

Generalized propensity score approach to causal inference with spatial interference

A. Giffin¹  | B. J. Reich¹  | S. Yang¹  | A. G. Rappold²

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

²Environmental Protection Agency, Chapel Hill, North Carolina, USA

Correspondence

Andrew Giffin, Department of Statistics, North Carolina State University, 2311 Stinson Drive, Raleigh, NC 27607, USA.
Email: giffin.andrew@gmail.com

Abstract

Many spatial phenomena exhibit interference, where exposures at one location may affect the response at other locations. Because interference violates the stable unit treatment value assumption, standard methods for causal inference do not apply. We propose a new causal framework to recover direct and spill-over effects in the presence of spatial interference, taking into account that exposures at nearby locations are more influential than exposures at locations further apart. Under the no unmeasured confounding assumption, we show that a generalized propensity score is sufficient to remove all measured confounding. To reduce dimensionality issues, we propose a Bayesian spline-based regression model accounting for a sufficient set of variables for the generalized propensity score. A simulation study demonstrates the accuracy and coverage properties. We apply the method to estimate the causal effect of wildland fires on air pollution in the Western United States over 2005–2018.

KEYWORDS

air pollution, causal inference, interference, spatial process, wildfire

1 | INTRODUCTION

Understanding spatial processes in the environmental and health sciences has taken on new importance as we grapple with emerging ecological and epidemiological issues. Much of the research in these areas are associative in nature despite the effects of interests being causal (Bind, 2019). This is not only a result of both the frequent necessity of using observational data, but also the difficulty of implementing causal inference tools on data that exhibit spatial dependence and, in particular, interference. Interference is the phenomenon in which exposures or treatments at one location may affect the response at other locations. Naturally, with spatially dependent processes, an

exposure may impact the response nearby, leading to interference.

An example of spatial interference is the relationship between wildland fires and air pollution. Treating wildland fires as the exposure and pollution as the response, it is clear that the exposure can substantially impact the response at the location of exposure and at distant locations. In this example, all available data are observational, and therefore isolating average causal exposure effects requires accounting for confounding variables. Even in the ideal case where all potential confounders are observed across locations, it is unclear how to condition on these confounders without knowing their specific spatial relationships with the exposure and response. Conditioning on

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2022 The Authors. *Biometrics* published by Wiley Periodicals LLC on behalf of International Biometric Society.

all locations, which is one way around this, is impractical for all but the smallest studies.

The difficulty that arises from interference in the context of spatially dependant processes is immediately apparent from the vantage of the potential outcomes framework developed by Rubin (1974). For a binary exposure without interference, there are two unit-level potential outcomes to consider. Under general exposure interference, there are 2^n unit-level potential outcomes to consider, where n is the total number of units, because each exposure permutation across all units represents a distinct exposure. In the case of geostatistical models that contain uncountably many spatial locations, the problem becomes even more intractable. For this reason, beginning with Cox (1958), much of the causal inference literature assumes away interference. The no-interference assumption is usually invoked as part of the stable unit treatment value assumption (Rubin, 1980).

Relaxations to the no-interference assumption generally involve placing assumptions on the form of interference. Partial interference, a term coined by Sobel (2006), was the first relaxation developed, specifically for modeling vaccination exposures that are known to induce herd immunity. This assumption defines disjoint groups or clusters a priori that may exhibit interference, but precludes interference between groups. This form of interference was originally considered with experimental data by Halloran and Struchiner (1991, 1995), but expanded to nonrandomized data by Hudgens and Halloran (2008), Tchetgen and VanderWeele (2012), Liu and Hudgens (2014), and Papadogeorgou et al. (2019). The dual nature of this form of interference allows for information on both the direct exposure effects as well as the indirect or spill-over effects from interference. Additionally, the deluge of network data has resulted in a literature that allows for interference along edges of a pre-specified graph (Aronow & Samii, 2017; Athey et al., 2018; Forastiere et al., 2018, 2020; Karwa & Airolidi, 2018). In a method largely analogous to ours, Forastiere et al. (2018) use a generalized propensity score method to estimate the effect of individual exposures, as well as the effects from “neighborhood” interference—that is, interference effects from subjects connected on the prespecified graph. Aronow and Samii (2017) define an exposure mapping similar to that of Manski (2013), which allows for arbitrary interference forms.

Spatially indexed data have been analyzed using both the partial interference and network interference strategies. For naturally clustered spatial data, the partial interference assumption can be used, for example, as in Zigler and Papadogeorgou (2021), Zigler et al. (2012), and Perez-Heydrich et al. (2014). Spatial data can also be simplified to the network setting. For areal data, this often entails

creating a graph with edges between neighboring units, as in Verbitsky-Savitz and Raudenbush (2012). This, however, discards information about the distance between units. Additional causal techniques for spatial data are outlined in Papadogeorgou et al. (2018), Papadogeorgou and Dominici (2020), and Schnell and Papadogeorgou (2020).

Despite these advances, there has been little exploration of strictly spatial forms of interference. To fill this gap in the literature, we propose a new framework to recover causal direct and spill-over effects in the presence of spatial interference, while taking into account the high dimensionality of the problem. We develop a generalized propensity score to account for spatial dependence in the distribution of exposure (Hirano & Imbens, 2004; Imbens, 2000; Imai & Van Dyk, 2004). To further reduce the size of the problem, we propose a model that accounts for a sufficient set of summary variables rather than the full generalized propensity score. Notably, the exposure mapping of Aronow and Samii (2017) can be given spatial form, but the mapping is assumed to be known a priori.

The proposed approach has a number of advantages over using a partial interference or network interference assumption. The partial interference assumption is only reasonable for limited cases when the data naturally cluster a significant distance apart. Moreover, the partial interference grouping must be specified a priori. The network interference assumption requires that spill-over treatments propagate along a prespecified network, which can allow for spill-over effects that are not monotonic over distance—a key feature of our spatial method. Of particular note, Forastiere et al. (2018) use a generalized propensity score similar to ours over a network. Using a “community detection” algorithm, the “neighborhood” that contributes to the spill-over effect can be assessed from the data, after which weighting between points can be given with a prespecified “exposure mapping” function. In contrast, our method simultaneously estimates the spatial range and the weighting of the nearby points, using an assumption that their spill-over effect is proportional to a kernel of distance. Finally, the highly general methodology of Aronow and Samii (2017) again assumes a known exposure mapping, which would not allow estimating the range of exposure. Our proposed method retains all spatial information, and allows for the kernel range to be estimated concurrently.

2 | POTENTIAL OUTCOMES, INTERFERENCE, AND IDENTIFICATION

Assume that data are available at n spatial locations $s \in \{s_1, \dots, s_n\} \subset \mathcal{D} \subset \mathbb{R}^2$. For spatial location s , define $X_s \in$

\mathbb{R}^p as the relevant covariates and $Y_s \in \mathbb{R}^1$ the response. We will consider both real-valued and binary exposures A_s . We use subscript D to refer to the full fields of random variables, for example, $X_D = \{X_s : s \in D\}$. Variables with subscript $_{-s}$ denote all locations in D excluding s . Lowercase letters refer to realizations of the variables. For notational convenience, in this section we assume a single realization of the process. We also apply this model for data with multiple time steps by assuming temporal independence and that the model has the same structure over time.

Without restrictions, the response Y_s is potentially a function of X_D and A_D at all locations, greatly increasing the number of potential outcomes. To make this manageable while still taking spatial interference into account, we assume that the potential outcome $Y_s(a_D)$ depends on exposure field a_D through two mechanisms; a direct exposure, a_s , and an indirect/spill-over exposure, $\tilde{a}_{\tau,s} = \int_{D \setminus s} \omega_\tau(\|s - s'\|) a_{s'} ds'$, where $\omega_\tau(\cdot) : \mathbb{R}^+ \mapsto [0, 1]$ is a kernel function with bandwidth $\tau > 0$. This constitutes a general class of interference structures, although we will focus primarily on a Gaussian kernel form: $\omega_\tau(d) = \exp\{-(d/\tau)^2\}$, which decays smoothly over space. This choice of kernel, which is an assumption on how dependency decays over space, will be relaxed in Section 7. Because only finitely many locations are observed in practice, the integral form of $\tilde{a}_{\tau,s}$ must be approximated with a sum. We approximate $\tilde{a}_{\tau,s}$ with the form $\tilde{a}_{\tau,s} = \sum_{s' \in \{s_1, \dots, s_n\} \setminus s} \omega_\tau(\|s - s'\|) a_{s'}$.

Implicitly, we assume that for any s and exposures a_D and a'_D , $Y_s(a_D) = Y_s(a'_D)$ if $a_s = a'_s$ and $\tilde{a}_s = \tilde{a}'_s$. This simplified exposure allows us to parsimoniously define the individual potential outcomes for all possible exposure fields a_D in terms of only the local direct and spill-over exposures: $Y_s(a_s, \tilde{a}_{\tau,s})$. This is a crucial assumption about the effect of exposure, which substantially reduces the dimension of the problem. In particular, it assumes that any exposure at a given distance from location s (where distance is interpreted by the kernel function) has the same effect on the response. This can be thought of as a spatial analog to the “stratified interference” assumption introduced by Hudgens and Halloran (2008), in which only the aggregated exposures within a cluster contribute to interference rather than the full set of individual exposures.

