Biometrics 66, 783–792
September 2010
DOI: 10.1111/j.1541-0420.2009.01323.x

Spatial Cluster Detection for Weighted Outcomes Using Cumulative Geographic Residuals

Andrea J. Cook,[1,2,*] Yi Li,[3] David Arterburn,[1] and Ram C. Tiwari[4]

[1]Biostatistics Unit, Group Health Research Institute, Seattle, Washington 98101, U.S.A.
[2]Department of Biostatistics, University of Washington, Seattle, Washington 98105, U.S.A.
[3]Department of Biostatistics, Harvard School of Public Health and the Dana Farber Cancer Institute, Boston, Massachusetts 02115, U.S.A.
[4]Office of Biostatistics, CDR, FDA, Silver Spring, Maryland 20993, U.S.A.
[*]*email:* cook.aj@ghc.org

Summary. Spatial cluster detection is an important methodology for identifying regions with excessive numbers of adverse health events without making strong model assumptions on the underlying spatial dependence structure. Previous work has focused on point or individual-level outcome data and few advances have been made when the outcome data are reported at an aggregated level, for example, at the county- or census-tract level. This article proposes a new class of spatial cluster detection methods for point or aggregate data, comprising of continuous, binary, and count data. Compared with the existing spatial cluster detection methods it has the following advantages. First, it readily incorporates region-specific weights, for example, based on a region's population or a region's outcome variance, which is the key for aggregate data. Second, the established general framework allows for area-level and individual-level covariate adjustment. A simulation study is conducted to evaluate the performance of the method. The proposed method is then applied to assess spatial clustering of high Body Mass Index in a health maintenance organization population in the Seattle, Washington, USA area.

Key words: Body mass index; Cumulative residuals; Generalized estimating equations; Socioeconomic status; Spatial cluster detection; Weighted linear regression.

1. Introduction

The increasing prevalence of obese/overweight individuals is a growing public health concern, causing a tremendous strain on the health care system. Identifying regions or neighborhoods that have elevated body mass index (BMI) compared to the rest of the area could help direct the distribution of resources for obesity prevention and treatment programs. Furthermore, if the elevated BMI regions can be explained by factors, such as area-level socioeconomic status (SES), walkability, and density of fast food establishments, these may shape an appropriate intervention program tailored for the residents in a particular geographic region. Typically, these types of data are not available at the individual's residence location level, but are aggregated to census tracts, counties, or states. To be able to conduct such analyses, a spatial cluster detection method is needed that handles weighted continuous outcomes, while being able to adjust for area-level covariates.

Currently, numerous spatial cluster detection methods are available for the analysis of individual level data using a wide variety of outcomes. For example, there are methods for binary outcomes to identify areas with elevated prevalence of disease and for count outcomes to identify excess rates of incidence or mortality (Turnball et al., 1990; Kulldorff, 1997; Tango, 2000; Duczmal and Assunção, 2004; Patil and Taillie, 2004; Tango and Takahashi, 2005; Kulldorff et al., 2006). There are also several methods for censored continuous outcomes, which explore potential spatial clusters for detection of time to early event (Cook, Gold, and Li, 2007; Huang, Kulldorff, and Gregorio, 2007). However, the only available method for aggregate continuous data is a recently accepted spatial scan statistic for weighted normal outcomes (Huang et al., 2009). This latter method is not able to incorporate covariate adjustment (area or individual-level) and is not a general approach for any weighted noncontinuous outcome, which is the key for our application of interest.

The application of interest is a prospective cohort study evaluating areas of elevated BMI of adult females from a mixed-model health plan and delivery system located in the Seattle area in western Washington State, USA. The purpose of this analysis is to assess whether there are areas of elevated BMI and if these elevated areas can be explained by area-level SES predictors. BMI is measured at the individual level, but spatial location is measured only at the aggregate, census-tract, level due to data availability. Furthermore, even though BMI is a continuous outcome, it is highly skewed with heavier tails at the upward end of the distribution than would be expected if it was normal distributed, that is, methods assuming a normal assumption are not valid. Our proposed spatial cluster detection approach is able to handle aggregate level, skewed data, while being able to adjust for both individual and area-level covariates, which is critical for this analysis.

We present in this article a general statistical approach to quantifying spatial cluster detection for weighted outcomes that can be continuous, but is applicable for most outcome types. The new method is presented in Section 2 and a simulation study evaluating its properties can be found in Section 3. In Section 4, we apply the proposed weighted cumulative geographic residual method to an analysis assessing elevated areas of BMI for adult females in western Washington State, USA. We conclude with a general discussion in Section 5.

2. Method

2.1 *Weighted Geographic Cumulative Residual Method*

We exemplify the development of our test statistic in the framework of a continuous outcome, though the formulation may be easily generalized to any binary/discrete data with proper link functions (e.g., Poisson data with a log link function). Suppose the outcome for region $i (i = 1, \ldots, n), Y_i$, is continuous, with a $1 \times p$ vector of population characteristics, $\mathbf{X}_i$, and the geographic centroid of the region is (s_i, r_i). Under the null hypothesis that the outcome is independent of geographic location (s_i, r_i), conditional on a given set of area-level covariates, $\mathbf{X}_i$, we assume that the continuous outcome follows the following model with only the first two moments specified:

$$Y_i = \mathbf{X}_i \boldsymbol{\beta} + e_i, \quad e_i \overset{ind}{\sim} (0, \sigma^2/w_i), \tag{1}$$

where $\boldsymbol{\beta}$ is a $p \times 1$ vector of unknown regression parameters, σ^2 is an unknown variance parameter, and $w_i > 0$ are the weights assigned to the geographic area $i (i = 1, \ldots, n)$. The error terms, e_i, are independent with mean 0 and variance σ^2/w_i. The weights, w_i, are assumed known and may represent the (extra) regional variability in the characteristic of interest Y_i. It may be taken to be the inverse of the local variance of Y or the population size, depending on the application. One can estimate $\boldsymbol{\beta}$ by $\hat{\boldsymbol{\beta}}$, and σ^2 by $\hat{\sigma}^2$, using the following estimating equations, $U_\beta(\boldsymbol{\beta}) = \sum_{i=1}^n U_{i\beta} = \sum_{i=1}^n w_i \mathbf{X}_i^T (Y_i - \mathbf{X}_i \boldsymbol{\beta})/\sigma^2 = 0$ and $U_\sigma(\sigma) = \sigma^2 - 1/n \sum_{i=1}^n w_i (Y_i - \mathbf{X}_i \boldsymbol{\beta})^2 = 0$ and simultaneously solving both equations. These estimating equations are derived assuming $e_i \overset{ind}{\sim} N(0, \sigma^2/w_i)$ and therefore are the most efficient estimators if the weighted normal model assumption is correct, but the theory holds without assuming normality. This makes the proposed method robust to distributional assumptions. From these estimating equations one can use the residuals, $\hat{e}_i = Y_i - \mathbf{X}_i \hat{\boldsymbol{\beta}}$, to test for spatial cluster patterns by finding areas with higher than expected sum of residuals. Sum of residuals is a natural test statistic to use because it has a defined distribution and it has monotonic properties that areas with higher sum of residuals indicate areas with higher then expected outcomes. Further, the sum may be preferred over a mean, because it weights those potential spatial clusters with more observations higher, which is preferred because they have more statistical information. Our approach follows the line of the goodness-of-fit testing initially proposed by Su and Wei (1991) for generalized linear models.

