# Discussion of the manuscript: Spatial+: a novel approach to spatial confounding

Georgia Papadogeorgou Department of Statistics, University of Florida

I congratulate Dupont, Wood and Augustin (DWA hereon) for providing an easy-to-implement method for estimation in the presence of spatial confounding, and for addressing some of the complicated aspects on the topic. The method regresses the covariate of interest on spatial basis functions and uses the residuals of this model in an outcome regression. The authors show that, if the covariate is not completely spatial, this approach leads to consistent estimation of the conditional association between the exposure and the outcome.

Below I discuss conceptual and operational issues that are fundamental to inference in spatial settings: (i) the target quantity and its interpretability, (ii) the non-spatial aspect of covariates and their relative spatial scales, and (iii) the impact of spatial smoothing. While DWA provide some insights on these issues, I believe that the audience might benefit from a deeper discussion. In what follows, I focus on the setting where a researcher is interested in interpreting the relationship between a given covariate and an outcome. I refer to the covariate of interest as the *exposure* to differentiate it from the rest.

## 1 What are we estimating?

First, I discuss the target of estimation in the manuscript and in the literature. This target is key in deciding what is the appropriate method for estimation.

I begin by introducing some notation. Assume that Z and Y are the exposure and outcome of interest, and covariates X include all the variables that confound this relationship. I use the subscript *i* to denote the value of the variables at location i = 1, 2, ..., n. In what follows, I assume that X is comprised of four random variables: (a) a measured variable C that can be spatial or not, and unmeasured variables of the following three types: (b)  $S_1$  that is completely spatial, (c)  $S_2^+$ that includes a spatial and an independent component, and (d) U that is not spatial. Similarly to DWA, we can think of  $S_2^+$  as a random variable that can be written as  $S_2^+ = S_2 + E$  for  $S_2$  being a completely spatial random variable and E representing a component without spatial structure.

For simplicity of discussion, I also assume that the following additive model holds:

$$Y_{i} = \beta_{0} + \beta_{1} Z_{i} + \beta_{2} C_{i} + \beta_{3} S_{i1} + \beta_{4} S_{i2}^{+} + \beta_{5} U_{i} + \epsilon_{i}, \qquad (1)$$

where  $\epsilon_i \stackrel{iid}{\sim} N(0, \sigma^2)$ . This model implies that any spatial structure in the outcome Y is due to spatial dependencies in its predictors, though this assumption could be relaxed.

### 1.1 The "unconditional" association

A researcher only has access to  $(C_i, Z_i, Y_i)$ , for i = 1, 2, ..., n, while the remaining three random variables in X are not measured. With these data available, a first-step analysis could regress the

outcome, Y on the exposure and the sole observed covariate, (Z, C). Even if not directly stated, the target of estimation in this regression is the relationship between the exposure and outcome of interest, while conditioning solely on the measured variable:

## $(Y,Z) \mid C.$ (Unconditional target)

Even though the remaining variables in X are not accounted for, this relationship can be of explicit scientific interest, especially if the measured variable C is an interpretable characteristic of the population at hand. For example, Z, Y and C can represent a subject's exposure to a given pollutant, their blood pressure levels, and age, respectively. In this case, the unconditional target represents the expected change in an individual's blood pressure for a unit increase in the pollutant, after adjusting for age. I refer to this target as the *unconditional* association as it does not condition on the spatial random variables. For *this* target of estimation, the OLS estimator will be *un*biased.

However, not conditioning on the remaining variables in X leads to complications for *inference* for the estimated coefficient. If at least one of the coefficients of the spatial predictors are nonzero ( $\beta_3 \neq 0$ , or  $\beta_4 \neq 0$ ), the residuals of this regression would be spatially correlated. Since correlation in the model residuals that is not accounted for would lead to incorrect inference on model coefficients, researchers often include a spatial random effect with covariance structure that aims to mimic the residual dependence. The goal of the random effect is to explain the spatial dependence in the outcome that is *not* already explained by (Z, C). However, if the exposure Z is, at least partially, spatial, and since the additional model components are themselves spatial, they are often spuriously associated due to structural dependencies (Hodges & Reich 2010, Lee & Ogburn 2020). Therefore, the random effect "involuntarily" captures part of the spatial variability in the outcome that was originally explained by the exposure, and the coefficient of the exposure in the model with the spatial random effect no longer has the same interpretation as the coefficient of the non-spatial model.