Identification of the exposure effects follows from the following assumptions:

Assumption 1 (Unconfoundedness). For all a_D , $Y_s(a_D) = Y_s(a_s, \tilde{a}_{\tau,s}) \perp\!\!\!\perp A_D \mid X_D$.

Assumption 2 (Positivity). \forall with $\text{pr}(X_D =) > 0$, $\text{pr}(A_D = a_D \mid X_D =) > 0, \forall a_D$.

Assumption 3 (Consistency). The potential outcome $Y_s(a_s, \tilde{a}_{\tau,s}) = Y_s$ when $A_s = a_s, \tilde{A}_s = \tilde{a}_{\tau,s}$.

For finite D , with only the assumptions above, exposures effects theoretically are identifiable. However, identification requires the number of repeated field observations to be at least 2^n , which is rare. To make the situation tractable, we make one additional assumptions about our data:

Assumption 4 (Marginal Structural Model). The potential outcomes model take the form

$$Y_s(a_s, \tilde{a}_{\tau,s}) = \beta_0 + \delta_1 a_s + \delta_2 \tilde{a}_{\tau,s} + h(X_D)_s + \epsilon_s, \quad (1)$$

where $h(X_D)$ is a general function of X_D , and ϵ_s is an error process that is independent of A_D and X_D .

We will focus exclusively on i.i.d. error, however, dependence (e.g., spatial) can be incorporated into ϵ_s . Here δ_1 and δ_2 quantify the direct and spill-over effects of exposure, respectively. Specifically, δ_1 and δ_2 represent the causal effect of a unit-increase in the direct and spill-over effects:

$$E\{Y_s(1, \tilde{a}_s) - Y_s(0, \tilde{a}_s)\} = \delta_1,$$

$$E\{Y_s(a_s, c + 1) - Y_s(a_s, c)\} = \delta_2.$$

Similarly, τ quantifies the range of the spill-over effect.

Under Assumptions 1–4, the coefficients in (1) are identifiable in the sense that

$$\begin{aligned} E\{Y_s(a_D) \mid X_D\} &= E\{Y_s(a_D) \mid X_D, A_s = a_s, \tilde{A} = \tilde{a}\} \\ &= E(Y_s \mid X_D, A_s = a_s, \tilde{A} = \tilde{a}). \end{aligned} \quad (2)$$

The first equality follows from Assumptions 1 and 2; the second from Assumptions 3 and 4.

It is instructive to consider the dependence that is created by Assumptions 1–4. X_D is unrestricted, and is therefore plausibly spatially correlated. Because the direct exposure mechanism is a function of X_D , A_D will likely reflect any spatial structure in X_D . Y_D may reflect both general spatial dependence from X_D as well as any induced spatial dependence from A_D .

3 | THE GENERALIZED PROPENSITY SCORE IS A BALANCING SCORE

The identification formula (2) implies that we can estimate δ_1, δ_2 , and τ using the regression model

$$Y_s = \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} + h(X_D)_s + \epsilon_s,$$

if $h(\cdot)$ is known and ϵ_s is a mean zero error process. In most cases, though, $h(X_D)$ is not known. The standard causal inference strategy at this point is to condition on X_D itself, if known. However, even when X_D is known, in the context of spatial analysis it is high dimensional. Specifically, for unit s it does not suffice to condition on X_s , but requires conditioning on X at all locations. With both high-dimensional confounders as well as our assumptions about the exposure mechanism, the natural path forward is to condition on the propensity of exposure (Rosenbaum & Rubin, 1983).

In a setting without interference, and thus only direct exposure effects, the standard propensity score e_s for binary exposures is defined as $e_s(X_D) = P(A_s = 1 | X_D)$. This is easily extended to real valued exposures using the form $e_s(X_D) = f(A_s = \eta | X_D), \eta \in \mathbb{R}$. In both cases, e_s simply summarizes the conditional distribution of exposure. The propensity score is an example of a balancing score: a function of the covariates that, once conditioned on, induces independence between the exposure and covariates. If all confounders are included in X , then e_s , rather than X_D , may be conditioned on for unbiased exposure effects. When X_D is high dimensional, as in our motivating example, this is a substantial dimension reduction.

Under interference, with exposure components a_s and $\tilde{a}_{\tau,s}$, the propensity score approach can still be utilized, by defining the propensity of exposure to be a summary of the conditional distribution of $(A_s, \tilde{A}_{\tau,s})$. To this end, we define $g_{\tau,s}$ to be the joint propensity of A_s and $\tilde{A}_{\tau,s}$:

$$g_{\tau,s}(X_D) = f(A_s = \eta, \tilde{A}_{\tau,s} = \nu | X_D), \quad \eta, \nu \in \mathbb{R}. \quad (3)$$

We refer to the bivariate density function $g_{\tau,s}$ as the generalized propensity score (Bia et al., 2014; Hirano and Imbens, 2004; Imbens, 2000; Imai & Van Dyk, 2004; Kluve et al., 2012; Lechner, 2001). Importantly, this form of $g_{\tau,s}$ allows for A_D to be correlated, which may cause dependence between A_s and $\tilde{A}_{\tau,s}$.

The key insight is that $g_{\tau,s}$ is a balancing score. This implies that, paired with our no unmeasured confounders assumption, the observed exposures and potential outcomes are independent conditional on $g_{\tau,s}$. This is the strategy that we use to recover unbiased estimates of our key coefficients δ_1 and δ_2 . Theorem 1 shows this formally, by extending the analogous result for propensity scores for continuous exposures by Hirano and Imbens (2004) to our score $g_{\tau,s}$.

Theorem 1 ($g_{\tau,s}$ is a balancing score). *Given Assumptions 1–4, then for all locations s and spill-over exposure levels ν ,*

$$Y_s(a_D) = Y_s(a_s, \tilde{a}_{\tau,s} = \nu) \perp\!\!\!\perp (A_s, \tilde{A}_{\tau,s}) \mid g_{\tau,s}(X_D).$$

The proof is provided in the [supporting information](#).

By Theorem 1, it suffices to adjust for $g_{\tau,s}$ to remove confounding bias. That is,

$$E\{Y_s(a_D) \mid g_{\tau,s}(X_D)\} = E\{Y_s \mid g_{\tau,s}(X_D), A_s = a_s, \tilde{A}_{\tau,s} = \tilde{a}_{\tau,s}\}.$$

This suggests that we can adjust for confounding by incorporating $g_{\tau,s}$ into the regression model.

4 | MODELING THE GENERALIZED PROPENSITY SCORE

Estimating $g_{\tau,s}$ is difficult. It is a bivariate distribution function over X_D , and non-parametric estimation of even univariate density functions suffers from dimensionality issues. To overcome this, we make the following dimension reduction assumption.

Assumption 5 ($g_{\tau,s}$ is a parametric distribution). $g_{\tau,s}$ is a bivariate parametric density with parameters $\bar{Z}_s = (Z_s^{(1)}, \dots, Z_s^{(K)})$ that are a functions of τ and X_D .

That is, the distribution of $(A_s, \tilde{A}_{\tau,s})$ can be completely summarized by the parameters \bar{Z}_s .

Example 1. If A_D are independent and Gaussian, then $\tilde{A}_{\tau,s}$ is itself Gaussian. Setting Z_s^1, \dots, Z_s^4 to be the mean and variance of both A_s and $\tilde{A}_{\tau,s}$ completely summarizes its distribution.

Corollary 1. *Given Assumptions 1–5, then for all locations s*

$$Y_s(a_D) = Y_s(a_s, \tilde{a}_{\tau,s}) \perp\!\!\!\perp (A_s, \tilde{A}_{\tau,s}) \mid \bar{Z}_s.$$

Corollary 1 follows immediately from Theorem 1 and states that conditioning on \bar{Z}_s is equivalent to conditioning directly on the distribution $g_{\tau,s}$, and so our Theorem 1 result of unconfoundedness given g_{τ} extends to the considerably more tractable situation of unconfoundedness given \bar{Z} .

Identification of δ_1 and δ_2 follows from the conditional independence in Corollary 1. Let * denote true values; variables without * being estimated values. Based on

$$\begin{aligned} E\{Y_s(a_D) \mid \bar{Z}_s\} &= \beta_0^* + \delta_1^* a_s + \delta_2^* \tilde{a}_{\tau,s} + E\{h(X_D)_s \mid \bar{Z}_s\} \\ &= E\{Y_s(a_D) \mid A_D = a_D, \bar{Z}_s\} \\ &= E\{Y_s \mid A_D = a_D, \bar{Z}_s\} \\ &= E\{\beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} \mid A_D = a_D, \bar{Z}_s\}, \end{aligned} \quad (4)$$

we must have $\delta_1 = \delta_1^*$ and $\delta_2 = \delta_2^*$.