We consider a two-dimensional moving block process over location $(x_1, x_2), Z_{loc}(x_1, x_2 \,|\, b)$, which depends on geographic locations for a fixed block size b as follows:

$$Z_{loc}(x_1, x_2 \,|\, b) = \frac{1}{\sqrt{n}} \sum_{i=1}^n W_i(x_1, x_2 \,|\, b) \hat{e}_i, \tag{2}$$

where $W_i(x_1, x_2 \,|\, b) = I(x_1 - b < s_i \le x_1 + b, x_2 - b < r_i \le x_2 + b) w_i$, a weighted location indicator function. For given location, $(x_1, x_2), \sqrt{n} Z_{loc}$ is the weighted sum of residuals from regions that are within a box with side length $2b$ around point (x_1, x_2). A spatial cluster would occur in areas with a higher intensity of an outcome that implies a larger value of $Z_{loc}(x_1, x_2 \,|\, b)$.

The exact distribution of $Z_{loc}(x_1, x_2 \,|\, b)$ cannot be solved analytically so we propose to use an asymptotic equivalent distribution to approximate the true distribution. We consider the following pseudo moving block process in $(x_1, x_2), \hat{Z}_{loc}(x_1, x_2 \,|\, b)$, as

$$\hat{Z}_{loc}(x_1, x_2 \,|\, b) = \frac{1}{\sqrt{n}} \sum_{i=1}^n [W_i(x_1, x_2 \,|\, b) \hat{e}_i + \nu(x_1, x_2 \,|\, b) \mathbf{I}^{-1}(\hat{\boldsymbol{\beta}}) U_{i\beta}(\hat{\boldsymbol{\beta}})] G_i, \tag{3}$$

where

$$\nu(x_1, x_2 \,|\, b) = -\sum_{i=1}^n W_i(x_1, x_2 \,|\, b) \partial \mu / \partial \boldsymbol{\beta} = -\sum_{i=1}^n W_i(x_1, x_2 \,|\, b) \mathbf{X}_i.$$

$\mathbf{I}(\boldsymbol{\beta}) = -\partial U_\beta / \partial \boldsymbol{\beta}$ and $G_i (i = 1, \ldots, n)$ are independent mean 0 and variance 1 random variables that are also independent of $(Y_i, \mathbf{X}_i, s_i, r_i)$. It follows that the asymptotic conditional distribution of the pseudo process $\hat{Z}_{loc}(x_1, x_2 \,|\, b)$ given the observed data $(Y_i, \mathbf{X}_i, s_i, r_i) (i = 1, \ldots, n)$ is equivalent to the limit distribution of $Z_{loc}(x_1, x_2 \,|\, b)$ assuming that geographic location, (s_i, r_i), is independent of continuous outcome, Y_i, after adjusting for covariates, $\mathbf{X}_i$, with the first two moments (1) being correctly specified. This result can be obtained by using the independence between the residuals and geographic location under the null hypothesis. Details of the proof are outlined in the Web Appendix.

This asymptotic result immediately allows us to approximate the null distribution of $Z_{loc}(x_1, x_2 \,|\, b)$ with a large number, say, N, realizations of $\hat{Z}_{loc}(x_1, x_2 \,|\, b)$, $(\hat{Z}_{1,loc}(x_1, x_2 \,|\, b), \ldots, \hat{Z}_{N,loc}(x_1, x_2 \,|\, b))$, by repeatedly simulating independent samples of $(G_1, \ldots, G_n)$, while fixing the data $(Y_i, \mathbf{X}_i, s_i, r_i) (i = 1, \ldots, n)$ at their observed values. However, for the particular purpose of spatial cluster detection, it is important to allow the data to depict the best cluster size. Therefore, we consider a finite vector of length M of varying cluster sizes, denoted by $\mathbf{b} = (b_1, \ldots, b_M)$, where each b_m denotes half of the edge length of the potential square cluster. Accordingly, we define a cluster detection test statistic to test existence of any spatial clusters,

$$S_{loc} = \sup \left[\sup_{x_1, x_2} Z_{loc}(x_1, x_2 \,|\, b_1), \ldots, \sup_{x_1, x_2} Z_{loc}(x_1, x_2 \,|\, b_M) \right].$$

Continuous mapping theorem will show that S_{loc} has the same limiting distribution as the following stochastic process, conditional on the observed data:

$$\hat{S}_{loc} = \sup \left[\sup_{x_1, x_2} \hat{Z}_{loc}(x_1, x_2 \,|\, b_1), \ldots, \sup_{x_1, x_2} \hat{Z}_{loc}(x_1, x_2 \,|\, b_M) \right].$$

Hence, the empirical p-values can be computed as p-value $= \frac{\sum_{j=1}^{N} I[S_{loc} \le \hat{S}_{j,loc}]}{N}$, where $\hat{S}_{j,loc}$ is the $\hat{S}_{loc}$ evaluated at the jth realization of $\hat{Z}_{j,loc}$. In practice, to obtain the observed test statistic, S_{loc}, and simulated test statistics, $\hat{S}_{j,loc}$, it is necessary to create a finite grid of values over x_1, x_2, and b to approximate the continuous stochastic processes.

This hypothesis test can be inverted to form confidence bands around $Z_{loc}(x_1, x_2 \mid b)$ to find the values of (x_1, x_2, b) that have significantly higher average outcome than expected assuming the null hypothesis and the first two moments being correctly specified (1). Explicitly, $\{(x_1, x_2, b) : Z_{loc}(x_1, x_2 \mid b) \ge \hat{S}_{(.95N)}\}$, where $\hat{S}_{(.95N)}$ is the 95th percentile of all $\hat{S}_{j,loc}$. Therefore multiple clusters can be easily detected utilizing this proposed test statistic.

This overcomes a potential limitation of this method in which potential clusters assume a square cluster grid instead of a potentially more intuitive circular cluster grid. The square grid indicator function, $W_i(x_1, x_2 \mid b)$, was required for the theory to estimate the distribution of the two-dimensional stochastic process, $Z_{loc}(\cdot, \cdot \mid b)$, under the null. The theory, presented in detail in the Web Appendix, requires the component $W_i(x_1, x_2 \mid b)$ to be a monotone function element-wise, or independent, on each dimension (x_1, x_2) and therefore cannot be a circle. Since this approach readily finds multiple clusters that are significant, it is then able to detect nonsquare clusters by defining the final detected cluster as all areas that are significant. So this method is not restricted to finding only square clusters.

