STATISTICS IN MEDICINE
Statist. Med. 2000; **19**:2399–2408

Evaluation of spatial filters to create smoothed maps of health data

Thomas O. Talbot[1,*,†], Martin Kulldorff[2], Steven P. Forand[1] and Valerie B. Haley[1]

[1]*Geographic Research and Analysis Section, Bureau of Environmental and Occupational Epidemiology, New York State Department of Health, 547 River Street, Room 200, Troy, NY 12180-2216, U.S.A.*
[2]*Division of Biostatistics, Department of Community Medicine and Health Care, University of Connecticut School of Medicine, Farmington, CT 06030-6205, U.S.A.*

SUMMARY

Spatial filters have been used as an easy and intuitive way to create smoothed disease maps. Birth weight data from New York State for 1994 and 1995 are used to compare the traditional filter type of fixed geographical size with a filter size of constant or nearly constant population size. The latter are more appropriate for mapping disease in geographic areas with widely varying population density, such as New York State. Issues such as the choice of population size for the filter, the scale of smoothing, the ability to detect true spatial variation and the ability to smooth over random spatial noise are evaluated and discussed. Copyright © 2000 John Wiley & Sons, Ltd.

1. INTRODUCTION

With the development of geographic information systems (GIS) and software for automated geocoding of addresses, local, state and federal health agencies are increasingly mapping incidence and mortality rates, disease prevalence and other health related outcomes, in order to look at their geographical variation. This is important, as part of geographical disease surveillance, to detect areas with high or low rates of adverse health outcomes or areas where the population may not be receiving adequate health care.

Maps are commonly produced by shading counties according to the rates of disease within each county. However, these maps are not useful for identifying high or low rate areas that are either localized within part of a county, or that cross county boundaries. If health outcomes are instead mapped for smaller areas, such as United States ZIP code areas, census tracts or census block groups, the estimated rates will be increasingly unstable as they are based on fewer and fewer individuals.

One can display data collected at smaller geographic areas and still maintain the stability of the estimated rates by constructing a smoothed map. One way to do this is to use Bayesian or

* Correspondence to: Thomas O. Talbot, Geographic Research and Analysis Section, Bureau of Environmental and Occupational Epidemiology, New York State Department of Health, 547 River Street, Room 200, Troy, NY 12180-2216, U.S.A.
† E-mail: tot01@health.state.ny.us

empirical Bayes methods, calculating the estimated rates for each smaller area by incorporating information about the observed data from neighbouring areas together with priors concerning the spatial variation of the rates [1–3].

Another approach is to use what has been called a spatial filter [4] or ratio smoother [5]. This technique can be applied to individual point data, as well as to data aggregated into small census areas. In its simplest form, the estimated rate at a particular location, or grid point, is defined as the observed rate within a fixed distance from the grid point. The circles of neighbouring grid points are set to overlap to allow neighbouring grid points to share observations. Once estimated rates are assigned to each grid point, contouring software can be used to create isarithmic maps in which regions with a constant range of values can be recognized. This enables one to create a continuous smoothed map of the data. Rushton and Lolonis [4] proposed this technique to map birth defect rates in Des Moines, Iowa. Kafadar [5] proposed a general class of such smoothers, where observations closer to the centre may be weighted more heavily than those further away, in order to map prostate cancer mortality in the United States.

Building upon the ideas of Rushton and Lolonis, and Kafadar, we propose a modified spatial filter for creating smoothed disease maps, where the spatial filter is defined in terms of constant or near constant population size rather than constant geographic size. This means that the circles are larger in the rural areas compared to urban areas. The distinction between fixed and variable geographic size filters corresponds to the distinction between fixed window widths and adaptive window widths in the density estimation literature [6–8].

The new methods with variable window size are evaluated and compared with the fixed geographical size approach, using low birth weight data from New York State for 1994 and 1995. It is shown that spatial filter techniques that use fixed geographical sizes are of limited use in areas with heterogeneous population densities, such as New York State. The filter size must be set very large in order to obtain sufficiently stable rates in the rural areas, but setting such a large filter size will mask differences in rates in urban areas. In contrast, a spatial filter based on a fixed or nearly fixed population size enables one to retain adequate resolution in the urban areas while at the same time providing stable rates in the rural areas.

A smoothed disease map is primarily a descriptive tool, showing areas with high and low rates after removing some of the random spatial noise. To complement the smoothed map of rates, Rushton and Lolonis [4] proposed a smoothed probability map, where the *p*-value for each individual circle is calculated and mapped in the same fashion as the rates. To evaluate the ability of the smoothed maps of rates to both detect true spatial variation and to smooth over random spatial noise, we compared them with a smoothed probability map where the *p*-values are adjusted for multiple testing (Talbot *et al*., manuscript in preparation). Moreover, we compared them with statistically significant clusters obtained using the spatial scan statistic [9,10], to determine whether the smoothed disease maps retained or smoothed out these non-random clusters.