In practice, we will condition on the components of $g_{\tau,s}$ using B-splines, by including them directly in the regression:

$$E\left(\beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} \mid A_D = a_D, \bar{Z}_s\right) \approx \beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} + spl\left(\bar{Z}_s\right) \approx \beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} + spl_1\left(Z_s^{(1)}\right) + \dots + spl_q\left(Z_s^{(K)}\right). \tag{5}$$

This allows for completely arbitrary dependence between \bar{Z} and Y . The second line of (5) implicitly assumes that the spline components enter additively, an assumption that can be tested. In the presence of nonadditivity, a tensor product of the components should be used that allows for general interactions, but at great computational cost (Wood, 2006).

5 | BAYESIAN INFERENCE AND COMPUTATIONAL ALGORITHM

The identification results (4) and (5) allow unbiased estimation of δ_1 and δ_2 using a regression of the observed response onto the direct and spill-over exposures as well as the spline estimates of $\bar{Z}_s(\tau)$. Implementing this involves three steps: Step 1 parametrizes and estimates the propensities $g_{\tau,s}$ of direct and spill-over exposure. Step 2 estimates a preliminary posterior for the range parameter τ , which must be done in a separate step for reasons discussed below. Step 3 estimates final posterior distributions for all parameters via Markov chain Monte Carlo sampling. This is illustrated in Figure 1.

The propensities of direct exposure that are tackled in Step 1 are first estimated by regressing A_s onto X . This requires parametrizing the form of $f(A)$, and identifying a correctly specified propensity score. The form of this score can vary in complexity. The simplest case is that of a local exposure assignment mechanism, that is, the distribution of A_s is influenced by X_s only. This would simply entail a regression on local covariates. A moderately complex case would allow for nearby X to inform the propensity of exposure. A very general case would allow A_D to be spatially dependent, conditional on X_D . That is, A_s would depend directly on nearby A .

Estimating the spill-over propensity component of Step 1 is similar. First, a family of parametric distributions must be identified. One intuitive method of doing this is to select several candidate distributions based on the form of A_s , and select among them by simulating values of \tilde{A}_s . For example, if A_s is binary, then the potential candidates for the distribution of \tilde{A} must be nonnegative and allow for point mass at zero. Obvious contenders are zero-inflated

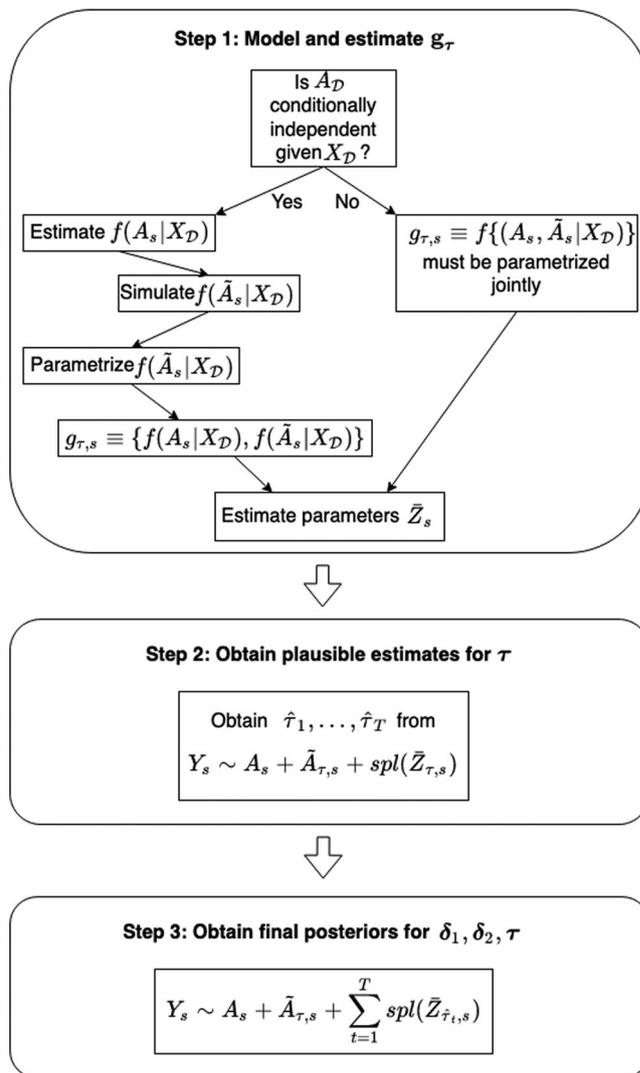


FIGURE 1 Schematic plot for computation of the proposed method. Here $f(\cdot)$ refers to a generic probability density/mass function, and spl refers to a B-spline basis

lognormal and zero-inflated Gamma distributions. A natural way to select between them is to simulate from the estimated propensities of A_s , to get simulated \tilde{A}_s values using different reasonable τ . The empirical distributions of these simulated \tilde{A}_s will often suggest one family of distributions. With a chosen distribution in hand, the parameters \bar{Z} at each location can be estimated directly from the field X_D and τ . Because these parameters will be conditioned on by entering into a splined regression, it is advantageous that their values have reasonable spread. To this end, one-to-one transformations of the parameters such as log and logit are helpful.

Step 2 involves identifying a plausible set of τ values to be used in Step 3. Because $\bar{Z}(\tau)$ represents a propensity score, estimating τ directly in the final model is problematic. It is clear from the definition of a propensity score that the response Y should not provide any information on the

propensity of exposure. However, estimating a response model that includes $\bar{Z}(\tau)$ directly does just that, since Y can influence $\bar{Z}(\tau)$ through τ . This problem is articulated in McCandless et al. (2010), Saarela et al. (2015), Saarela et al. (2016) Zigler et al. (2013), and Zigler (2016), while steps can be taken to mitigate feedback from Y to \bar{Z} issues remain.

Our solution to this issue takes inspiration from the standard two-step propensity score treatment in which propensity scores are first estimated and treated as fixed, and then conditioned on in an outcome model. Because τ is unknown, estimating $\bar{Z}(\tau)$ in advance is impossible. However, estimating the model with feedback shown below does give approximate estimates of τ . From this approximate posterior of τ , a set of reasonable τ values (τ_1, \dots, τ_T) covering the plausible range of τ can be identified. Then $\bar{Z}(\tau_1), \dots, \bar{Z}(\tau_T)$ can be pre-computed and conditioned on simultaneously in the response model in Step 3. Because each of these $\bar{Z}(\tau_t)$ are computed before the response model, the feedback issue is resolved.

Therefore, in Step 2 we estimate

$$Y_s = \beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} + \sum_{k=1}^K \text{spl}_{tk} \left(Z_s^{(k)}(\tau) \right) + \epsilon_s, \quad (6)$$

where ϵ_s is distributed independent Normal($0, \sigma^2$). An attempt to cut the feedback from Y to \bar{Z} is made by estimating τ in the Metropolis step using only \tilde{A}_{τ} while holding $\bar{Z}(\tau)$ fixed. A recommended plausible set for τ might then be $\{\hat{\tau}, \hat{\tau} \pm 2s, \hat{\tau} \pm 4s\}$, where $\hat{\tau}$ and s are the posterior mean and standard deviation of τ in (6).

Finally in Step 3 each fixed $\bar{Z}(\tau_t)$ enters the final model as

$$Y_s = \beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_{\tau,s} + \sum_{t=1}^T \sum_{k=1}^K \text{spl}_{tk} \left(Z_s^{(k)}(\tau_t) \right) + \epsilon_s. \quad (7)$$

This model produces accurate posteriors on all variables. Although each τ_t is fixed within the \bar{Z} terms, τ can still vary within $\tilde{A}_{\tau,s}$. For the spline terms in (6)–(7), we use B-spline expansions taken at fixed intervals over the variables' range of values (Eilers & Marx, 1996; Ngo & Wand, 2004). Before running the regression, we remove any redundant $Z_s^{(k)}$ that are collinear. Specifically, the spline columns that the R “lm” function selects to remove due to collinearity when running the analogous regression, we remove as well (R Core Team, 2018). All regression coefficients are estimated using Gibbs sampling; τ , which now enters only through \tilde{A}_{τ} , uses a Metropolis step. If Assumptions 1–5 hold, we recover unbiased estimate of the exposure

effects. Comparing the forms of the assumed true model (1) and the estimated model (7) shows that we have essentially replaced the unknown $h(X_D)$ with flexible functions of \bar{Z} .

6 | SIMULATION STUDY

We examine the performance of this method using simulated data, which take inspiration from the wildfire/air pollution data in Section 7. Since we use a binary exposure in Section 7 to indicate the presence of a fire, we use $A_s \in \{0, 1\}$ here. In addition, we assume A_s at different locations is independent conditional on local X_s . This precludes the more complex cases of independence conditional on X_D or conditional dependence. Doing this allows for more straightforward modeling of $g_{\tau,s}$, as shown in 6.1.