The weighted cumulative geographic residual method can be applied to aggregate data in which the outcome is continuous, such as mean body mass index (BMI) in a region or 5-year mortality of breast cancer in a region. We are able to use the population-size or inverse of variance weights for parameter estimation through the weighted linear regression model and directly in our weighting function for spatial cluster detection using the residuals. The method is able to adjust for area-level covariates, but adjustment for both individual and area-level covariates has not been incorporated. The next section will propose a simple two-stage approach to extend the method to a general covariate adjustment.

2.2 *Incorporating the Adjustment of Individual-Level Covariates*

Often datasets evaluating spatial clustering only have data available at an aggregate level including spatial locations, outcomes, and covariates. However, there are situations, such as the application of this article for the health care population, in which spatial locations may be available at the aggregate-level only, but outcomes and covariates are available at both the individual and aggregate level. Therefore it is important to extend our method to be able to address the question of whether spatial clustering exists after adjusting for individual-level covariates and if it still persists after adjusting for area-level covariates, such as SES.

We propose a simple two-stage approach to account for individual-level covariates. The first stage is the model that regresses individual-level covariates, $\mathbf{X}_{ij}^I$, on the outcome Y_{ij}, while still allowing for differences across regions,

$$Y_{ij} = \boldsymbol{\beta}_I \mathbf{X}_{ij}^I + \mathbf{U}_i + e_{ij}, \quad (4)$$

where subscript ij denotes individual $j (j = 1, \ldots, n_i)$ in region $i (i = 1, \ldots, n), \mathbf{X}_{ij}^I$ is a $p_I \times 1$ vector of individual-level covariates, $\mathbf{U}_i$ is the parameter of interest indicating region-specific effect, and e_{ij} is the residual error with mean 0 and variance σ^2.

Stage two applies the weighted cumulative residual method discussed in Section 2.1 using derived outcomes and weights from stage one. The derived outcome from stage one can be thought of as an individual-level covariate adjusted outcome and we can derive this outcome using two different approaches; (1) $\mathbf{U}_i$ is fixed effect or (2) $\mathbf{U}_i$ is a random effect. If the region effect is fixed then $\mathbf{U}_i$ is the mean effect of region i after adjusting for individual-level covariates, $\mathbf{X}_{ij}^I$. The individual-level adjusted outcome is $V_i = \hat{\mathbf{U}}_i + \hat{\boldsymbol{\beta}}_I \bar{\mathbf{X}}_{ij}^I$, where $\hat{\mathbf{U}}_i$ and $\hat{\boldsymbol{\beta}}_I$ are estimates from stage one assuming normal distribution estimating equations for e_{ij} distribution and $\bar{\mathbf{X}}_{ij}^I$ is the vector of the mean values of each covariate in the dataset. A potential efficient weight structure for stage two would be the inverse variance estimates of the fixed effect estimated parameter, $V_i, w_i = 1/\hat{\mathrm{V}}\mathrm{ar}(\hat{\mathbf{U}}_i + \hat{\boldsymbol{\beta}}_I \bar{\mathbf{X}}_{ij}^I)$.

Estimation of V_i and its variance may have issues for regions with small sample sizes. Therefore, when the fixed effect approach is not estimable we propose an alternate approach using a random effect framework to obtain individual-level adjusted outcomes. Now assume that $\mathbf{U}_i \overset{ind}{\sim} (\boldsymbol{\beta}_R, \sigma_R^2)$ and $e_{ij} \overset{ind}{\sim} (0, \sigma^2)$. For estimation one can choose any distribution, but it is most common to assume a normal distribution. Inference can be sensitive to distribution choice so it may be important to assess several distributions. Given the distribution assumption on $\mathbf{U}_i$ and e_{ij}, the empirical Bayes estimates for $\mathbf{U}_i$ can be obtained using the standard generalized linear mixed model (GLMM) estimation framework. Then, we propose using the following individual-level adjusted outcome:

$$B_i = \sum_{j=1}^{n_i} \left[Y_{ij} - \hat{\boldsymbol{\beta}}_I \mathbf{X}_{ij}^I - \hat{\mathbf{U}}_i \right] / n_i.$$

Assuming no relationship between aggregate spatial location and outcome Y_{ij} given $\mathbf{X}_{ij}^I$ the $E(B_i) = 0$ and $\mathrm{Var}(B_i) = (\sigma^2 + \sigma_R^2)/n_i$. Efficient weights to use for stage two would be inverse of the variance of $B_i, w_i = n_i/(\sigma^2 + \sigma_R^2)$.

We have proposed two different approaches for adjusting for individual-level covariates. The fixed effect approach is preferable since it limits modeling assumptions. Specifically, when using the random effect approach, estimation is sensitive to the random effects distribution assumed, which has been noted by others (Heagerty and Kurland, 2001; Litiere, Alonso, and Molenberghs, 2008). For similar reasons, we have proposed a two-stage approach instead of a single-stage approach, that is, using residuals directly from a GLMM that adjusts for both area and individual-level covariates as have been proposed by others for goodness-of-fit tests (Pan and Lin, 2005). Particularly, since we are in the situation of spatial clustering it may actually be inappropriate to apply a GLMM, which does not take into account spatial correlation, but only regional correlation, since it may lead to biased parameter estimates due to residual confounding (Wakefield, 2003). Therefore, the general modeling approach presented in this article tries to present methods that limit model assumptions. The next section evaluates the performance of the proposed cumulative geographic residual method in a simulation study.

3. Simulation Study

We conducted a simulation study calculating the Type I error and power for the proposed cumulative geographic residual test for weighted outcomes. For computational efficiency, we allowed a finite range for half edge length, b, of 0.5 to 3 sequenced by 0.1 and simulated $1000\hat{Z}_{loc}$ per dataset. We ran 1000 simulated datasets per calculation.

3.1 *Type I Error*

We first ran Type I error calculations to evaluate the validity of our proposed cluster detection method. The cumulative residual method is based on how well $\hat{Z}_{loc}$ approximates Z_{loc} that requires a sufficient sample size (e.g., number of regions). Our type I error simulation assumed a 10×10 unit-less square study area that we broke into 25, 36, 49, 100, and 225 square regions and varied the weights per region from 1, 20, 40, 60, and 80. The simulated dataset assumed that $Y_i \sim N(0, 1)$ with weight w_i. We defined Type I error as the proportion of simulations that detect a significant ($\alpha = 0.05$) cluster. See Web Table 1 for details, but Type I error increased to 0.05 as number of regions and weight per region increased and reached the correct level by 80 regions.

We then ran a simulation evaluating the effect on the Type I error when we assume the weights are known/fixed, when they actually are variable. For simplicity, we assume a 10×10 unitless grid with 100 equally sized square regions. We simulate the data assuming $Y_i \sim N(0, 1)$ with weight $w_i \sim \nu + \sigma * \mathit{Uniform}(0, 1)$. We varied ν to be 1, 10, 20, and 30 and σ to be 1, 5, 10, 20, and 30. Results are displayed in Web Table 2, but there were no obvious patterns on the effect of Type I error when assuming the weights are fixed when they are variable (Type I error ranged form 0.041 to 0.069).