2. MATERIALS AND METHODS

2.1. Low birth weights in New York State

In order to evaluate the spatial filter smoothing techniques, we used birth weight data from New York State for the years 1994–1995. All multiple birth deliveries, and births in out-of-state hospitals, were excluded from the analyses. Births with a report weight of less than 500 g or more than 7500 g were also excluded since such low or high weights are likely to be due to reporting errors. These errors in the birth weight reporting led to the exclusion of 1005 births or 0.19 per cent of the total births used in the analysis. Most birth records contained the ZIP code of the mother's residence at birth. A small portion of the birth records, 0.25 per cent, did not, and these were also excluded from the analyses.

Birth weights of less than 2500 g were defined as a low birth weight (LBW). In the data set, there were a total of 525 143 births, including 32 592 LBWs, which is 6.2 per cent of the total births. The variable of interest is the proportion of LBWs, and whether this proportion varies geographically across the state.

Population-weighted centroids were calculated for each ZIP code using 1990 census data at the lowest level of census geography available, the census block [11]. Each of the approximately 275 000 census blocks in New York State has latitude, longitude, and a 100 per cent count of persons. The blocks were geocoded by ZIP code using 1995 ZIP code boundaries obtained from Geographic Data Technologies [12]. For each ZIP code, the latitude was calculated as the mean of the block latitudes, weighted by the number of persons in each block. The ZIP code longitude was calculated similarly.

For each of 1601 area ZIP codes in New York State, we summed the number of LBWs as well as the total number of births. All births in a particular ZIP code area were then assigned to the geographical co-ordinates of the population-weighted ZIP code centroid.

2.2. *Spatial filter smoothing methods*

We examined three smoothing techniques based on different spatial filters, all using a circular window. For all methods we used a 1 km grid for the circle centroids. The spacing of the grid points is less important than the size of the filter as long as the distance between grid points is much less than the distance between ZIP code centroids. The grid we used contained approximately 129000 points, which covered the state along with a 1 km buffer around the state.

The three methods evaluated were as follows:

(a) *A spatial filter with constant geographic size* [4,5]. This method captures all LBWs and births within a fixed geographical radius of the grid point.

The filter size was set to 5, 10 and 57.7 km respectively. The latter is the smallest radius that will capture at least 250 births at each grid point. This is rather large due to the sparsely populated areas in the Adirondack region, where as many as 23 ZIP codes were needed in order to capture 250 births.

(b) *A spatial filter with nearly constant population size*. This method sets, *a priori*, a minimum number of births to capture at each grid point. At the same time, the minimum radius is set to 0.75 km in order to ensure that no ZIP codes are missed due to the 1 km spacing of the grid points. Where more than one centroid is within the 0.75 km radius, these ZIP codes are treated as one.

The nearest ZIP code centroid to the grid point is located. If the number of births is less than the minimum number, then the next nearest ZIP code centroid is located and the number of births are added to those from the previous ZIP codes. This process is continued until the total number of births captured is greater than or equal to the minimum number of births. At this point, the total number of births and LBWs captured in the selected ZIP codes are assigned to the grid point. The number of births captured will exceed the predetermined minimum if the last ZIP code area captured has more births than needed to reach the minimum population size.

This method was tested using five population based filter sizes, which were set at minimum number of 100, 250, 500, 1000 and 5000 births, respectively. For example, if 250 births were captured, the expected number of LBWs is 15.5 at the grid point if the rate within the circle is the same as the overall New York State rate of 6.2 LBWs per 100 live births.

(c) *A spatial filter with constant population size*. This method is the same as the previous method, with one exception. If the number of births captured exceeds the preset population size when a new ZIP code area is captured, then only a percentage of the births in that area are included, so as to obtain the exact number of desired births. The number of LBWs included from that last ZIP code area is determined using the same percentage. This process of including only a proportion of LBWs and births in the last area was also used by Turnbull *et al.* as part of a statistical cluster detection test, called the Cluster Evaluation Permutation Procedure [13].

Constant filter sizes of 100, 250, 500, 1000 and 5000 births were examined using this method.

For all methods, the standard incidence ratio was calculated for each grid point. The study area was covered by a fixed number of evenly spaced grid points I. Let n_i be the number of births captured at grid point i and let N be the total number of births in the state. Likewise let c_i be the number of cases of LBW captured at grid point i while C is the total number of LBW cases in the state. The standardized incidence ratio (SIR) at grid point i is then

$$\mathrm{SIR}_i = (c_i/n_i)/(C/N)$$

A number of examples from our data are presented in Table I to illustrate how methods (b) and (c) differ. The first example shows that if the first ZIP code includes more than 250 people, then the two methods produce identical rates since they both rely on data from only one ZIP code. In the second example, several ZIP codes were captured and the total number of births is exactly 250 so the methods produce the same rate. The next two examples are extreme cases where there are differences in our data. One of these shows that the first ZIP codes captured by this grid point had 230 births with a LBW rate of 2.6 per cent. The next ZIP code captured had 510 births with a LBW rate of 7.3. In method (b), the LBW rate in the last ZIP captured has a greater influence on the rate for this grid point since it contains 69 per cent of the births captured, giving a rate of 5.8 per cent. Method (c) only uses information on 20 births from the last ZIP code captured, which is 8 per cent of the births used to calculate the rate for that grid point. The next example in Table I is similar, except that the direction of the difference is reversed. The difference in rates was found to be less than 1 per cent for 95 per cent of the grid points while less than 1 per cent of the grid points had differences of more than two percentage points in LBW rates.