The restricted spatial regression (RSR) approach ensures that the random effect is orthogonal to the exposure of interest. That means that the target of estimation for RSR is the same as the one from the non-spatial model, and RSR is truly useful, applicable, and *un*biased for the relationship between the exposure and the outcome conditional *only* on the measured variables.

#### 1.2 The "spatially-conditional" association

The target of estimation for DWA is different. Instead of focusing on the (Y, Z) relationship when conditioning solely on C, DWA aim to also adjust for all spatial confounders. Informally, the target is set to learning

$$(Y, Z) \mid C$$
, all spatial confounders, (Spatially-conditional target)

though as we will see below, it truly achieves an alternative quantity. Therefore, in the presence of unmeasured spatial confounders, the targets of estimation for RSR and Spatial+ are different, the coefficient of the exposure in the two models has a different interpretation, and the methods should not be directly compared.

Since confounders are predictors of the outcome, an outcome model that does not adjust for spatial confounders will exhibit residual spatial dependence. At the same time, since they are predictors of the exposure, adding components in the regression model that are orthogonal to the exposure (as in RSR) will not capture them. However, since conditioning on any spatial variable that is *not* a confounder does not change the relationship under study, the spatially-conditional target above can be re-written as

 $(Y, Z) \mid C$ , all spatial variables.

Writing the target of estimation explicitly as conditioning on any spatial variable clarifies two aspects of the manuscript. (1) If the exposure Z is fully spatial and it does not include any independent components, then the spatially-conditional relationship cannot be learnt without additional assumptions, since the exposure is perfectly collinear with the conditioning set.<sup>1</sup> (2) If Z includes an independent component, the spatially-conditional relationship between the exposure and the outcome can be learnt while focusing on the *non*-spatial aspect of the data. The former is central in DWA's manuscript. The latter motivates three of the approaches discussed in the manuscript: the spatial model that includes basis functions directly in the outcome model, the gSEM model that controls for space both in the exposure and in the outcome model, and the proposed Spatial+ model that regresses the exposure itself on the basis functions and uses the exposure model residuals in the outcome model. When no spatial smoothing is performed, all of these methods adjust for any completely spatial structure in the data, and are therefore unbiased for the spatially-conditional association, as expected, and as illustrated in the simulations. The issues that arise with smoothing are discussed in Section 2.

The latter observation is also key in Schnell & Papadogeorgou (2020), where the authors link this observation to the assumption of positivity in causal inference, and discuss that estimation of the exposure coefficient can only be accurately achieved if the exposure itself varies within levels of the spatial confounder. In the setting of DWA where the spatial confounder can vary at *any* spatial frequency, this observation by Schnell & Papadogeorgou (2020) is translated to Z including an independent component.

As discussed in Section 1.1, RSR studies the exposure-outcome relationship when conditioning only on the measured variables. As discussed above, Spatial+ aims to extend the conditioning set to include spatial confounders. Even though conditioning on a more complete set of confounders is generally preferred, I illustrate that the estimated quantities from Spatial+ are not always interpretable. Take for example the true outcome model in Equation (1), and assume that both spatial variables are important ( $\beta_3 \neq 0, \beta_4 \neq 0$ ) while the independent variable is not ( $\beta_5 = 0$ ). Then, we can re-write this model as

$$Y_{i} = \beta_{0} + \beta_{1}Z_{i} + \beta_{2}C_{i} + \beta_{3}S_{i1} + \beta_{4}(S_{i2} + E_{i}) + \epsilon_{i},$$