3.2 *Power: No Covariates*

The next set of power calculations follow the lines of the simulations performed for the spatial scan statistic for weighted outcomes presented by Huang et al. (2009). The spatial scan statistic for weighted outcomes applies a likelihood-ratio statistic assuming a weighted normal linear regression model assessing if the outcome mean inside the potential cluster is greater than the outcome mean outside the potential cluster, $\mu_I > \mu_O$. Potential cluster areas are defined as all circular areas, with varying radius, consuming up to 50% of the study area. To calculate if the maximum cluster, defined as suprema of the likelihood-ratio statistic over all potential clusters, has a significantly higher outcome mean than the rest of the study population, a permutation test is performed to hold the overall Type I error level (i.e., handles multiple comparison issue).

For the power calculation, we begin by assuming a 10×10 unitless study area representing 100 geographic units. The true cluster area Z^* is defined as the circular region centered at grid (6,3) with a radius of 2. All 13 geographic regions that have the center of their region included within the circle are included in Z^*. We defined power as the proportion of simulations that detect a significant ($\alpha = 0.05$) cluster. We define sensitivity as, of those simulations that detect a significant cluster, the proportion that includes at least one region in Z^*. We define accuracy as, of those simulations that find a significant cluster, the number of detected regions that are part of Z^* out of the total number of regions included in the detected cluster. For both sensitivity and accuracy, we only included the highest significant cluster and not all clusters significant at 0.05 level.

For the first power calculation, we generated the dataset assuming that the outcome distribution outside the cluster area is $Y_i \mid Z^{*c} \sim N(0, 1)$ with weight $w_i \mid Z^{*c} = \eta_0$ and within the cluster area is $Y_i \mid Z^* \sim N(c\sqrt{(2)}, \sigma_Z^2)$ with weight $w_i \mid Z^* = \eta_Z$. We vary the magnitude of c, σ_Z, η_0, and η_Z to depict the performance of the test for different scenarios. Note that when $\sigma_Z = 1$, the variance for the difference in means, $\mu_{Z^*} - \mu_{Z^{*c}}$, is $1 + 1 = 2$ with corresponding standard error of $\sqrt{2}$. Therefore c can be interpreted as the number of standard error units of the difference between means within versus outside Z^*.

Table 1 displays the results of the power calculation. As expected, the power increases with increased effect size and becomes above 90% power after $c > 1.3$. The proposed cumulative geographic residual method has lower power compared to the weighted spatial scan statistic, but this is mainly due to the design of the power calculation. If the true cluster is a circle the weighted spatial scan statistic will have higher power, sensitivity, and accuracy than the proposed cumulative residual method since the cumulative geographic residual method detects square clusters and not circles like the spatial scan statistic. However, the advantage of the cumulative geographic residual test is that it can find multiple clusters which the spatial scan statistic cannot. Therefore, the clusters detected are not necessarily squares, but can be a combination of overlapping square areas.

When the weight of the observations within Z^* increases while the weight of observations outside Z^* stays constant the power substantially decreases. This is expected since the residual of the regions within Z^* will go to zero since the estimates of the model will give more weight to the values within Z^*. Similarly when the weight of the observations outside Z^* increases while the weight of the observations inside Z^* stays constant the power decreases. Now the estimates of the model give more weight to the values outside Z^* so the residuals within Z^* may be larger, but the weighted indicator variable $W_i(x_1, x_2 \mid b)$ down weights the observations within Z^* making it less likely to find a significant cluster. The power stays relatively constant with increased weight on all observations. Therefore, there is no clear benefit to use weights unequal to 1 unless there is truly differential weighting throughout the regions of interest.

The next power calculation evaluates the robustness of the cumulative geographic residual test when the outcome is not normal and follows the power calculation presented by Huang et al. (2009). We apply the method assuming a double exponential(DoubleE), logistic, uniform, lognormal, and Poisson distributions. For the DoubleE, logistic, and uniform distributions, we assume a mean 0 and variance 1 outside Z^* and mean $c\sqrt{2}$ and variance 1 inside Z^* and equal weights of 1 for all geographic areas. For the lognormal distribution, we assumed the mean outside Z^* to be 2 instead of 0 since $Y_i > 0$ and the mean inside Z^* to be $2 + c\sqrt{2}$ and variance of

Table 1
Power calculations of the weighted cumulative geographic residual test and weighted spatial scan statistic for different weights and effect size

	Weight		Var		CumGeoRes			Spatial Scan[a]		
Scenario	η_0	η_Z	σ_Z^2	c	Pow	Sens	Accuracy	Pow	Sens	Accuracy
Change	1	1	1	0.5	0.20	0.76	0.45	0.25	0.60	0.50
Effect	1	1	1	1.0	0.68	0.90	0.46	0.88	0.92	0.89
Size	1	1	1	1.1	0.77	0.91	0.46	0.93	0.82	0.84
	1	1	1	1.2	0.86	0.93	0.45	0.98	0.88	0.89
	1	1	1	1.3	0.90	0.95	0.43	0.99	0.91	0.92
	1	1	1	1.4	0.94	0.97	0.41	0.99	0.93	0.94
	1	1	1	1.5	0.96	0.98	0.40	1.00	1.00	1.00
	1	1	1	2.0	1.00	1.00	0.33	1.00	1.00	1.00
	1	1	1	3.0	1.00	1.00	0.30	1.00	1.00	1.00
Change	1	2	1	1.5	1.00	1.00	0.32	1.00	0.99	0.99
Z^{a}	1	10	1	1.5	1.00	1.00	0.33	1.00	0.97	0.95
Weight	1	100	1	1.5	0.17	0.49	0.56	0.34	0.65	0.56
	1	1000	1	1.5	0.03	0.37	0.67	0.12	0.43	0.40
Change	2	1	1	1.5	0.57	0.88	0.47	0.99	0.99	0.98
Z^{*c}	10	1	1	1.5	0.07	0.57	0.32	0.08	0.56	0.44
Weight	100	1	1	1.5	0.05	0.25	0.13	0.05	0.21	0.14
	1000	1	1	1.5	0.05	0.22	0.13	0.05	0.19	0.40
Change	1	1	2	1.5	0.91	0.97	0.42	0.99	0.95	0.98
Z^{a}	1	1	10	1.5	0.55	0.89	0.48	0.76	0.66	0.90
Variance	1	1	100	1.5	0.15	0.78	0.57	0.44	0.41	0.75
	1	1	1000	1.5	0.08	0.77	0.61	0.34	0.35	0.67
Change	10	10	1	1.0	0.70	0.89	0.47	0.88	0.75	0.76
All	100	100	1	1.0	0.66	0.89	0.46	0.85	0.73	0.74
Weights	5000	5000	1	1.0	0.67	0.91	0.46	0.84	0.72	0.73
	10,000	10,000	1	1.0	0.65	0.90	0.47	0.88	0.76	0.77

[a] Results for spatial scan statistic when available were done and reported by Huang et al. (2009). Otherwise, were performed through simulations as outlined for the cumulative geographic residual method.