The three different spatial filter approaches took approximately 20 minutes each using a Pentium II 450 MHz computer to calculate SIRs for all grid points. Using the SIR values at these grid points, contour-mapping software [14,15] was used to create the isarithmic maps. Because the data points were very closely and regularly spaced, SIR values at all grid points fell within the correct contours.

The smoothed maps of LBW rates were compared with the results of analysing the same data using the spatial scan statistic [9,10], in order to determine whether areas with statistically

Plate 1. Map of New York State showing the locations of the major cities and county boundaries. The New York City area is shown in the inset.

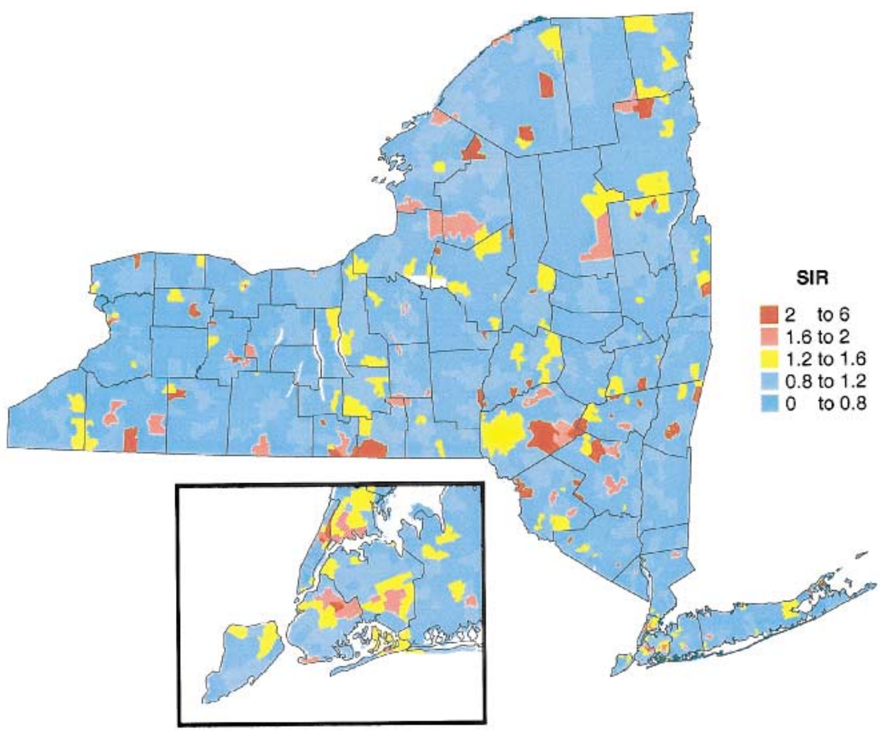

Plate 2. Low birth weight standard incidence ratios by ZIP code (unsmoothed).

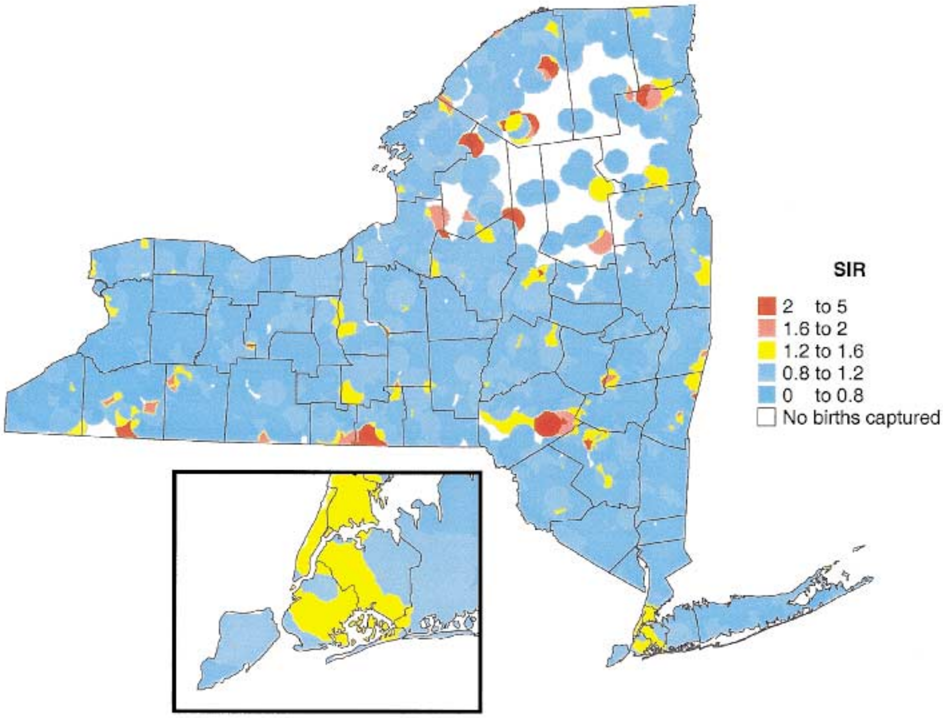

Plate 3. Low birth weight standard incidence ratios using a fixed filter size of 10.0 km at each grid point.

