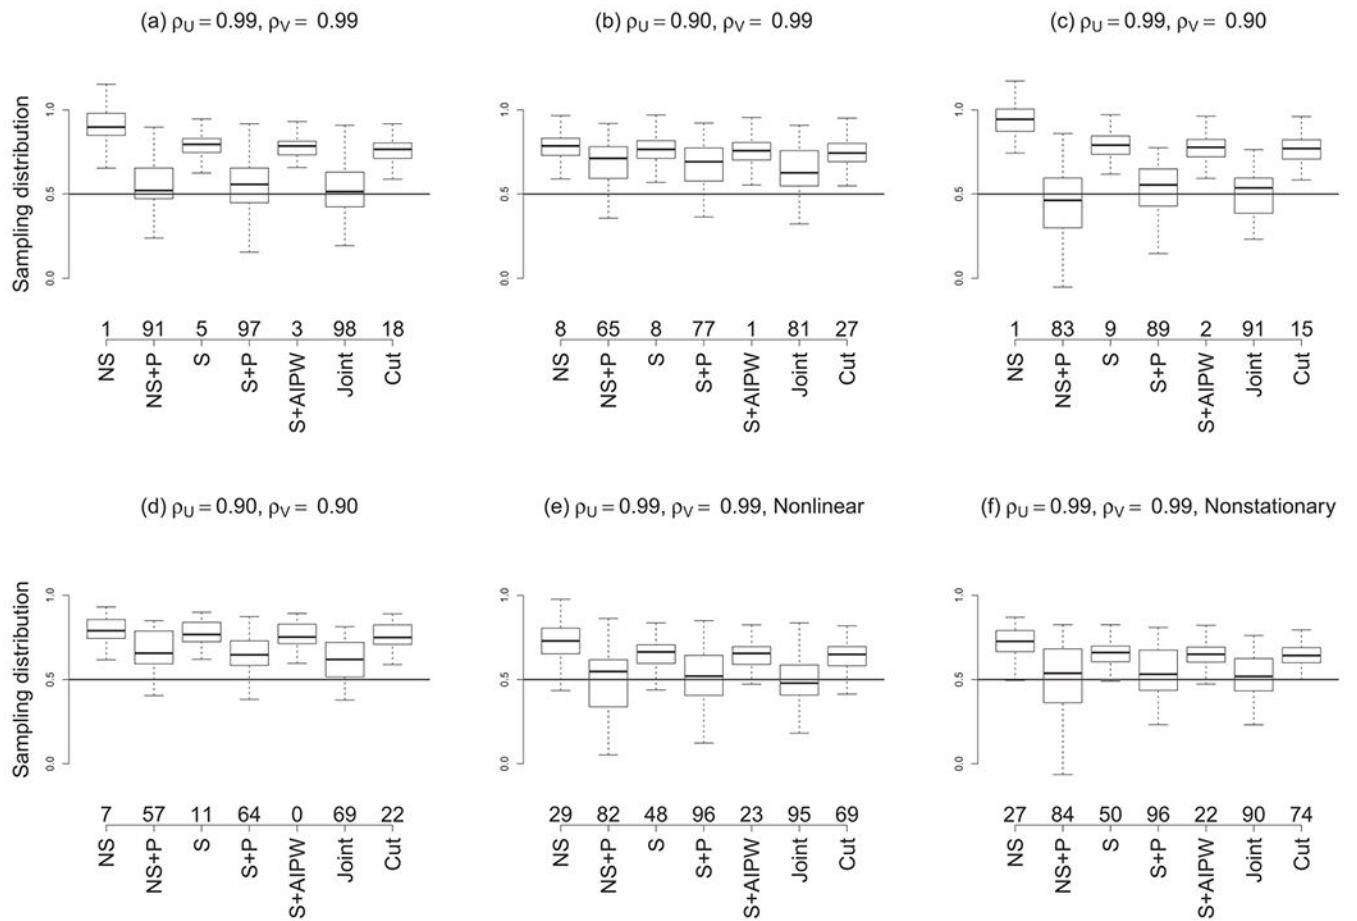
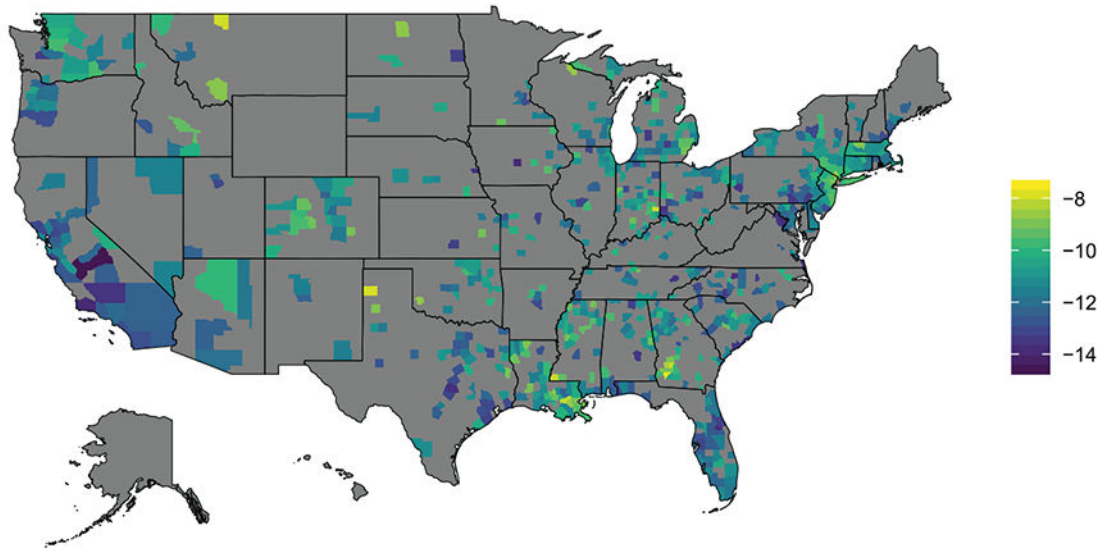
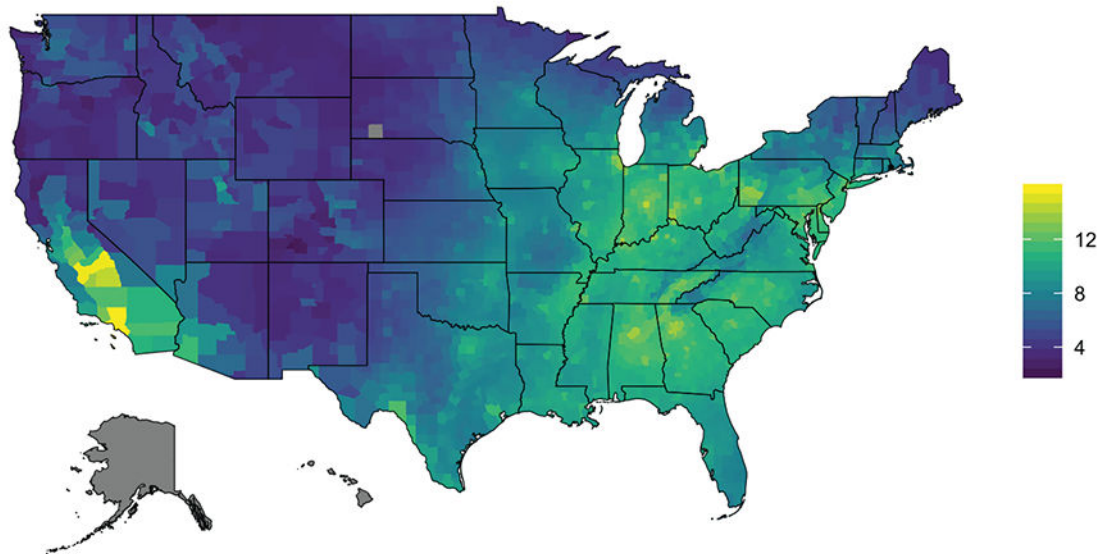