Model Framework: $Y_i \mid Z_i^* \sim N\left(c\sqrt{2}Z_i^*, \sigma_Z^2 Z_i^* + Z_i^{*c}\right), w_i \mid Z_i^* = \eta_0 Z_i^{*c} + \eta_Z Z_i^*$, and $Z_i^{*c} = 1 - Z_i^*$

$$\text{power} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05)$$

$$\text{sensitivity} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05) I\left(\left\{CritReg_j \cap Z^*\right\} \neq \emptyset\right)$$

$$\text{accuracy} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05)\frac{1}{|\,CritReg_j\,|}\sum_{k=1}^{|\,CritReg_j\,|} I\left(\left\{CritReg_{jk} \in Z^*\right\} \neq \emptyset\right)$$

1 in both regions. For the Poisson distribution, we assumed a mean/variance outside Z^* of 1 and within Z^* as $1 + c\sqrt{2}$.

Table 2 displays the results of the power calculation evaluating the robustness to distribution assumptions. The power is lower for the DoubleE, lognormal, and Poisson distributions and equivalent for the uniform and logistic distributions compared to the normal distribution. Overall the power does not seem largely affected by the distribution assumptions, which is similar to the results of the weighted spatial scan statistic.

3.3 *Power: Area-Level Covariates*

We then conducted a simulation to evaluate the effect of area-level covariate adjustment on a cluster detection. We assumed that $E(Y_i \mid Z_i^*, X_i) = c\sqrt{2}Z_i^* + \beta X_i$ and $\text{Var}(Y_i) = 1$ where $w_i = 1$ for all $i(i = 1, \ldots, 100)$, Z_i^* is an indicator if region i is within Z^* and X_i is a continuous area-level covariate with $E(X_i \mid Z_i^*) = \gamma Z_i^*$ and $\text{Var}(X_i) = 1$, independent of Z_i^*. We ran power calculations varying β (dependence of Y_i on X_i) and γ (dependence of X_i on Z_i^*). We chose X_i to be a continuous covariate because often area-level covariates are continuous such as percent white or median household income. These variables can be standardized to have a variance 1 and a range of means and therefore the simulation is similar to what may be observed in practice. The simulation of each dataset was a two-step process as follows: Step 1: Simulate $X_i \sim N(\gamma Z_i^*, 1)$ and Step 2: Simulate $Y_i \sim N(c\sqrt{2}Z_i^* + \beta X_i, 1/w_i)$ independently for $i = 1, \ldots, 100$.

The first five columns of Table 3 display the results when $c = 1$ and therefore when the outcome, Y_i, and area-level covariate, X_i, both depend on spatial location, Z_i^*. Unadjusted refers to when the residuals, $\hat{e}_i$, used for the spatial

Table 2

Power calculations of the weighted cumulative geographic residual test (CGR) and weighted spatial scan statistic (SS) for different distributions

	Distribution											
	Normal		DoubleE		Logistic		Uniform		Lognormal		Poisson	
c	CGR	SS	CGR	SS	CGR	SS	CGR	SS	CGR	SS	CGR	SS
0.5	0.20	0.25	0.14	0.23	0.22	0.25	0.22	0.30	0.17	0.13	0.21	0.19
1.0	0.68	0.88	0.40	0.86	0.66	0.86	0.65	0.90	0.59	0.71	0.56	0.63
1.5	0.96	1.00	0.72	1.00	0.95	1.00	0.96	1.00	0.88	0.99	0.88	0.95
2.0	1.00	1.00	0.94	1.00	1.00	1.00	1.00	1.00	0.91	1.00	0.98	0.99
3.0	1.00	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.92	1.00	1.00	1.00

[a]Results for Spatial Scan Statistic were done and reported by Huang et al. (2009).
Model Framework: $w_i = 1$.
For Normal, DoubleE, logistic, and uniform: $E(Y_i \mid Z_i^*) = c\sqrt{(2)}$ and $\text{Var}(Y_i \mid Z_i^*) = 1$.
For lognormal: $E(Y_i \mid Z_i^*) = 2 + c\sqrt{2}$ and $\text{Var}(Y_i \mid Z_i^*) = 1$.
For Poisson: $E(Y_i \mid Z_i^*) = 1 + c\sqrt{2}$ and $\text{Var}(Y_i \mid Z_i^*) = 1 + c\sqrt{2}$.

$$\text{power} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05)$$

cluster detection result from a model without adjusting for $X_i(Y_i = \beta_0 + e_i, e_i \overset{ind}{\sim} (0, \sigma^2/w_i))$ and adjusted refers to when the model is adjusted for $X_i(Y_i = \beta_0 + \beta_x X_i, e_i \overset{ind}{\sim} (0, \sigma^2/w_i))$. If unadjusted for the area-level covariate, the power increases if there is a positive $(\gamma > 0, \beta > 0)$ relationship between the area-level covariate and outcome. This is expected since in this case, $E(Y_i \mid Z_i^*) = (c\sqrt{2} + \beta\gamma)Z_i^*$, directly depends on the values of γ and β. If adjusted for X_i, power is decreased with a stronger association between adjusted area-level covariate, X_i, and $Z_i^*(\gamma \to \infty)$. It does not seem to be dependent on the association between X_i and Y_i. The next simulation assessed the key concept of an area-level covariate explaining away the spatial clustering. Specifically, assume the same distributions on X_i and Z_i^* as have been discussed above. If there is no additional relationship between the outcome and location, $c = 0$, other than through the dependence of X_i on Z_i^*, defined by γ, the concept of explaining covariate adjustment is two part. First, a spatial cluster needs to be detected when not adjusting for X_i since $E(Y_i \mid Z_i^*) = \beta\gamma Z_i^*$. Second, after adjusting for X_i, there should be no significant spatial clusters detected $(Z_i^* \perp Y_i \mid X_i \Rightarrow E(Y_i \mid Z_i^*, X_i) = E(Y_i \mid X_i) = \beta X_i)$. Given both of these steps are true indicates that the initial spatial cluster detected in Step 1 is, at least partially, explained by the area-level covariate, X_i.