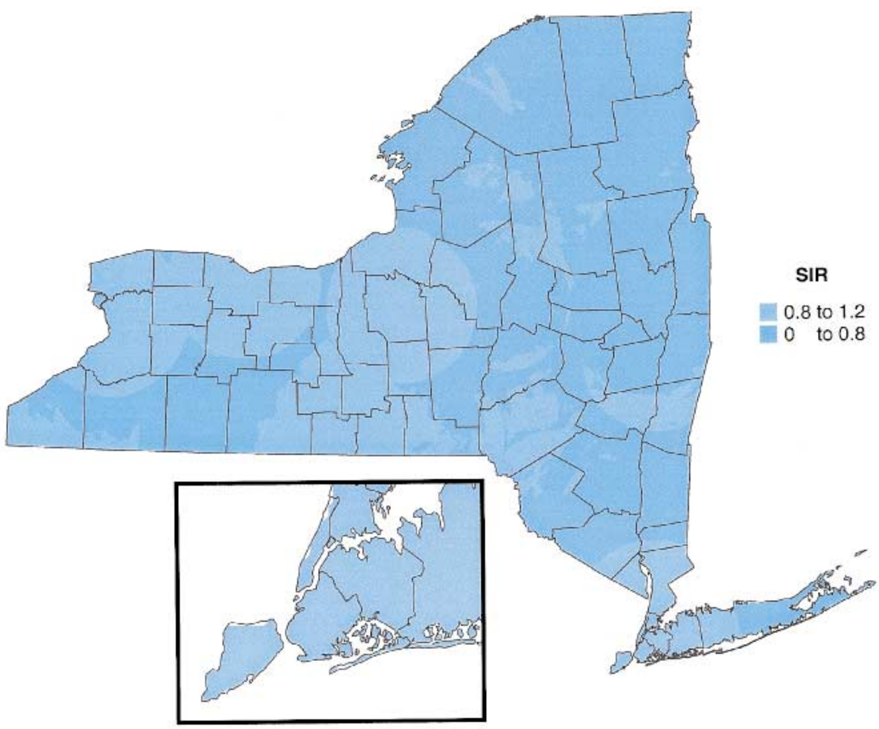

Plate 4. Low birth weight standard incidence ratios using a fixed filter size of 57.7 km at each grid point.

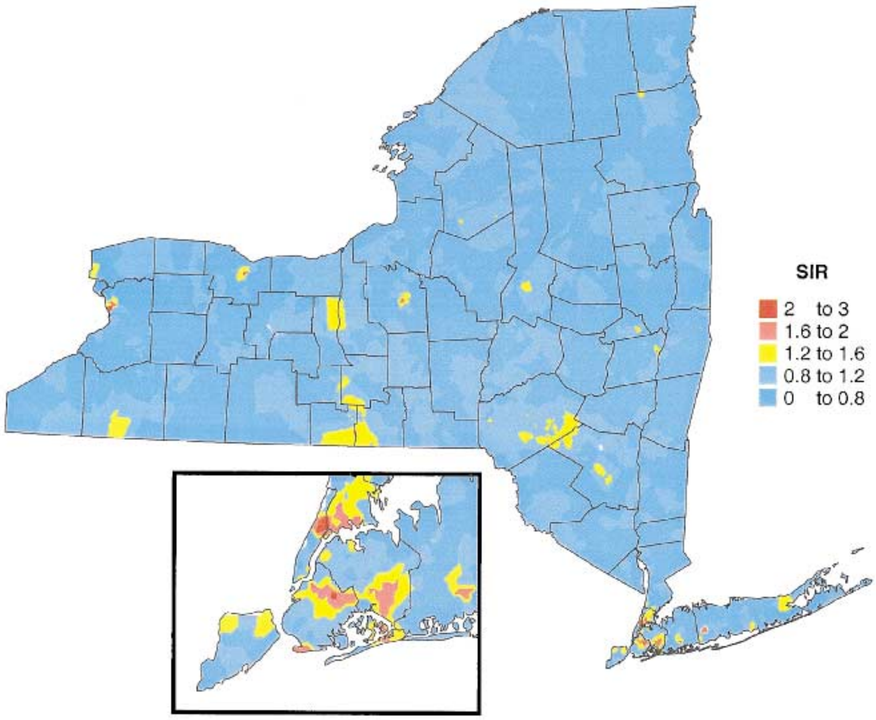

Plate 5. Low birth weight standard incidence ratios using a variable filter size set to capture at least 250 births at each grid point.