We generate the data as follows. Fields X_D , A_D , and Y_D are generated on $n^{1/2} \times n^{1/2}$ grids, with $n = 25, 100$ on the unit square $[0, 1] \times [0, 1]$. We generate $N = 100$ independent repeated observations of the fields for each dataset. Thus each complete dataset involves $n \times N$ different data points. The single covariate $X_s \in \mathbb{R}^1$ is a mean zero, variance one, Gaussian process and with isotropic exponential covariance and spatial range 0.6. The binary direct exposure A_s is determined locally and distributed independently Bernoulli $\{\text{expit}(X_s - 3)\}$. The spill-over exposure takes the form $\tilde{A}_{\tau,s} = \sum_{s'} \omega_{\tau}(\|s - s'\|) A_{s'}$, with ω_{τ} a Gaussian kernel with $\tau = 0.3$. Several confounders $h(X_D)_s$ are investigated: a weighted average W_s is taken of the X_D values using a Gaussian kernel with $\tau = 0.5$ and weights normalized to sum to 1. Simulations are run with $h(X_D)$ set to W_s , $-(W_s)^3$, and $\exp(W_s)$. Additionally, a final $h(X_D)$ is included, which consists of two small Gaussian kernels ($\tau = 0.15$), that are centered two grid cells diagonally to the northeast and southwest of s , respectively: $\tilde{A}_s = \sum_{s' \in \{s_1, \dots, s_n\}} \{\omega_{\tau}(\|s - s' + u\|) a_{s'} + \omega_{\tau}(\|s - s' - u\|) a_{s'}\}$, where $u = (1/9, 1/9)$ is the distance between adjacent points and ω_{τ} is a Gaussian kernel. This is included to assess sensitivity to anomalous $h(X_D)$.

Lastly, Y_s follows the form of (1), with $\beta_0 = 0$, $\delta_1 = \delta_2 = 1$, and ϵ_s independently distributed standard normal. Each setting is repeated 500 times.

6.1 | Estimation

Following the three steps outlined in Section 5, we first parametrize and estimate $g_{\tau,s}$. Because A_s is assumed to be conditionally independent given X_s , we can estimate Z

components for the distributions of A and \tilde{A} separately. A_s is binary, so we assume it has a Bernoulli distribution with the correctly specified propensity in which $\text{logit}\{E(A_s)\}$ is affine in X_s . Its distribution is then captured with the standard propensity score $Z_s^{(1)} = \text{pr}(A_s = 1 | X_s)$. Estimates are obtained with a simple logistic regression from A_s onto X_s , with $Z_s^{(1)}$ set to the log of the fitted values.

We then choose a parametric form for the distribution of $\tilde{A}_{\tau,s}$. From our estimated $Z_s^{(1)}$, we use different plausible τ values to generate simulated A , which we then use to get an empirical distribution of simulated \tilde{A} . Examination of these distributions leads us to choose a zero-inflated lognormal distribution for \tilde{A}_s :

$$\text{pr}(\tilde{A} = 0 | X_s) = p_0, \quad \text{pr}(\tilde{A} = v | \tilde{A} > 0, X_s) = \frac{1}{v\sigma\sqrt{2\pi}} \exp\left\{-\frac{(\log v - \mu)^2}{2\sigma^2}\right\}.$$

Rather than use the three parameters p_0 , μ , and σ^2 for our $Z_s^{(2)}$, $Z_s^{(3)}$, and $Z_s^{(4)}$, we choose three more stable one-to-one transformations: $\text{logit}(p_0)$, $\text{log}\{E(\tilde{A})\}$, and $\text{log}\{\text{Var}(\tilde{A})\}$.

In place of Step 2 the (τ_1, \dots, τ_T) values used are $\{0.25, 0.35, 0.45, 0.55\}$, which surround but do not contain the true $\tau = 0.3$. Rather than re-estimate these values with each simulation repetition, we use this set to ensure comparability across repetitions. Finally, Step 3 uses Markov chain Monte Carlo to estimate all variables in (7). Further details are provided in the Supporting Information.

In addition to the proposed generalized propensity score model, we estimate three comparison models: (i) the oracle model, $E(Y_s) = A_s + \tilde{A}_{\tau,s} + h(X_D)_s$, is the true model that includes otherwise unknown $h(X_D)$ as a covariate, (ii) the local only model, $E(Y_s) = A_s + \tilde{A}_{\tau,s} + \sum_j \text{spl}_j(X_s^j)$, conditions on local covariates using splines, and (iii) the naive model, $E(Y_s) = A_s + \tilde{A}_{\tau,s}$, simply regresses the outcome onto the exposures, but does not incorporate any causal conditioning.

6.2 | Simulation results

Tables 1 and 2 show the simulation bias and coverage for the 10×10 grids. The Naive model does very poorly in all scenarios, indicating substantial confounding between A and Y . The generalized propensity score model performs substantially better than both the Local only and the Naive models, although, intuitively, the Local only model does show reasonable direct effect estimates. In most cases, the generalized propensity score model performs comparably to the Oracle model. Results for the 5×5 grids are similar. A sensitivity analysis that examines robustness to incorrectly specified error and poorly fitting propensity score models is given in Web Appendix D.

7 | ESTIMATING THE CAUSAL EFFECT OF WILDLAND FIRES ON AIR POLLUTION

Wildland fires release harmful particles and gasses impacting air quality near the fire and downwind (Larsen et al., 2018). Fine particulate matter smaller than $2.5 \mu\text{m}$ ($\text{PM}_{2.5}$) have been linked to adverse cardiorespiratory health outcomes (Brook, 2007; Corrigan et al., 2018; Dominici et al., 2006; Rappold et al., 2012; Weber et al., 2016). For these reasons, understanding the causal effect of wildland fires on air pollution across space is of significant interest. Specifically, we are interested in the time-averaged causal effect of wildfires on ambient $\text{PM}_{2.5}$ concentrations across Western United States from 2005 to 2018.

7.1 | Data

The response Y consists of 24-hour average $\text{PM}_{2.5}$ concentrations measured in $\mu\text{g}/\text{m}^3$ at 416 measurement sites, some of which are plotted in Figure 2. Observations are collected every one, three or six days depending on the station. The data are publicly available and provided by the Environmental Protection Agency (Environmental Protection Agency, 2019). For each location, the long-term mean is subtracted.

The dates and locations of fires are compiled from a mix of satellite data and incident reports reported to the recently closed Geospatial Multi-Agency Coordination (GeoMAC) Wildland Fire Support program (Geospatial Multi-Agency Coordination, 2019). Moving forward these datasets will be available through the National Interagency Fire Center (National Interagency Fire Center, 2020). Because the focus of our analysis is on $\text{PM}_{2.5}$ only fires larger than 1000 acres are included in the analysis. Among the 3,930 fires, 34.8% of fires are missing either a start or end date. For these fires, we impute missing values by modeling fire duration as a linear function of $\text{log}(\text{area burned})$.

Finally, 11 confounders X^1, \dots, X^{11} are included in the exposure balancing score. These include the four components of the National Fire Danger Rating System (energy release component, burning index, ignition component, and spread index) that are used to monitor daily risk of fire in the United States. The other variables used in the balancing score include elevation, daily temperature, relative humidity, wind speed, precipitation level, the Keetch–Byram drought index, and the numeric day of the year. These variables are compiled and made available through the Wildland Fire Assessment System (Wildland Fire Assessment System, 2019). A snapshot of the exposure,

TABLE 1 Simulation bias for 10×10 grids multiplied by 1000, with standard errors

$h(X_D)_s$	Model	δ_1	δ_2	τ
W_s	Oracle	0.2 (1.8)	-0.7 (0.9)	0.1 (0.2)
	Generalized propensity score	1.4 (1.9)	0.5 (1)	0 (0.2)
	Local Only	2.6 (2)	-72.1 (1.1)	69.1 (0.4)
	Naive	236.7 (1.9)	52.7 (1.5)	79.2 (0.5)
$-(W_s)^3$	Oracle	0.1 (1.8)	-0.7 (0.9)	0.1 (0.2)
	Generalized propensity score	1.1 (1.9)	-0.1 (1)	0.1 (0.2)
	Local Only	1 (2.1)	28.5 (1.3)	-39.7 (0.4)
	Naive	-205.9 (2.8)	-104.1 (1.6)	-53.1 (0.6)
$\exp(W_s)$	Oracle	0.3 (1.8)	-0.6 (0.9)	0.1 (0.2)
	Generalized propensity score	1.5 (1.9)	0.9 (1.1)	0.1 (0.3)
	Local Only	2.8 (2.3)	-101.4 (2.1)	120.7 (1.4)
	Naive	381.3 (2.8)	105.3 (2.3)	114.2 (1.3)
2 circles	Oracle	-0.9 (1.8)	0.5 (0.9)	0.1 (0.2)
	Generalized propensity score	-0.9 (2)	-0.1 (1.2)	0.5 (0.3)
	Local Only	0.1 (2.1)	-39 (1.2)	62.5 (0.4)
	Naive	244.1 (2)	94.1 (1.5)	72.4 (0.5)