The next set of simulations evaluates the concept of the spatial cluster being explained by the area-level covariate. Results are shown in Web Table 3, but are briefly summarized here. Our first simulation assessed step 1, the unadjusted analyses, when there is spatial clustering indirectly induced by $X_i(E(Y_i \mid Z_i^*) = \beta\gamma Z_i^*)$. We varied β (-2 to 2 sequenced by 1) and $\gamma = 0, 0.5, 1$. When $\beta \leq 0$ and $\gamma = 0.5$ or 1.0 there is no power to detect spatial clustering and the power is approximately the Type I error, 0.05. When we allowed $\beta > 0$, the power increased as β increased with a maximum power of 0.373 when $\beta = 4$ and $\gamma = 1$. Therefore, the power to detect a significant spatial cluster is relatively low, but it does increase as expected with more positive dependence between Z_i^* and $X_i(\gamma > 0, \gamma \to \infty)$ and stronger positive association between X_i and $Y_i(\beta > 0, \beta \to \infty)$. For step 2, the adjusted analyses, where there is independence between Z_I^* and Y_i conditional on $X_i(Y_i \mid X_i \perp Z_i^*)$, the power should equal the Type I error rate of 0.05. The adjusted simulations showed that the second step of explaining away the significant clustering is working and the power is at approximately 0.05, equal to the Type I error rate. Therefore, the concept of explaining away spatial clustering through covariate adjustment seems to be viable, but a dependence between location and covariate needs to be strong and positive $(\gamma > 0, \gamma \to \infty)$.

3.4 *Type I Error and Power: Individual-Level*

The next set of simulations evaluates the performance of the fixed effect individual-level covariate adjustment approach and compares this method to adjusting for composite area-level covariates. The simulation of each dataset is a two-step process as follows: Step 1: Simulate $X_{ij}^I \sim N(\gamma Z_i^*, 1)$ and Step 2: Simulate $Y_{ij} \sim N(\sqrt{2}cZ_i^* + \beta X_{ij}^I, 1)$ independently for $i = 1, \ldots, 100$ and $j = 1, \ldots, n_i$. We vary the effect of Z_i^* on X_{ij}^I (γ), the effect of X_{ij}^I on Y_{ij} (β), and the number of individuals in an area (n_i). For the composite area-level covariate adjustment, we will use the area-level outcome $\bar{Y}_i = \sum_{j=1}^{n_i} Y_{ij}/n_i$, the area-level covariate $\bar{X}_i = \sum_{j=1}^{n_i} X_{ij}^I/n_i$, and the area-level covariate adjustment method presented in Section 2.1.

The first simulation assessed the scenario when there is no direct relationship between cluster location, Z^*, and individual-level outcome, Y_{ij}(c $= 0$) but a relationship between individual-level covariate, X_{ij}^I, and Z^* and a relationship between X_{ij}^I and individual outcome, Y_{ij}. Under this condition if you adjust for X_{ij}^I using the fixed effect adjustment presented in Section 2.2, the detection of spatial clusters should be held at the Type I error level of 0.05. Further if you use the area-level adjusted approach on the composite area-level outcomes, the type I error should also be 0.05. The Type I error was held at approximately 0.05 (independent

Table 3

Power calculations of the cumulative geographic residual test for adjusted and unadjusted area-level and individual-level covariate analyses when spatial clustering exists independent of covariate and outcome relationship ($c = 1$)

		Area-level Spatial Cluster				Individual-Level Spatial Cluster			
Type of Adjustment:		None	Area-Level			None	Individual-Level	None	Area-Level
	β	Pow	Pow	Sens	Accuracy	Pow	Pow	Pow	Pow
Independence	−2	0.191	0.698	0.884	0.466	0.944	1.000	0.941	1.000
$\gamma = 0$	−1	0.369	0.705	0.899	0.457	1.000	1.000	1.000	1.000
	0	0.710	0.701	0.906	0.459	1.000	1.000	1.000	1.000
	1	0.403	0.646	0.912	0.466	0.999	1.000	0.998	1.000
	2	0.173	0.663	0.908	0.461	0.936	1.000	0.943	1.000
Moderate	−2	0.064	0.668	0.893	0.463	0.148	1.000	0.146	1.000
Dependence	−1	0.185	0.652	0.891	0.470	0.963	1.000	0.957	1.000
$\gamma = 0.5$	0	0.687	0.688	0.890	0.460	1.000	1.000	1.000	1.000
	1	0.649	0.650	0.893	0.464	1.000	1.000	1.000	1.000
	2	0.433	0.673	0.888	0.462	1.000	1.000	1.000	1.000
Strong	−2	0.042	0.578	0.858	0.468	0.053	1.000	0.055	0.898
Dependence	−1	0.074	0.610	0.879	0.465	0.318	1.000	0.325	0.887
$\gamma = 1.0$	0	0.703	0.603	0.871	0.473	1.000	1.000	1.000	0.894
	1	0.863	0.618	0.877	0.458	1.000	1.000	1.000	0.909
	2	0.741	0.591	0.875	0.454	1.000	1.000	1.000	0.882

Model Framework:

Area-level Spatial Cluster:

$X_i \mid Z_i^* \sim N(\gamma Z_i^*, 1), Y_i \mid X_i, Z_i^* \sim N(c\sqrt{2}Z_i^* + \beta X_i, 1)$, and $w_i = 1 (i = 1, \ldots, 100)$

Individual-level Spatial Cluster:

$X_{ij} \mid Z_i^* \sim N(\gamma Z_i^*, 1), Y_{ij} \mid X_{ij}, Z_i^* \sim N(c\sqrt{2}Z_i^* + \beta X_{ij}, 1)(i = 1, \ldots, 100)(j = 1, \ldots, 10)$

Note: Area-level adjustment uses outcome $\bar{Y}_i = \sum_{j=1}^{10} Y_{ij}/10$, area-level covariate $\bar{X}_i = \sum_{j=1}^{10} X_{ij}/10$, and $w_i = 10$.

$$\text{Pow} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05)$$

$$\text{Sens} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05) I\left(\left\{CritReg_j \cap Z^*\right\} \neq \emptyset\right)$$

$$\text{accuracy} = \frac{1}{1000}\sum_{j=1}^{1000} I(P_j < 0.05) \frac{1}{|CritReg_j|} \sum_{k=1}^{|CritReg_j|} I\left(\left\{CritReg_{jk} \in Z^*\right\} \neq \emptyset\right)$$

adjustment: range: 0.034 to 0.060, area-level adjustment: range: 0.036 to 0.070) for a range of β (−2 to 2 sequenced by 1) and γ (0 to 1 sequenced by 0.5). Therefore, the proposed method is appropriate for individual-level covariate adjustment and when only composite area-level covariates and outcomes are available.

We then assessed the power for the proposed fixed effect individual-level and the composite area-level covariate adjustment approaches. Results are displayed in the last four columns of Table 3. The individual-level approach is more powerful then the area-level adjusted approach and the loss of power increased as the dependence between X_{ij}^I and Z_i^* increased.

The next section will assess the performance of the proposed methods evaluating spatial clustering of BMI after adjusting for individual and area-level socioeconomic covariates.

4. Application

We applied the proposed weighted cumulative geographic residual method to an observational study conducted at Group Health Cooperative (GHC), a mixed model health plan and delivery system serving approximately 300,000 members in western Washington State, USA. Females of 18 years and older, who had been continuously enrolled in GHC for at least six months as of July 31, 2006 (with less than a 60-day lapse in membership) and were residents of King County, WA, were eligible for the inclusion. Nursing staff and/or medical assistants obtained weight measurements during routine clinical care and entered these into the electronic medical record (EMR). We extracted any adult height and the most recent measurement of weight from the EMR during the preceding two-year period. Among the 57,499 females in our cohort, 19,451(33.8%) did not have a weight and height measurement in the EMR, 2627(4.6%) did not have valid address information to identify their US census tract location; these females were excluded. We also excluded a census tract with only three women residing in it due to model estimation issues giving a total of 35,418 (61.6%) women from 372 census tracts retained for all subsequent analyses.