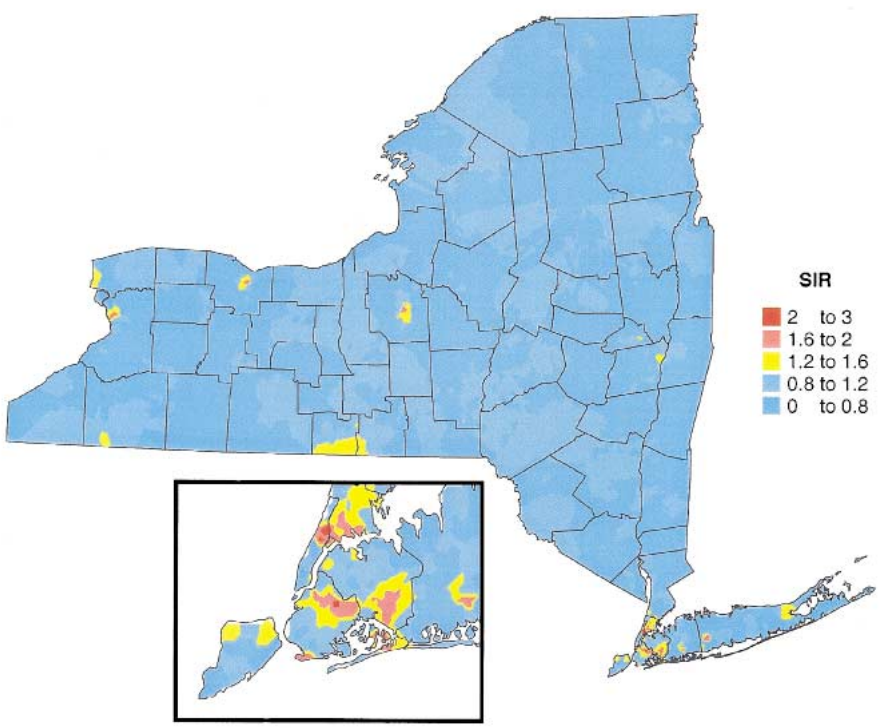

Plate 6. Low birth weight standard incidence ratios using a variable filter size set to capture at least 500 births at each grid point.

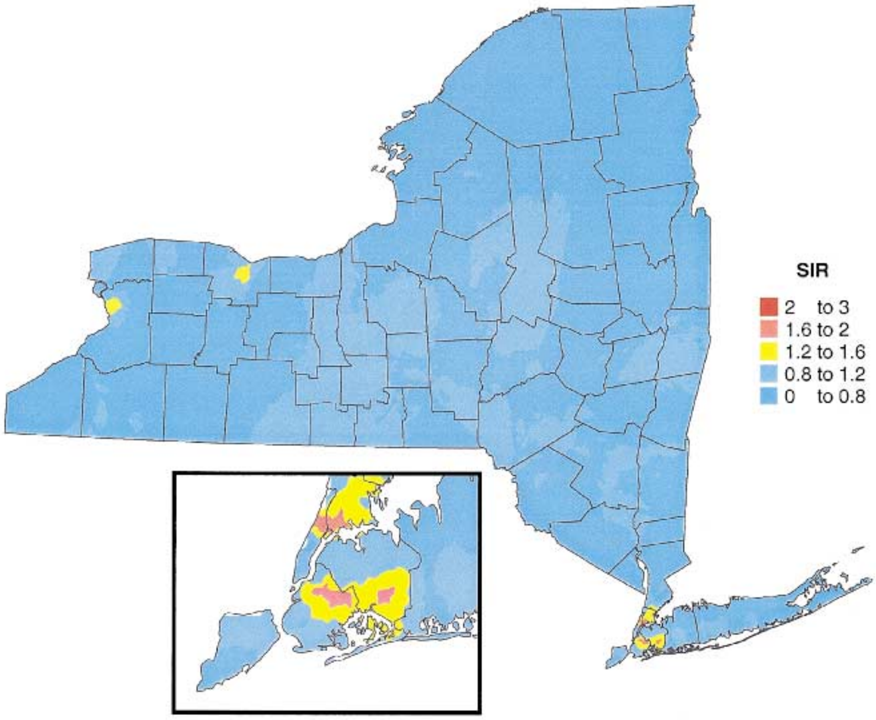

Plate 7. Low birth weight standard incidence ratios using a variable filter size set to capture at least 5000 births at each grid point.