TABLE 2 Simulation coverage for 10×10 grids, with standard errors

$h(X_D)_s$	Model	δ_1	δ_2	τ
W_s	Oracle	95 (1)	94.6 (1)	93.6 (1.1)
	Generalized propensity score	93.8 (1.1)	94.4 (1)	93.6 (1.1)
	Local Only	93.2 (1.1)	8.8 (1.3)	0 (0)
	Naive	0 (0)	27.6 (2)	0 (0)
$-(W_s)^3$	Oracle	95.2 (1)	94.6 (1)	93.8 (1.1)
	Generalized propensity score	95.2 (1)	93.8 (1.1)	92.8 (1.2)
	Local Only	93.8 (1.1)	77.2 (1.9)	0 (0)
	Naive	2 (0.6)	6 (1.1)	0 (0)
$\exp(W_s)$	Oracle	95.4 (0.9)	94.2 (1)	93.8 (1.1)
	Generalized propensity score	93 (1.1)	90.2 (1.3)	90.2 (1.3)
	Local Only	90.8 (1.3)	4.6 (0.9)	0 (0)
	Naive	0 (0)	8.2 (1.2)	0 (0)
2 circles	Oracle	95.6 (0.9)	95 (1)	93.4 (1.1)
	Generalized propensity score	94.6 (1)	89.8 (1.4)	90.4 (1.3)
	Local Only	95.2 (1)	47.4 (2.2)	0 (0)
	Naive	0.2 (0.2)	4 (0.9)	0 (0)

response, energy release component, ignition component, Keetch-Byram drought index, and relative humidity are shown in Figure 2.

Our analysis treats each daily air observation as the center of a 9×9 grid, with a height and width of 9 degrees latitude/longitude. For each such grid, only the center grid cell has a response Y_s value. However, all 81 grid cells have covariates X_s^j and direct exposure A_s values. Each grid cell receives direct exposure $A_s = 1$ if there was at least one fire in the cell on that particular day; 0 otherwise. Each X_s^j value is taken to be the mean of the observed

covariates in each cell/day combination. For cell/days with no observed values, a value is imputed from nearby cells using a kernel smoother as implemented in the “fields” R package (Nychka et al., 2014). The end result is 592,274 observed grids, each of which contains 9×9 grids for A_D and X_D^j , $j = 1, \dots, 11$, as well as a centered Y_s value. Finally, any grid cells whose centers extend outside of the Western United States are disregarded and excluded from analysis. In this context, the direct effect of exposure consists of the causal effect on Y_s from a fire in the same grid cell ($A_s = 1$), whereas the indirect effect consists of the causal effect on

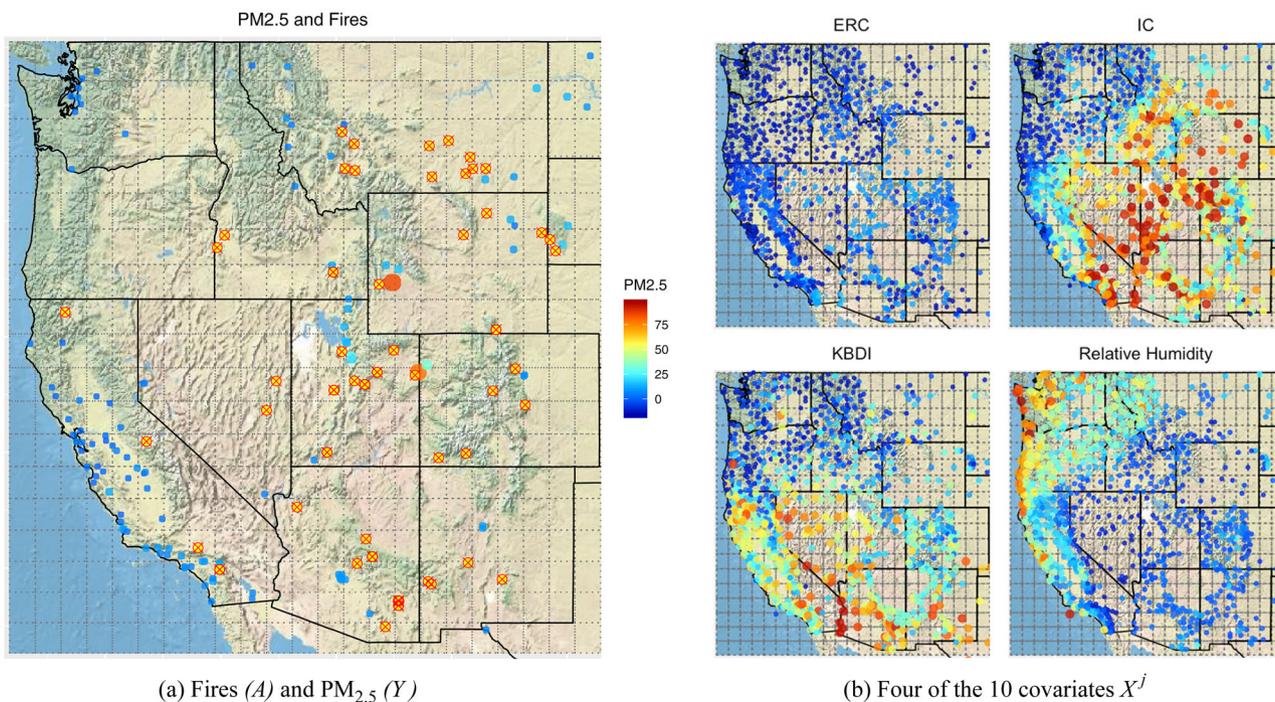


FIGURE 2 Data snapshot on July 1, 2012. Energy Release Component (ERC) and Ignition Component (IC) are two of National Fire Danger Rating System Components; KBDI refers to the Keetch-0Byram drought index. (a) Fires are shown as cross-hatched circles and $PM_{2.5}$ locations are shown as solid circles. This figure appears in color in the electronic version of this article, and any mention of color refers to that version

Y_s from $A_{s'}$ in other cells ($s \neq s'$). As in the simulation study, each of these grids is treated as independent. In addition to the generalized propensity score model, we estimate a model that conditions on the local covariates only, using splines. Though pre-processing and imputing the covariates $X^{(j)}$ is inescapable given shape of the data, we do not, strictly speaking, account for this imputation (and likely smoothing) in our analysis. This potentially reduces the accuracy of our propensity score, and thereby increases the uncertainty in our final estimates. Our standard errors also do not reflect this potential uncertainty.

We use the same form of $g_{\tau,s}$ as given in Section 6. A_s at different locations are assumed to be conditionally independent given X_s , which allows us to estimate separate components for A_s and $\tilde{A}_{\tau,s}$. The propensity component $\log\{E(A_s)\}$ is estimated as a linear model of 5-element B-splines of X_s^1, \dots, X_s^{11} , and the propensity of \tilde{A} is assumed to be zero-inflated lognormal. Conditioning on local X_s only is justified because we posit that local X_s contains the vast majority of information about the propensity of fire, with distant locations giving minimal information. Finally, we note that though the data are longitudinal, we assume i.i.d. error. Beyond including the numeric day of the year as a covariate, each 9×9 grid is considered independent over time.

Before proceeding, we briefly review the key assumptions that our method makes, and why they are reasonable

for this application. We assume all relevant confounders are stored in X . To ensure this is approximately accurate, we have included as many potential confounders as possible. Our analysis assumes the errors are independent over time and space. To satisfy this assumption, we have removed station-level means and included day of the year effects and several relevant covariates to the mean model. Finally, we make the assumption that the exposures A are independent given X_s . This is reasonable because the chance of a fire locally is determined by the conditions locally, and not within neighboring cells. The cells used are large enough that the spread of fire from cell to cell will be rare.

7.2 | Results

Table 3 shows the results. The causal direct effect estimate given by the generalized propensity score model is $1.03 \mu\text{g}/\text{m}^3$ of $PM_{2.5}$, or 11.9% of the annual mean $PM_{2.5}$ observed throughout. The range parameter τ is estimated to be 1.53 degrees of latitude/longitude, suggesting that fires impact up to roughly 3 degrees away. The estimate of 0.13 for δ_2 represents the height of spill-over kernel at its peak. All of δ_1 , δ_2 , and τ are highly significant. The estimated direct effect from the local-only model is 12% larger than the estimate from the generalized propensity score

TABLE 3 Posterior mean (95% credible interval)

	Direct Effect (δ_1)	Spill-over Effect (δ_2)	Bandwidth (τ)
Local Only	1.15 (1.05, 1.25)	0.11 (0.09, 0.12)	17.37 (8.28, 42.12)
Generalized propensity score	1.03 (0.93, 1.14)	0.13 (0.03, 0.25)	1.53 (1.17, 2.88)