Location was only available at the census-tract level so using a weighted spatial cluster detection method was deemed

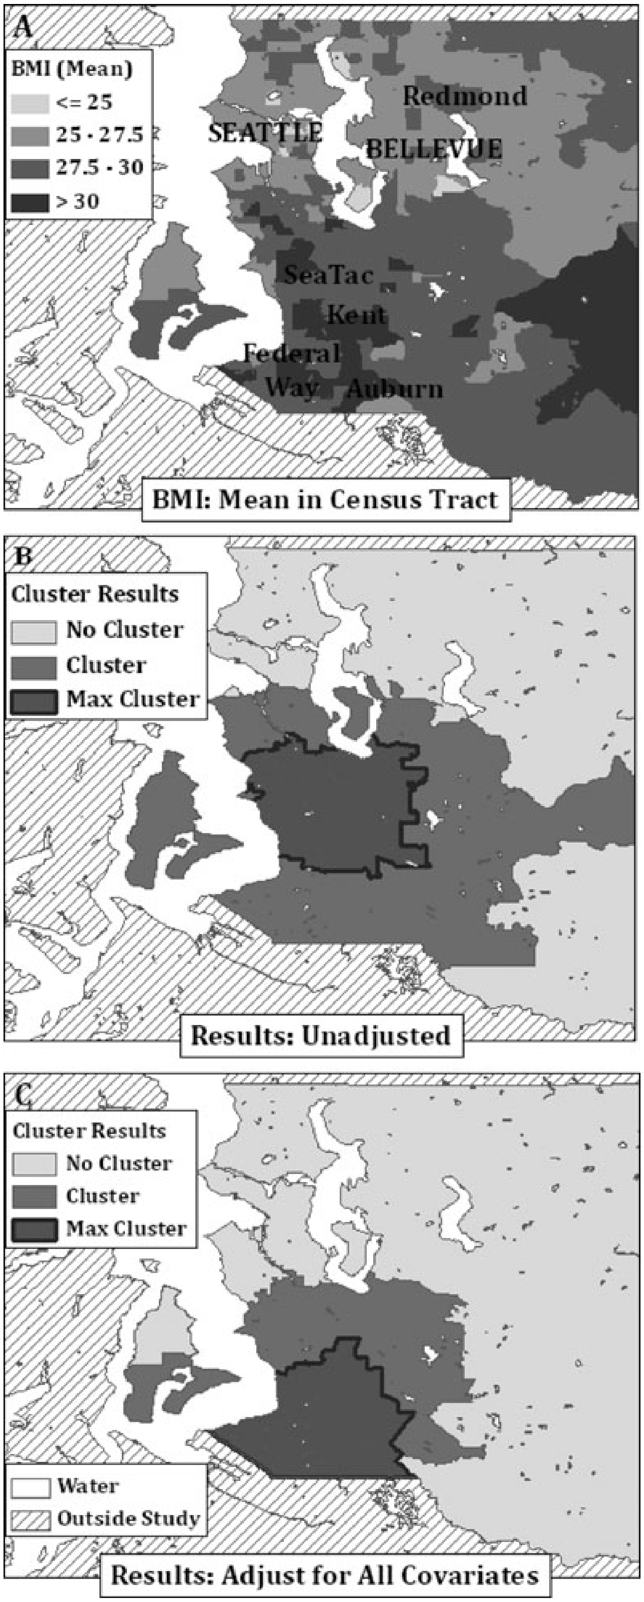

Figure 1. Assessing spatial clustering of female BMI cluster in King County, WA. (A): The raw mean census tract BMI. (B) and (C): The areas with statistically significant spatial clustering from (B) the unadjusted analyses and (C) adjusting for all individual and area-level covariates.

most appropriate for this analysis. The outcome of interest was BMI, which ranged from 15.0 kg/m^2 to 99.30 kg/m^2 with a median of 26.0 kg/m^2 and a mean of 27.5 kg/m^2. Figure 1 Part A shows the variability of the mean BMI within a census tract across the study region. First, it was of interest to assess if there were any elevated regions of BMI in the area and then to assess if the spatial clustering persisted after adjusting for important individual and area-level demographic and socioeconomic status (SES) variables. For our area-level SES variables, we used median household income, percent white race, percent of adult males with a bachelor's degree or higher, and percent below household poverty level, all obtained through the 2000 US census. The only available individual-level covariates were female's age and type of insurance (corporate, Medicare, or Medicaid). Maps of the spatial distribution of the area-level covariate are shown in the Web Figure 1.

The number of observations per census tract varied across the study area and ranged from 12 to 272 with a median of 85 people. For all analyses, we used the fixed effect approach for individual-level covariate adjustment and ranged the half edge length, b, from 1 to 5 miles sequenced by 0.25 miles and ran 1000 simulations of the G_i's.

The first analysis assessed if there existed any spatial clusters when not adjusting for covariates. We found significant spatial clusters in the southern western region of the area as displayed in Figure 1(B). We then ran a series of analyses adjusting for individual-level covariates and then all area-level SES variables. When only adjusting for individual-level covariates the spatial cluster remained statistically significant in a similar location as without adjusting for covariates. When further adjusting for area-level covariates there still existed significant spatial clustering but in a smaller region as shown in Figure 1(C). The location is moved to the most southern part of the study area and displays spatial clustering that is not explained away by our current area-level SES variables.

The final region with persistent clustering of elevated BMI, after adjusting for all individual and area-level covariates, is a region that is less urban compared to the other areas initially detected, including Federal Way, a large suburb of Auburn, which is a mixture of suburban and rural. The region has also a higher mixture of races including more recent immigrants from South Asia. Therefore, there are many possible unmeasured reasons that this area shows persistent clustering such as walkability of the neighborhood (i.e., sidewalks and safety), fast food density, and cultural differences. The link to BMI/obesity and SES in the United States has been well researched in the literature (King et al., 2005; Wang et al., 2007; Dubowitz et al., 2008), particularly for the Caucasian and African-American populations, but how obesity relates to neighborhood infrastructure such as walkability and differences between ethnic groups have been less established (McGinn et al., 2007; Mujahid et al., 2008).

The results from this analysis may not be directly generalizable to Washington State, or even King County, since it only includes an insured population that would have had at least one primary care visit. This is still an important analysis for the health care system to try to understand why BMI is elevated. Finding explanations for why specific areas have elevated BMI could potentially lead to targeted weight reduction programs.