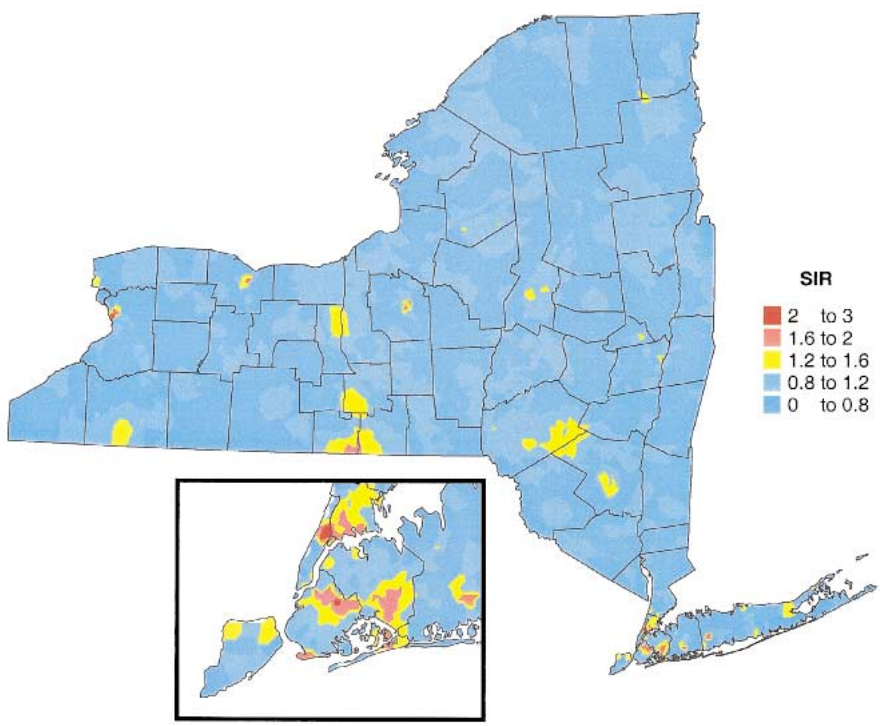

Plate 8. Low birth weight standard incidence ratios using a variable filter size set to capture exactly 250 births at each grid point.

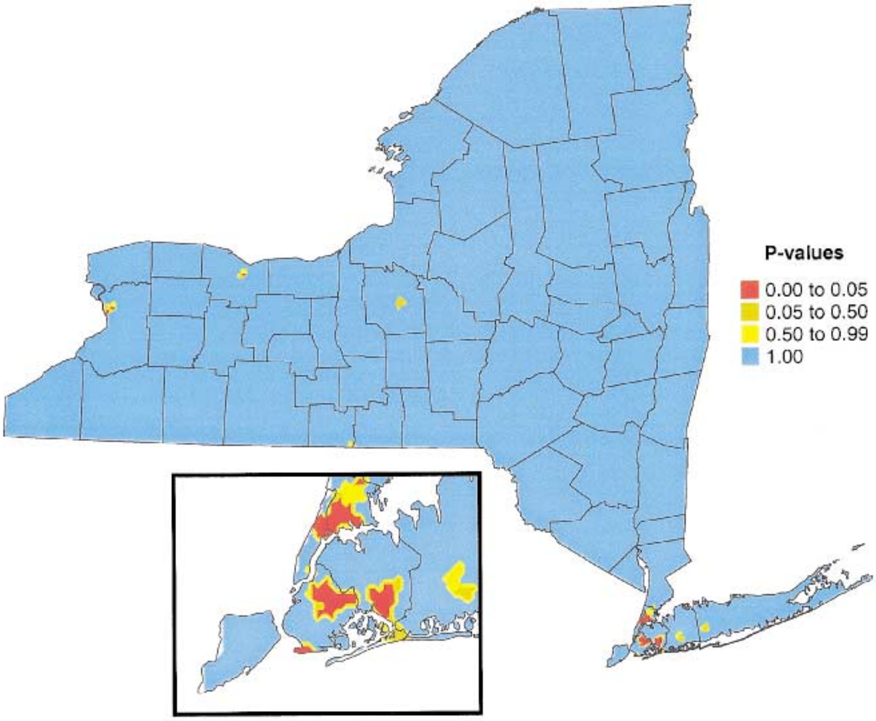

Plate 9. Smoothed probability map of low birth weight using the spatial scan statistic and a filter size set to capture at least 250 births.