Figure 2. Simulation study results. The boxplots summarise the sampling distribution of the causal estimates across data sets and the solid line at 0.5 is the true value. The scenarios vary by the spatial dependence parameter of the confounder (ρ_U) and treatment (ρ_V) variables, and whether the joint model is misspecified. The competing methods are defined in Section 2.8. The empirical coverage of 95% credible intervals for the causal effect are given above the model labels.

(a) Log mortality rate

(b) Average PM2.5 concentration ($\mu\text{g}/\text{m}^3$)**Figure 3.**

Plots of the COVID-19/PM2.5 data. Panel (a) plots the sample log COVID-19 mortality rate, $\log(Y_i/N_i)$, through 12 May 2020 with gray denoting no observed deaths ($Y_i = 0$); panel (b) maps the long-term (2000–2016) average fine particulate matter (PM2.5) concentration. Alaska and Hawaii are excluded from the study.






Figure 4.


Causal effect estimate for the COVID-19/PM2.5 analysis. Posterior distribution of the log relative risk of an increase of $1 \mu\text{g}/\text{m}^3$ in long-term average (PM2.5) (β) on a county's COVID-19 mortality rate. The four models are defined by whether they are non-spatial ('NS') or spatial ('S') and whether they include a spatial propensity score ('+ P').




(a) No interference



(b) General interference



(c) Partial interference



(d) Spatial network interference

Figure 5. Variable dependencies under different forms of interference. Spatial location is indicated horizontally. Indirect effects are shown as dashed lines, and confounding relationships are shown as solid lines at a location and dotted lines across locations. A is the treatment, Y is the outcome and X is the observed confounder.




Figure 6.

Illustration of regression discontinuity. The treatment region \mathcal{A} is the region above the curve, the points are the sample locations s_i with samples with $A_i = 1$ filled and the background color is the mean function $A(s)\beta + U(s)$ where $A(s)$ indicates that $s = (s_1, s_2) \in \mathcal{A}$.