where  $S_2$  and E decompose  $S_2^+$  in its spatial and independent components. The target of estimation that conditions on measured variables and *all* spatial confounders would be

$$(Y,Z) \mid C, S_1, S_2^+ \equiv (Y,Z) \mid C, S_1, S_2, E$$

Therefore, this target would represent the exposure-outcome relationship when conditioning on the variables  $C, S_1, S_2^+$ . However, the E part of  $S_2^+$  is *not* spatial and cannot be captured by any spatial basis function. Therefore, even though the implicit target would condition on *all* spatial variables, the method instead estimates

$$(Y,Z) \mid C, S_1, S_2.$$
 (Spatially-conditional quantity achieved)

Since the achieved target conditions solely on *part* of a confounder, the estimated conditional exposure-outcome relationship has a complicated interpretation. Whether this relationship is more or less interpretable than the unconditional target of Section 1.1 is, I believe, up for debate!

<sup>&</sup>lt;sup>1</sup>To make progress in this setting, one might need to make assumptions on the relative spatial scales of the spatial confounders and the exposure, though I ignore this scenario here. This setting is investigated in Schnell & Papadogeorgou (2020), and a related discussion in terms of the spatial scales of the spatial confounders themselves can be found in Section 1.3.

In the previous blood pressure example, let's think of  $S_1$  as representing outdoor temperature which is completely spatial, and  $S_2^+$  as representing eating habits (such as intake of salt, or fried food). Even though eating habits are partially spatial due to local culture, they also include an independent, subject-specific component. Then, the coefficient estimated by Spatial+ would describe the expected change in blood pressure for a unit increase in the pollutant when adjusting for age, temperature, and *part* of someone's eating habits.

#### **1.3** Assumptions for interpretable spatially-conditional target

Therefore, the achieved target of estimation for the Spatial+ model often has a complicated interpretation, largely due to variables that include both spatial and independent components. So when can Spatial+ be used to acquire interpretable conditional associations?

The first case corresponds to the scenario where there are no spatial confounders of the exposureoutcome relationship that include an independent component (Papadogeorgou et al. 2019). This assumption is implicitly made in equation (1) of DWA. If all covariates in the analysis are either fully spatial or not spatial at all (if, for example,  $S_2^+ = S_2$ ), then Spatial+ estimates the exposureoutcome relationship conditional on all spatial variables, which has a clear interpretation. In turn, if there exist covariates that include both spatial and independent components such as  $S_2^+$ , these variables must not confound the exposure-outcome relationship when conditioning on fully spatial variables. This can be satisfied if

$$E \perp\!\!\!\perp Y \mid Z, C, S_1, S_2, \quad \text{or} \quad E \perp\!\!\!\perp Z \mid C, S_1, S_2$$

where  $\perp$  is used to denote independence. However, since all of these spatial variables are unmeasured, assuming that none of the spatial confounders includes an independent component remains, exactly that, an assumption.

So what happens when we are not willing to make the assumption that all spatial confounders are fully spatial? Progress can still be made in this setting if we are willing to make an alternative assumption on the relative spatial scales of spatial confounders, and update Spatial+ accordingly. Remember that the issue arises because  $S_1$  and  $S_2^+$  are both spatial and unmeasured, which forbids us from directly differentiating between the two. To make progress, we could make an assumption which allows us to separate the one spatial variable from the other. Such an assumption could be the following:  $S_2^+$  varies only in higher spatial frequencies than  $S_1$ . If this assumption holds, then including spatial basis functions that represent only low-frequency spatial components would capture  $S_1$  fully, while not capturing any aspect (spatial or independent) of  $S_2^+$  (Antonelli et al. 2017, Keller & Szpiro 2020). This could be acheived, for example, using wavelet or fourier basis functions for which the frequency of the basis functions used is directly controllable. Under this assumption and the corresponding adjustment of Spatial+, the method would be estimating the exposure-outcome relationship when conditioning on C and  $S_1$  only, which would have a clear interpretation.