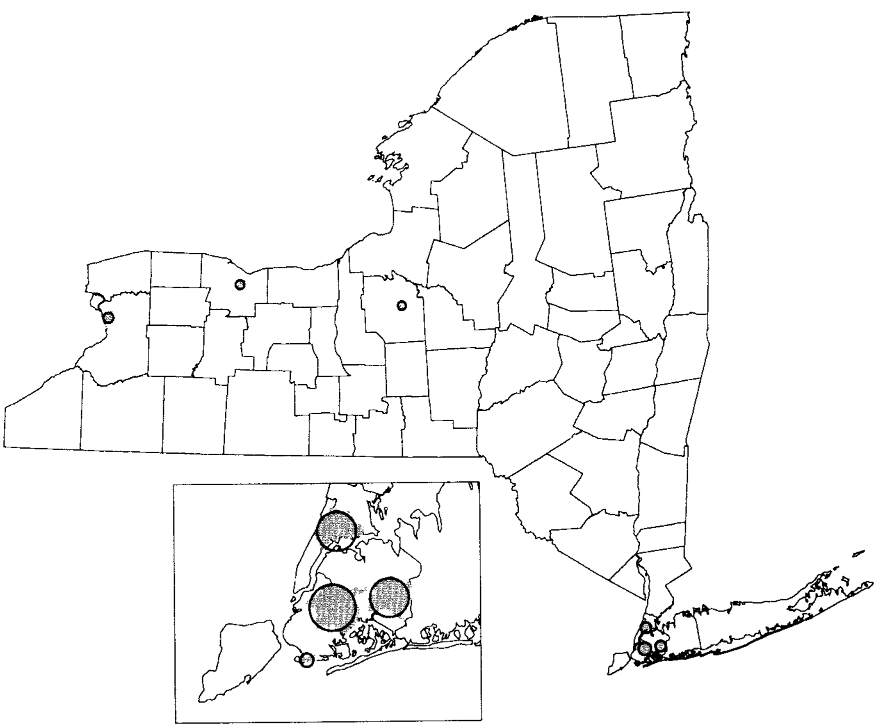

Plate 10. Most likely clusters of low birth weight deliveries using a one-sided spatial scan statistic. Only elevated clusters with significantly high rates are shown ($p<0.05$). Restrictions: no cluster can contain more than 10 per cent of births.

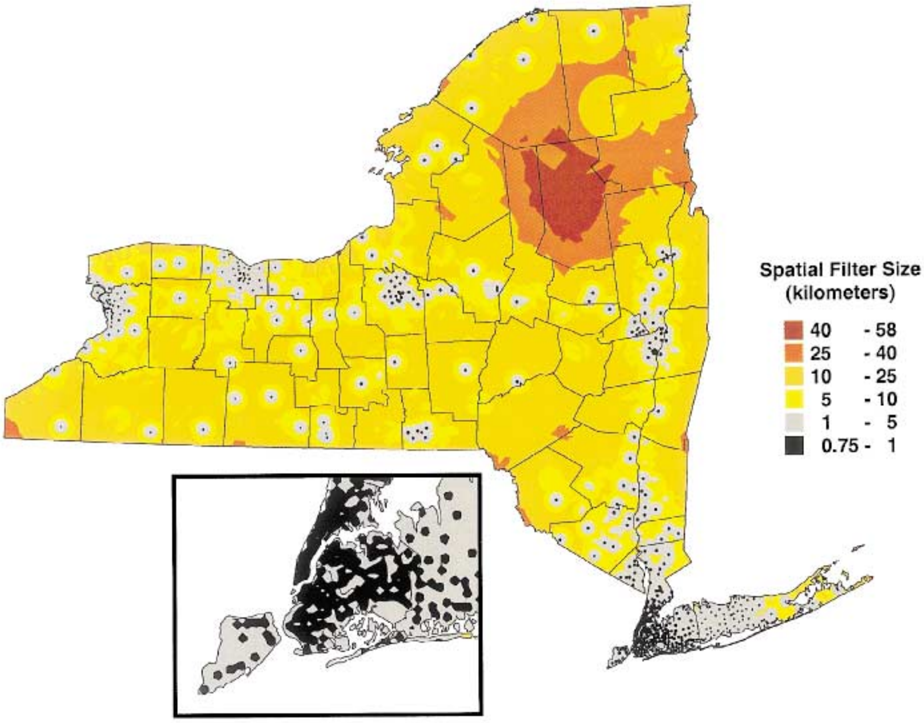

Plate 11. Spatial filter size needed for each grid point to capture at least 250 births.

Table I. Examples of differences between nearly constant population and constant population spatial filters methods.

Examples	Total before last ZIP code captured			Last ZIP code captured			Constant method	Nearly constant method	Rate difference
	Births	LBW	LBW Rate %	Births	LBW	LBW Rate %	LBW Rate %	LBW Rate %	
First ZIP code births > = 250	1023	102	10.0	1023	102	10.0	10.0	10.0	0
Several ZIP codes (births equal exactly 250)	190	6	3.2	60	2	3.3	3.2	3.2	0
Last ZIP code captured has higher rate	230	6	2.6	510	37	7.3	3.0	5.8	2.8
Last ZIP code has lower rate	245	17	6.9	599	19	3.2	6.9	4.3	− 2.6

significant rates were retained or smoothed out. The spatial scan statistic is a test for spatial randomness to determine the location of any area of the state with an elevated rate of LBW that is statistically significant after adjusting for multiple testing inherent in the many possible locations and sizes of the area. The test uses an infinite number of overlapping circles of many different sizes and locations. The maximum size was restricted to 10 per cent of all births in New York State. For each circle, the likelihood to find the observed number of LBWs inside and outside the circle, respectively, is calculated, and the circle with the highest likelihood is the most likely cluster. This is the circle with the rate that is least likely to have occurred by chance. There may also exist secondary clusters of interest. A p-value can be determined for each circle through Monte Carlo hypothesis testing. Since the maximum likelihood is calculated over all circles, the test accounts for multiple testing. Calculations were done using the SaTScan software [16]. In addition to obtaining the most likely clusters, the spatial scan statistic was also used to create a smoothed probability map adjusted for multiple testing (Talbot *et al*., manuscript in preparation). For each grid point, the p-value was obtained for the circle with a population size of 250 births centred at that grid point. These p-values, adjusted for multiple testing, were then mapped using contouring software.