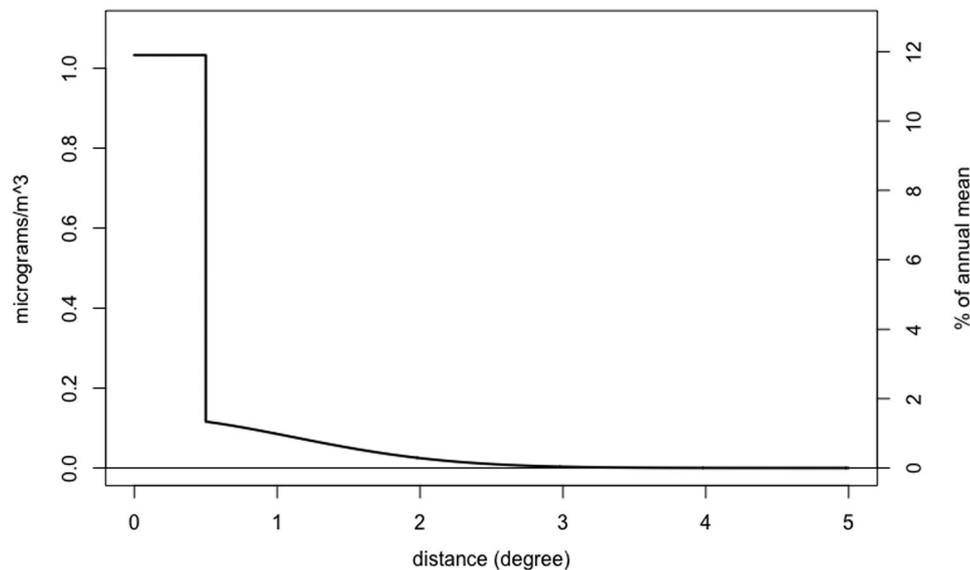


FIGURE 3 Causal effect of a fire on $PM_{2.5}$ by distance, as measured in degrees of latitude and longitude. The left axis shows the raw causal increase in $PM_{2.5}$; the right axis shows this as a percentage of annual mean $PM_{2.5}$ levels

model, and the local-only model has an implausibly large and imprecise estimate of τ .

Figure 3 illustrates the implied causal effect of fire at different distances from the generalized propensity score model. Taking the center of a grid cell as our vantage point, the direct effect of one or more fires in the same grid cell has a time-averaged causal increase of $1.03 \mu\text{g}/\text{m}^3$ of $PM_{2.5}$, which corresponds to the step from 0 to 0.5 in the east/west or north/south direction; slightly more than 0.5 when at an angle. As the fire gets progressively further away, the causal effect decays smoothly until it approaches 0 roughly 3 grid cells away. Intuitively this kernel extending out from 0 is completely determined by τ and δ_2 : τ corresponds to the width of the kernel; δ_2 is the height of the kernel at its peak. Because a significant amount of the fire exposure data was imputed, the analysis was repeated without including fires that had missing start or end dates. This was done by setting $A_s = 0$ for these fires. This analysis gave broadly similar results: $\hat{\delta}_1 = 1.18$ (1.05, 1.31), $\hat{\delta}_2 = 0.19$ (0.07, 0.34), $\hat{\tau} = 1.26$ (1.08, 1.56). The fact that the estimates for δ_1 and δ_2 are higher here may indicate that some fire-days in the original analysis may have been erroneously imputed, diluting the estimated exposure effects.

The wildfire analysis makes several simplifications that are important to consider. First, treating A_s as binary sac-

rifices information on the number and size of fires in a given grid cell. Extending this method to incorporate information on the size of the fire would preserve information. Moreover, we assume τ , δ_1 , and δ_2 are fixed, although it is possible that they naturally vary across different fires and locations. However, there is not enough information in the data to identify these differences. Additionally, we do not consider time-varying effects, as we focused on the contribution to time-averaged $PM_{2.5}$ levels. Another important simplification is the treatment of separate days as independent. There are temporal trends in the exposure, response, and covariates, and our assumption of independence may inflate the amount of information that our data appear to have. To test the sensitivity to the form of the kernel function in the definition of the indirect exposure, we refit the wildfire data with two additional kernels. However, the direct and spill-over effect estimates are similar for all kernels. These alternative results are provided in Web Appendix D.

8 | DISCUSSION

The generalized propensity score method presented here establishes a new framework to recover causal direct and spill-over effects in the presence of spatial interference, as

well as estimate the range of the interference. The inherent dimensionality issues of the problem are dealt with via a novel propensity score approach, which uses a Bayesian spline-based regression model and a dimension reduction approximation to make the problem tractable. However, there are several critical yet strong assumptions that must hold for our method to perform well. The method hinges on a correctly specified propensity score $g_{\tau,s}(X_D)$ as well as a correctly specified potential outcomes model in (1). This includes accommodating conditionally dependent A_D , and correctly characterizing the spatial dependence on A_s from nearby X . Moreover, the no unmeasured confounders assumption is always a strong, but necessary, assumption for causal inference on observational data. In practice, considerable effort should be made to include any potential confounders for this reason. We rely crucially on the assumption that the distribution of exposures (A_s, \tilde{A}_s) can be encapsulated with the parameters \bar{Z}_s of the propensity score $g_{\tau,s}$. This will rarely be completely accurate in practice, so effort should be made to select an appropriate parametric form for $g_{\tau,s}$. We assume that the exposure effects and range are constant across space and time, due to the difficulty of generalizing this assumption. However, allowing for spatially varying exposure effects and particularly range are a promising direction for future research.

DATA AVAILABILITY STATEMENT

The data that support the findings in this paper are openly available in Harvard Dataverse at <https://doi.org/10.7910/DVN/11GSTG> (Giffin, 2021).

ORCID

A. Giffin  <https://orcid.org/0000-0001-6487-7766>

B. J. Reich  <https://orcid.org/0000-0002-5473-120X>

S. Yang  <https://orcid.org/0000-0001-7703-707X>

REFERENCES

- Aronow, P.M. & Samii, C. (2017) Estimating average causal effects under general interference, with application to a social network experiment. *The Annals of Applied Statistics*, 11(4), 1912–1947.
- Athey, S., Eckles, D. & Imbens, G.W. (2018) Exact p-values for network interference. *Journal of the American Statistical Association*, 113(521), 230–240.
- Bia, M., Flores, C.A., Flores-Lagunes, A. & Mattei, A. (2014) A stata package for the application of semiparametric estimators of dose-response functions. *The Stata Journal*, 14(3), 580–604.
- Bind, M.-A. (2019) Causal modeling in environmental health. *Annual Review of Public Health*, 40, 23–43.
- Brook, R.D. (2007) Is air pollution a cause of cardiovascular disease? Updated review and controversies. *Reviews on Environmental Health*, 22(2), 115–138.
- Corrigan, A.E., Becker, M.M., Neas, L.M., Cascio, W.E. & Rappold, A.G. (2018) Fine particulate matters: the impact of air quality standards on cardiovascular mortality. *Environmental Research*, 161, 364–369.
- Cox, D.R. (1958) *Planning of experiments*. Hoboken, NJ: Wiley.
- Dominici, F., Peng, R.D., Bell, M.L., Pham, L., McDermott, A., Zeger, S.L., & Samet, J.M. (2006) Fine particulate air pollution and hospital admission for cardiovascular and respiratory diseases. *JAMA*, 295(10), 1127–1134.
- Eilers, P.H. & Marx, B.D. (1996) Flexible smoothing with b-splines and penalties. *Statistical Science*, 89–102.
- Environmental Protection Agency (2019) Air Quality Index, Pre-Generated Data Files: Daily Summary Data – Particulates: PM2.5 FRM/FEM Mass (88101). https://aqs.epa.gov/aqsweb/airdata/download_files.html [Accessed 2019].
- Forastiere, L., Airolidi, E.M. & Mealli, F. (2020) Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116, 1–18.
- Forastiere, L., Mealli, F., Wu, A. & Airolidi, E. (2018) Estimating causal effects under interference using Bayesian generalized propensity scores. Preprint, arXiv:1807.11038.
- Geospatial Multi-Agency Coordination (2019) Historic Fire Data. https://rmgsc.cr.usgs.gov/outgoing/GeoMAC/historic_fire_data [Accessed 2019].
- Giffin, A. (2021) Wildfire data for generalized propensity score approach to causal inference with spatial interference. *Biometrics*. <https://doi.org/10.7910/DVN/11GSTG>.
- Halloran, M.E. & Struchiner, C.J. (1991) Study designs for dependent happenings. *Epidemiology*, 331–338.
- Halloran, M.E. & Struchiner, C.J. (1995) Causal inference in infectious diseases. *Epidemiology*, 6(2), 142–151.
- Hirano, K. & Imbens, G.W. (2004) The propensity score with continuous treatments. *Applied Bayesian Modeling and Causal Inference from Incomplete-Data Perspectives*, 226164, 73–84.
- Hudgens, M.G. & Halloran, M.E. (2008) Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482), 832–842.
- Imai, K. & Van Dyk, D.A. (2004) Causal inference with general treatment regimes: generalizing the propensity score. *Journal of the American Statistical Association*, 99(467), 854–866.
- Imbens, G.W. (2000) The role of the propensity score in estimating dose-response functions. *Biometrika*, 87(3), 706–710.
- Karwa, V. & Airolidi, E.M. (2018) A systematic investigation of classical causal inference strategies under mis-specification due to network interference. Preprint, arXiv:1810.08259.
- Kluve, J., Schneider, H., Uhlendorff, A. & Zhao, Z. (2012) Evaluating continuous training programmes by using the generalized propensity score. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 175(2), 587–617.
- Larsen, A.E., Reich, B.J., Ruminski, M. & Rappold, A.G. (2018) Impacts of fire smoke plumes on regional air quality, 2006–2013. *Journal of Exposure Science & Environmental Epidemiology*, 28(4), 319.
- Lechner, M. (2001) Identification and estimation of causal effects of multiple treatments under the conditional independence assumption. In: Lechner, M., & Pfeiffer, F. (eds.) *Econometric evaluation of labour market policies*. Berlin: Springer, pp. 43–58.
- Liu, L. & Hudgens, M.G. (2014) Large sample randomization inference of causal effects in the presence of interference. *Journal of the American Statistical Association*, 109(505), 288–301.