#### 1.4 Spatial confounding when estimating causal effects

If the set of covariates X includes all the confounders of the exposure-outcome relationship, then an outcome model which adjusts for the exposure and all of the confounders and is correctly specified can be the base for estimating causal effects (Robins 1986, Snowden et al. 2011, g-computation). I focus first on linear additive models such as the one in Equation (1). In this case, the coefficient of the exposure,  $\beta_1$ , can be interpreted causally as the effect of increasing the exposure Z by one unit on the expected outcome Y. However, the coefficient will not have the same, causal interpretation

if we do not adjust for the variables  $S_1$ ,  $S_2$ + and U, which are unmeasured, nor if we only adjust for a subset of them. Since U is a completely independent random variable, we cannot hope to identify the causal effect of the exposure on the outcome in this manner. What's more, adjusting for *some* of the confounders, and not all, does *not* necessarily mean that the bias for estimating causal effects is reduced. Instead, it could *increase*, if confounding from spatial and independent variables are in opposite directions! Therefore, if unmeasured independent variables confound the exposure-outcome relationship, adjusting for spatial unmeasured confounders does not necessarily bring us "closer" to estimating causal effects.

However, if the true outcome model is non-linear, or includes exposure-covariate interactions, the complications are not resolved, even if we assume the absence of independent, unmeasured confounders (for example, if U does not exist, and  $S_2$ + is fully spatial). In order to use the outcome model for estimating causal effects, g-computation teaches us that we need to integrate over the distribution of the confounders, and the coefficient of the exposure is not necessarily causally interpretable (though log-linear models are an exception for a specific type of causal contrast (Schnell & Papadogeorgou 2020)). When the spatial confounders are unmeasured, and their distribution cannot be learnt from data, integrating over their distribution cannot be achieved directly, and requires additional assumptions (see, for example Christiansen et al. 2020). Recently, Reich et al. (2021) wrote a review article regarding spatial causal inference, including a thorough discussion on spatial confounding.

## 2 Spatial smoothing

Based on the discussion of Section 1, we have established that the non-spatial and RSR models estimate a different quantity compared to the rest. We have also established that regressing away all spatial aspect of the outcome or the exposure (through, for example, *unpenalized* splines) would adjust fully for the spatial component of all spatial confounders. Assuming that this quantity has a useful interpretation, here I focus on investigating the following questions: *what happens with spatial smoothing*, and *how much should predictive accuracy matter?* 

Since spatial smoothing excludes the higher frequency components, it implies that the achieved target represents the exposure-outcome association conditional on the spatial variables that vary at the included, lower frequencies *only*. If all spatial confounders vary at lower frequencies, the estimated quantity will truly align with the spatially-conditional association. However, even though spatial smoothing is expected to lead to lower prediction MSE, it is not immidiately clear whether it should be used for coefficient estimation. In the presence of high frequency spatial confounders, spatial smoothing might lead us to "miss" these confounders.

I investigate two scenarios where the Spatial and Spatial+ models might miss strong, high frequency spatial confounders, and return biased estimates of the spatially-conditional association.<sup>2</sup> These spatial variables can be strong confounders in at least two different ways: (1) they can be important predictors of the exposure and weak predictors of the outcome, or (2) they can be weak predictors of the exposure, and important predictors of the outcome. In the former case, since the high-frequency spatial variables are strong predictors of the exposure, one hopes that the choice of penalty in the exposure model for Spatial+ will not penalize these high frequency components heavily. The reverse is true for the Spatial model, for which we would expect that it will penalize the high frequency spatial basis functions more heavily, as these are weak predictors of the outcome. Therefore, in this scenario, we would expect Spatial+ to perform better than the Spatial model for estimating the spatially-conditional association.

<sup>&</sup>lt;sup>2</sup>Remember that the Spatial model employs spatial basis functions in the outcome model.

spatial confounders are weak predictors of the exposure, it is expected that the chosen penalty for the exposure model according to Spatial+ will penalize the high frequency components heavily, returning biased estimates of the spatially-conditional association. In contrast, we would expect the Spatial model that includes spatial basis functions in the outcome model directly to perform more accurately.