We also applied the unweighted normal spatial scan statistic using the SaTScan$^{\text{TM}}$ software (Kulldorff and Information Management Services, 2006). Software for the weighted normal spatial scan is not yet available. Results are displayed in Web Figure 2. The spatial scan found a much larger region ($\approx 50\%$), which included the entire southern half of the study area. It also found several small statistically significant secondary clusters in the rest of the study area. Some of these regions are based on weights from <20 females. This example shows the issues with not incorporating weights, and the results are not directly comparable with the results presented for the weighted cumulative residual method.

5. Discussion

In this article, we have proposed a robust to distributional assumptions spatial cluster detection method for weighted data that can be applied for a wide variety of outcomes. It was shown to have good power to detect single spatial clusters even when the outcome is not continuous. However, the simulation study did show that if interest is in detecting circular spatial clusters without adjusting for covariates then the weighted spatial scan statistic is more powerful. The new method is able to handle both individual- and aggregate-level covariate adjustment easily and to assess whether covariates are able to explain the detected spatial clusters. There are no previous methods available for weighted outcomes that can adjust for covariates or detect multiple clusters.

The proposed method improved upon the method published by Cook et al. (2007) using cumulative martingale residuals for censored outcome data, as it incorporates weighted outcomes of any type that can be applied to both point and aggregate data. This new method is more general than what was previously proposed since only the first two moments of the outcome are assumed. Further, by flexibly handling differential weighting of observations this method can be applied to a larger array of applications. The method proposed by Cook et al. (2007) may be thought of as a special case of the general methodology proposed in this article.

Throughout most of this article, including the BMI application, we have assumed the weights are known and not estimated. We ran a small simulation study assessing the Type I error when weights were variable and did not find a large effect on the Type I error. However, there may be more of an effect for certain scenarios including an effect on power and extending the method to incorporate extra variability from estimating the weights following the framework presented by Houseman, Coull, and Ryan (2006) may be a potential for future work.

Another potential for the development of new statistical methods would be the extension of this method to handle longitudinal data and to be able to assess if a spatial cluster persists in the same location over time. This would answer questions such as if a region consistently has higher mean BMI compared to other regions. This type of method could be used simultaneously with the proposed method in this article to assess cross-sectional and longitudinal location of spatial clusters over time.

6. Supplementary Material

The Web Appendix, Tables, and Figures referenced in Sections 2, 3, and 4 are available under the Paper Information link at the *Biometrics* website `http://www.biometrics.tibs.org`. R code (R Development Core Team, 2009) implementing the new weighted cumulative geographic residual method is available at the author's website, `http://faculty.washington.edu/acook/software.html`.

Acknowledgements

The authors would like to thank the associate editor and reviewers for their helpful feedback. This work was supported for by NIH Roadmap 5P20RR020774 (AJC) and NIH RO1 CA95747(YL).

References

Cook, A., Gold, D., and Li, Y. (2007). Spatial cluster detection for censored outcome data. *Biometrics* **63,** 540–549.

Dubowitz, T., Heron, M., Bird, C., Lurie, N., Finch, B., Basurto-Davila, R., Hale, L., and Escarce, J. (2008). Neighborhood socioeconomic status and fruit and vegetables intake among whites, blacks, and Mexican Americans in the United States. *American Journal of Clinical Nutrition* **87,** 1883–1891.

Duczmal, L. and Assunção, R. (2004). A simulated annealing strategy for the detection of arbitrarily shaped spatial clusters. *Computational Statistics and Data Analysis* **45,** 269–286.

Heagerty, P. and Kurland, B. (2001). Misspecified maximum likelihood estimates and generalised linear mixed models. *Biometrika* **88,** 973–985.

Houseman, E., Coull, B., and Ryan, L. (2006). A functional-based distribution diagnostic for a linear model with correlated outcomes. *Biometrika* **93,** 911–926.

Huang, L., Kulldorff, M., and Gregorio, D. (2007). A spatial scan statistic for survival data. *Biometrics* **63,** 109–118.

Huang, L., Tiwari, R., Kulldorff, M., Zou, J., and Feuer, E. (2009). Weighted normal spatial scan statistic for heterogenous population data. To appear in *Journal of the American Statistical Association.*

King, W., Belle, S., Brach, J., Simkin-Silverman, L., Soska, T., and Kriska, A. (2005). Objective measures of neighborhood environment and physical activity in older women. *American Journal of Preventive Medicine* **28,** 461–469.

Kulldorff, M. (1997). A spatial scan statistic. *Communications in Statistics* **26,** 1481–1496.

Kulldorff, M., Huang, L., Pickle, L., and Duczmal, L. (2006). An elliptic spatial scan statistic. *Statistics in Medicine* **25,** 3929–3943.

Kulldorff, M. and Information Management Services, I. (2006). *SaTScan*$^{\text{TM}}$ *v7.0: Software for the spatial and space-time scan statistics.* Available at `http://www.satscan.org/`.

Litiere, S., Alonso, A., and Molenberghs, G. (2008). The impact of a misspecified random-effects distribution on the estimation and the performance of inferential procedures in generalized linear mixed models. *Statistics in Medicine* **27,** 3125–3144.

McGinn, A., Evenson, K., Herring, A., Huston, S., and Rodriguez, D. (2007). Exploring associations between physical activity and perceived and objective measures of the built environment. *Journal of Urban Health* **84,** 162–184.

Mujahid, M., Roux, A., Shen, M., Gowda, D., Sanchez, B., Shea, S., Jacobs, D. J., and Jackson, A. (2008). Relation between neighborhood environments and obesity in the multi-ethnic study of atherosclerosis. *American Journal of Epidemiology* **167,** 1349–1357.

Pan, Z. and Lin, D. (2005). Goodness-of-fit methods for generalized linear mixed models. *Biometrics* **61,** 1000–1009.

Patil, G. and Taillie, C. (2004). Upper level set scan statistic for detecting arbitrarily shaped hotspots. *Environment and Ecology Statistics* **11,** 183–197.

R Development Core Team (2009). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.

Su, J. and Wei, L. (1991). A lack-of-fit test for the mean function in a generalized linear model. *Journal of the American Statistical Association* **86,** 420–426.

Tango, T. (2000). A test for spatial disease clustering adjusted for multiple testing. *Statistics in Medicine* **19,** 191–204.

Tango, T. and Takahashi, K. (2005). A flexibly shaped spatial scan statistic for detecting clusters. *International Journal of Health Geographics* **4,** 11.

Turnball, G., Iwano, E., Burnett, W., Howe, H., and Clark, L. (1990). Monitoring for clusters of disease: Application to leukemia incidence in upstate New York. *American Journal of Epidemiology* **132,** 136–143.

Wakefield, J. (2003). Sensitivity analyses for ecological regression. *Biometrics* **59,** 9–17.

Wang, M., Kim, S., Gonzalez, A., MacLeod, K., and Winkleby, M. (2007). Socioeconomic and food-related physical characteristics of the neighborhood environment are associated with body mass index. *Journal of Epidemiology and Community Health* **61,** 491–498.

Received August 2008. *Revised June* 2009.
Accepted June 2009.