- Manski, C.F. (2013) Identification of treatment response with social interactions. *The Econometrics Journal*, 16(1), S1–S23.
- McCandless, L.C., Douglas, I.J., Evans, S.J. & Smeeth, L. (2010) Cutting feedback in Bayesian regression adjustment for the propensity score. *The International Journal of Biostatistics*, 6(2).
- National Interagency Fire Center (2020) Wildfire Open Data – Historical GeoMAC fire perimeter datasets. <https://data-nifc.opendata.arcgis.com/search?q=geomac> [Accessed 2020].
- Ngo, L. & Wand, M.P. (2004) Smoothing with mixed model software. *Journal of Statistical Software*, 9. <https://doi.org/10.18637/jss.v009.i01>
- Nychka, D., Furrer, R. & Sain, S. (2014) fields: tools for spatial data. r package version 7.1.
- Papadogeorgou, G., Choirat, C. & Zigler, C.M. (2018) Adjusting for unmeasured spatial confounding with distance adjusted propensity score matching. *Biostatistics*, 20(2), 256–272.
- Papadogeorgou, G. & Dominici, F. (2020) A causal exposure response function with local adjustment for confounding: estimating health effects of exposure to low levels of ambient fine particulate matter. *Annals of Applied Statistics*, 14(2), 850–871.
- Papadogeorgou, G., Mealli, F. & Zigler, C.M. (2019) Causal inference with interfering units for cluster and population level treatment allocation programs. *Biometrics*, 75(3), 778–787.
- Perez-Heydrich, C., Hudgens, M.G., Halloran, M.E., Clemens, J.D., Ali, M., & Emch, M.E. (2014) Assessing effects of cholera vaccination in the presence of interference. *Biometrics*, 70(3), 731–741.
- R Core Team (2018) *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rappold, A.G., Cascio, W.E., Kilaru, V.J., Stone, S.L., Neas, L.M., Devlin, R.B. & Diaz-Sanchez, D. (2012) Cardio-respiratory outcomes associated with exposure to wildfire smoke are modified by measures of community health. *Environmental Health*, 11(1), 71.
- Rosenbaum, P.R. & Rubin, D.B. (1983) The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1), 41–55.
- Rubin, D.B. (1974) Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5), 688.
- Rubin, D.B. (1980) Randomization analysis of experimental data: the Fisher randomization test comment. *Journal of the American Statistical Association*, 75(371), 591–593.
- Saarela, O., Belzile, L.R. & Stephens, D.A. (2016) A Bayesian view of doubly robust causal inference. *Biometrika*, 103(3), 667–681.
- Saarela, O., Stephens, D.A., Moodie, E.E. & Klein, M.B. (2015) On Bayesian estimation of marginal structural models. *Biometrics*, 71(2), 279–288.
- Schnell, P.M. & Papadogeorgou, G. (2020) Mitigating unobserved spatial confounding when estimating the effect of supermarket access on cardiovascular disease deaths. *The Annals of Applied Statistics*, 14(4), 2069–2095.
- Sobel, M.E. (2006) What do randomized studies of housing mobility demonstrate? Causal inference in the face of interference. *Journal of the American Statistical Association*, 101(476), 1398–1407.
- Tchetgen, E.J.T. & VanderWeele, T.J. (2012) On causal inference in the presence of interference. *Statistical Methods in Medical Research*, 21(1), 55–75.
- Verbitsky-Savitz, N. & Raudenbush, S.W. (2012) Causal inference under interference in spatial settings: a case study evaluating community policing program in Chicago. *Epidemiologic Methods*, 1(1), 107–130.
- Weber, S.A., Insaf, T.Z., Hall, E.S., Talbot, T.O. & Huff, A.K. (2016) Assessing the impact of fine particulate matter (PM_{2.5}) on respiratory-cardiovascular chronic diseases in the New York City Metropolitan area using Hierarchical Bayesian Model estimates. *Environmental Research*, 151, 399–409.
- Wildland Fire Assessment System (2019) WFAS Archive. www.fs.fed.us/land/wfas/archive [Accessed 2019].
- Wood, S.N. (2006) Low-rank scale-invariant tensor product smooths for generalized additive mixed models. *Biometrics*, 62(4), 1025–1036.
- Zigler, C.M. (2016) The central role of Bayes' theorem for joint estimation of causal effects and propensity scores. *The American Statistician*, 70(1), 47–54.
- Zigler, C.M., Dominici, F. & Wang, Y. (2012) Estimating causal effects of air quality regulations using principal stratification for spatially correlated multivariate intermediate outcomes. *Biostatistics*, 13(2), 289–302.
- Zigler, C.M. & Papadogeorgou, G. (2021) Bipartite causal inference with interference. *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 36(1), 109.
- Zigler, C.M., Watts, K., Yeh, R.W., Wang, Y., Coull, B.A. & Dominici, F. (2013) Model feedback in Bayesian propensity score estimation. *Biometrics*, 69(1), 263–273.

SUPPORTING INFORMATION

Web Appendices and Tables referenced in Sections 3, 6, and 7, as well as the R code, are available with this paper at the Biometrics website on Wiley Online Library.

Table 3: Simulation bias (standard errors) under alternative propensity score forms and correctly/incorrectly specified response form.

Table 4: Simulation coverage (standard errors) under alternative propensity score forms and correctly/incorrectly specified response form.

Table 5: Posterior mean (95% interval) exposure effects for generalized propensity score model, with alternative kernels.

Data S1

How to cite this article: Giffin, A., Reich, B.J., Yang, S., & Rappold, A.G. (2022) Generalized propensity score approach to causal inference with spatial interference. *Biometrics*, 1–12.

<https://doi.org/10.1111/biom.13745>

Supporting Information for Generalized propensity score approach to causal inference with spatial interference by A. Giffin, B. J. Reich, S. Yang, and A. G. Rappold

Web Appendix A

Proof of Theorem 1:

Claim 1: g_τ is a balancing score.

Proof 1 *By the definition of a propensity score, $g_{\tau,s}(X_{\mathcal{D}})$ has the property that $\text{pr}\{(A_s = \eta, \tilde{A}_s = \nu) \mid X_{\mathcal{D}}, g_{\tau,s}\} = \text{pr}\{(A_s = \eta, \tilde{A}_s = \nu) \mid X_{\mathcal{D}}\}$ which implies $X_{\mathcal{D}} \perp\!\!\!\perp (A_s, \tilde{A}_s) \mid g_{\tau,s}$. And thus $g_{\tau,s}$ is a balancing score for our covariates $X_{\mathcal{D}}$. As noted by Hirano & Imbens (2004) this balancing is a characteristic of $g_{\tau,s}$, and does not rely on any unconfoundedness in the response yet.*

Claim 2: for all levels ν ,

$$\text{pr} \left[A_s = \eta, \tilde{A}_{\tau,s} = \nu \mid Y_s(a_s = \eta, \tilde{a}_s = \nu), g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \right] = \text{pr}\{A_s = \eta, \tilde{A}_s = \nu \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}})\}.$$

(To ease notation, now let $\bar{A}_s = (A_s, \tilde{A}_{\tau,s})$ and $\bar{a}_s = (a_s, \tilde{a}_{\tau,s})$.)

Proof 2 *We can then write*

$$\begin{aligned}
\text{pr} \{ \bar{A}_s = (\eta, \nu) \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} &= f_{\bar{A}_s} \{ \eta, \nu \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} \\
&= \int f_{\bar{A}_s} \{ \eta, \nu \mid X_{\mathcal{D}}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} \, dF_{X_{\mathcal{D}}} \{ X_{\mathcal{D}} \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} \\
&= \int f_{\bar{A}_s}(\eta, \nu \mid X_{\mathcal{D}}) \, dF_{X_{\mathcal{D}}} \{ X_{\mathcal{D}} \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} \\
&= \int g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \, dF_{X_{\mathcal{D}}} \{ X_{\mathcal{D}} \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \} \\
&= g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}),
\end{aligned}$$

$$\begin{aligned}
\text{pr} [\bar{A}_s = (\eta, \nu) \mid Y_s \{ \bar{a}_s = (\eta, \nu) \}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}})] \\
&= f_{\bar{A}_s} [\nu \mid g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}), Y_s \{ \bar{a}_s = (\eta, \nu) \}] \\
&= \int f_{\bar{A}_s} [\eta, \nu \mid X_{\mathcal{D}}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}), Y_s \{ \bar{a}_s = (\eta, \nu) \}] \, dF_{X_{\mathcal{D}}} [X_{\mathcal{D}} \mid Y_s \{ \bar{a}_s = (\eta, \nu) \}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}})] \\
&= \int f_{\bar{A}_s}(\eta, \nu \mid X_{\mathcal{D}}) \, dF_{X_{\mathcal{D}}} [X_{\mathcal{D}} \mid Y_s \{ \bar{a}_s = (\eta, \nu) \}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}})] \\
&= \int g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}) \, dF_{X_{\mathcal{D}}} [X_{\mathcal{D}} \mid Y_s \{ \bar{a}_s = (\eta, \nu) \}, g_{\tau,s}(\eta, \nu, X_{\mathcal{D}})] \\
&= g_{\tau,s}(\eta, \nu, X_{\mathcal{D}}).
\end{aligned}$$

Combining these gives Claim 2, which then implies our result.