Even though it is prudent not to overly generalize the message from this discussion, I believe that deeper investigation is warranted on how the methods perform under different confounding scenarios. That is essential under the light of similar observations from other parts of the literature that focus on estimation rather than prediction, and encourage the use of *both* an exposure and an outcome model simultaneously (e.g., Wang et al. 2012, Belloni et al. 2014, Shortreed & Ertefaie 2017, Antonelli et al. 2019). Specificially in the setting of spatial confounding, the use of both models is proposed in at least two approaches (gSEM, investigated in the manuscript; Thaden & Kneib (2018) and Schnell & Papadogeorgou (2020)). If we view gSEM as a combination of the Spatial and Spatial+ models in that it adjusts for spatial effects in both the exposure and the outcome model, one might expect gSEM to be more robust to high frequency spatial confounding than any of the two methods separately.

Finally, the work referenced above has been crucial in illuminating that *prediction* and *estimation* are two separate goals, and that methodology that is appropriate for one can be inappropriate for the other. In light of this work, even though the methods' predictive accuracy can inform us of how well each model fits the data, it should *not* be used as a model-choice criterion when we are interested in interpreting estimated parameters. In fact, a model with lower AIC might often have higher bias for the coefficient of interest.

## References

- Antonelli, J., Parmigiani, G. & Dominici, F. (2019), 'High-dimensional confounding adjustment using continuous spike and slab priors', *Bayesian analysis* 14(3), 805.
- Antonelli, J., Schwartz, J., Kloog, I. & Coull, B. A. (2017), 'Spatial multiresolution analysis of the effect of pm2. 5 on birth weights', *The annals of applied statistics* 11(2), 792.
- Belloni, A., Chernozhukov, V. & Hansen, C. (2014), 'Inference on treatment effects after selection among high-dimensional controls', The Review of Economic Studies 81(2), 608–650.
- Christiansen, R., Baumann, M., Kuemmerle, T., Mahecha, M. D. & Peters, J. (2020), 'Towards causal inference for spatio-temporal data: Conflict and forest loss in colombia', *arXiv preprint* arXiv:2005.08639.
- Hodges, J. S. & Reich, B. J. (2010), 'Adding spatially-correlated errors can mess up the fixed effect you love', *The American Statistician* 64(4), 325–334.
- Keller, J. P. & Szpiro, A. A. (2020), 'Selecting a scale for spatial confounding adjustment', Journal of the Royal Statistical Society. Series A, (Statistics in Society) 183(3), 1121.
- Lee, Y. & Ogburn, E. L. (2020), 'Network dependence can lead to spurious associations and invalid inference', Journal of the American Statistical Association pp. 1–15.
- Papadogeorgou, G., Choirat, C. & Zigler, C. M. (2019), 'Adjusting for unmeasured spatial confounding with distance adjusted propensity score matching', *Biostatistics* 20(2), 256–272.

- Reich, B. J., Yang, S., Guan, Y., Giffin, A. B., Miller, M. J. & Rappold, A. (2021), 'A review of spatial causal inference methods for environmental and epidemiological applications', *International Statistical Review*.
- Robins, J. (1986), 'A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect', *Mathematical modelling* 7(9-12), 1393–1512.
- 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.
- Shortreed, S. M. & Ertefaie, A. (2017), 'Outcome-adaptive lasso: variable selection for causal inference', *Biometrics* **73**(4), 1111–1122.
- Snowden, J. M., Rose, S. & Mortimer, K. M. (2011), 'Implementation of g-computation on a simulated data set: demonstration of a causal inference technique', *American journal of epidemiology* 173(7), 731–738.
- Thaden, H. & Kneib, T. (2018), 'Structural equation models for dealing with spatial confounding', The American Statistician **72**(3), 239–252.
- Wang, C., Parmigiani, G. & Dominici, F. (2012), 'Bayesian effect estimation accounting for adjustment uncertainty', *Biometrics* 68(3), 661–671.