Web Appendix B

Bayesian estimation details for simulation:

Uninformative priors are used for all parameters except τ which receives a mildly informative prior. Markov chain Monte Carlo iterations begin at maximum likelihood values for all parameters except τ , which requires an initial estimate. A burn-in length of 7,500 iterations is used, after which we sample 22,500 iterations. Gibbs sampling is used for all parameters except τ , which we transform and sample using Metropolis sampling, with an adaptive tuning scheme during the

burn-in. Specifically, we use a normal proposal distribution for $\log(\tau - \frac{1}{d})$, where d is the number of grid cells along each axis. This prevents the τ samples from becoming pathologically small, in which case the kernel cannot reach the neighboring cells and δ_2 becomes arbitrary large. The comparison models are estimated with similar parameter settings.

For convenience, define β as the vector of β_0 , δ_1 , δ_2 , and the spline coefficients; let $\mu_s = \beta_0 + \delta_1 A_s + \delta_2 \tilde{A}_s + \sum_{j=1}^J b_j^{(0)} B_j^{(0)}(e_s) + \sum_{t=1}^T \sum_{k=1}^q \sum_{j=1}^J b_{j,t}^{(k)} B_{j,t}^{(k)}\{Z_s^{(k)}(\tau_t)\}$; let M be the matrix with columns $1_{(nN)}$, $A_{\mathcal{D}}^{vec}$, $\tilde{A}(\tau)^{vec}$, and the B-splines bases; and let $\Sigma_0 = \text{diag}(1000, \dots, 1000)$. We then specify

$$Y_s \mid \tau, \beta, \sigma_\epsilon^2 \sim \text{Normal}\{\mu_s(\tau, \beta), \sigma_\epsilon^2 I\},$$

$$\log(\tau - 1/d) \sim \text{Normal}(-1, 1),$$

$$\sigma_\epsilon^2 \sim \text{InverseGamma}(0.001, 0.001),$$

$$\beta \sim \text{Normal}(0, \Sigma_0),$$

$$\beta \mid \tau, \sigma_\epsilon^2 \sim \text{Normal}\{(\Sigma_0^{-1} + M^T M / \sigma_\epsilon^2)^{-1} M^T Y_{\mathcal{D}}^{vec} / \sigma_\epsilon^2, (\Sigma_0^{-1} + M^T M / \sigma_\epsilon^2)^{-1}\},$$

$$\sigma_\epsilon^2 \mid \beta, \tau \sim \text{InverseGamma}\{0.001 + (nN)/2, 0.001 + (Y_{\mathcal{D}}^{vec} - M\beta)^\top (Y_{\mathcal{D}}^{vec} - M\beta)/2\},$$

$$\log(\tau - 1/d) \mid \beta, \sigma_\epsilon^2 \propto \text{Normal}_Y(\mu^{vec}, \sigma_\epsilon^2 I) \times \text{Normal}_{\log(\tau - 1/d)}(0, 100).$$

Web Appendix C: Simulation sensitivity analysis

To see how robust the method is to incorrect response form and poor-fitting propensity scores, we conduct a small sensitivity analysis. Our method assumes an i.i.d. standard normal error, but the “true” error term is alternated between correctly specified i.i.d. standard normal error, as well incorrectly specified i.i.d. error with a centered exponential distribution and spatially dependent error generated from a Gaussian process with the same spatial range as the exposure ($\tau = 0.3$). All “true” error terms have mean zero and variance one. We also alternate between having our method utilize a propensity score with a zero-inflated lognormal and exponential distribution. The former likely has a better fit to the distribution of \tilde{A} , whereas the exponential distribution has fewer

parameters and likely has a worse fit. Because we do not generate the \tilde{A} data with a propensity score, we cannot call these correctly/incorrectly specified. For all settings the “2 circles” $h(X_{\mathcal{D}})$ is used and 500 repetitions run.

Tables 3 and 4 give the bias and coverage under these different scenarios. The point estimation appears largely robust to the deviations from the original settings. Bias does not appear markedly worse in any of the settings, with the exception of the exponential propensity score/incorrectly specified exponential error term, in which δ_1 has some negative bias. The coverage tends to be reasonable for all settings with i.i.d. error. However, coverage is decidedly low for δ_2 and τ under the spatial error settings. Of course, if exploratory analysis suggests the residuals are dependent, then spatial basis functions, covariates or correlation could be added to the Bayesian hierarchical model to improve uncertainty quantification.

Table 3: Simulation bias (standard errors) under alternative propensity score forms and correctly/incorrectly specified response form. Simulations use 10×10 grids and the “2 circles” $h(X_{\mathcal{D}})$ as described above. The zero-inflated lognormal (ZILN) propensity score is the recommended propensity score; the exponential propensity score is assumed to be a poor fit for the distribution of \tilde{A} . The method assumes i.i.d. standard normal error. All “true” error terms have mean zero and variance one. Bias and standard errors are multiplied by 1,000.

Setting	δ_1	δ_2	τ
ZILN prop. score, i.i.d. normal error	-0.9 (2)	-0.1 (1.2)	0.5 (0.3)
ZILN prop. score, i.i.d. exponential error	1.1 (1.9)	-0.4 (1.1)	0.3 (0.3)
ZILN prop. score, spatial normal error	-1.3 (2)	-1.7 (1.9)	0.5 (0.6)
Exponential prop. score, i.i.d. normal error	-1.5 (2)	1.8 (1.2)	-0.2 (0.3)
Exponential prop. score, i.i.d. exponential error	-5.8 (1.9)	-0.8 (1.2)	0.1 (0.3)
Exponential prop. score, spatial normal error	0.1 (1.9)	0.1 (1.8)	0.6 (0.6)

Web Appendix D

Wildfire analysis under alternative kernels:

To test the sensitivity to the form of the kernel function in the definition of the indirect exposure, we refit the wildfire data with two additional kernels. The Gaussian kernel is defined as

Table 4: Simulation coverage (standard errors) under alternative propensity score forms and correctly/incorrectly specified response form. Simulations use 10×10 grids and the “2 circles” $h(X_{\mathcal{D}})$ as described above. The zero-inflated lognormal (ZILN) propensity score is the recommended propensity score; the exponential propensity score is assumed to be a poor fit for the distribution of \tilde{A} . The method assumes i.i.d. standard normal error. All “true” error terms have mean zero and variance one.

Setting	δ_1	δ_2	τ
ZILN prop. score, i.i.d. normal error	94.6 (1)	89.8 (1.4)	90.4 (1.3)
ZILN prop. score, i.i.d. exponential error	96.2 (0.9)	92 (1.2)	91.8 (1.2)
ZILN prop. score, spatial normal error	94.6 (1)	70.2 (2)	52 (2.2)
Exponential prop. score, i.i.d. normal error	94.6 (1)	87.4 (1.5)	90.4 (1.3)
Exponential prop. score, i.i.d. exponential error	94.2 (1)	90.2 (1.3)	90.8 (1.3)
Exponential prop. score, spatial normal error	94.8 (1)	69.8 (2.1)	57.4 (2.2)

$\exp\{-(d/\tau)^2\}$; the exponential kernel is defined as $\exp(-d/\tau)$; and the uniform kernel is defined as $\mathbf{1}(d \leq \tau)$, where d is euclidean distance. However, the direct and spill-over effect estimates are similar for all kernels.

Table 5: Posterior mean (95% interval) exposure effects for generalized propensity score model, with alternative kernels.

Kernel	Direct Effect (δ_1)	Spill-over Effect (δ_2)	Bandwidth (τ)
Gaussian	1.03 (0.93, 1.14)	0.13 (0.03, 0.25)	1.53 (1.17, 2.88)
Exponential	1.04 (0.94, 1.15)	0.04 (0.01, 0.12)	4.04 (1.19, 15.51)
Uniform	1.03 (0.92, 1.13)	0.06 (0.03, 0.08)	2.33 (1.73, 2.51)

References

Hirano, K. & Imbens, G. W. (2004). The propensity score with continuous treatments. *Applied Bayesian modeling and causal inference from incomplete-data perspectives*, 226164, 73–84.