3. RESULTS

As a geographical reference, Plate 1 shows the county boundaries in New York State, together with the five main urban areas: New York City; Albany; Syracuse; Rochester, and Buffalo. The population density varies greatly in New York State, from the very sparsely populated Adirondack mountain region in the northeastern part of the state, to the very densely populated New York City in the south. For example, Manhattan has more than 20 000 people per square kilometre while several rural towns in northern New York State have less than 0.5 persons per square kilometre. This makes it an ideal geographical region for evaluating how spatial filter smoothing techniques perform when there is heterogeneity in the population density.

An unsmoothed map is provided in Plate 2, showing the SIRs for each individual ZIP code area. There are several areas of high incidence scattered around the map with no readily apparent pattern. Since the rates are unsmoothed, the SIRs are very unstable for all ZIP code areas except those with the very largest population. If there is any true spatial variation of importance, it is hidden, and impossible to distinguish from the random spatial noise.

Plate 3 shows the smoothed map using a fixed geographic filter size of 10 km. The white shading on the map corresponds to areas that could not be estimated because there were no ZIP code centroids within 10 km of the grid points. In the most rural areas of the state, this map is very similar to the unsmoothed map in Plate 2, with equally unstable rates. In the urban and semi-urban areas though, a considerable amount of smoothing has been achieved, providing more stable rates than in Plate 2. In the most densely populated areas such as New York City, the rates may be overly smoothed, masking any real differences in rates within the city.

To overcome the problem of unstable estimates in rural areas, it is necessary to increase the size of the spatial filter. The map in Plate 4 was generated using a fixed filter size of 57.7 km in order to reduce random noise for all grid points. Very little spatial variation in rates remains. The filter size was set so high to remove random noise in the rural areas, that it masks potentially true differences which may exist in the urban areas. Overall, this map does not convey much useful information in determining the geographical pattern of LBW rates.

The smoothed SIR map using the spatial filter with nearly constant population size of 250 births is shown in Plate 5. Many of the areas that had a high SIR in the unsmoothed ZIP code map (Plate 2) are no longer seen. This indicates that these elevations were perhaps due to random fluctuations of rates because of small numbers of total births within a ZIP code area. Some elevated areas do exist though, most notably in parts of New York City, Rochester, Buffalo and Syracuse, where the LBW rate is in some cases twice the state average.

To evaluate the effect of changes in filter size, we created maps with different population filter sizes. Maps produced using nearly equal filter sizes of 500 and 5000 are shown in Plates 6 and 7. Unusually high rates are moderated as the population filter increases. This is particularly noticeable in the rural areas where the unsmoothed ZIP code rates were unstable due to the small numbers of births. The high rates, which remain as the filter size increases to 250, 500 and 1000 births, respectively, indicate that these differences are less likely due to chance. Once the filter size becomes too large, this technique loses the ability to detect elevated rates in all but the most densely populated areas of New York City. This is apparent in Plate 7 in which the filter is set to capture a minimum of 5000 births.

The smoothed map using method (c), where exactly 250 births were used to calculate the SIR for each grid point, is shown in Plate 8. This map is almost identical to the one using method (b) with 250 births (Plate 5). The two methods were also similar for the other filter sizes. The same general areas appear high on both maps, but there are some minor differences.

Methods (b) and (c) were also compared, by plotting the difference in the LBW rates across the state (map not shown). There were very few differences in the very densely populated areas of New York City and in the very rural areas of the Catskill and Adirondack Mountain regions. This is because almost all the New York City ZIP code areas have over 250 births, in which case the rates are the same for both methods. In the very rural areas there are very few births in each ZIP code. This means that the last ZIP code captured has a small impact on the overall rate within the circle, and that the total number of births is unlikely to diverge too much from the minimum specified. Small differences occurred in the suburban and semi-rural areas of upstate New York.

The smoothed probability map using the spatial scan statistic is shown in Plate 9. The maps are shaded by the *p*-values calculated for the grid points. The most likely clusters, irrespective of size, are shown in Table II and Plate 10. The general areas that have a rate that is statistically significant adjusting for the multiple testing, also show up as high rate areas on the smoothed maps of LBW rates with a window of constant and nearly constant population size (Plates 5, 6 and 8). Hence, these maps do not smooth over the areas with statistically significant rates except when the filters are set too large. We see in Plate 7 that by using a large filter size of 5000 births, the small statistically significant areas in Coney Island and Syracuse are smoothed over. When using a fixed geographic filter size though, areas with a statistically significant elevated rate are sometimes smoothed out (Plate 4), and when they are not, there are other areas with non-significantly elevated rates that dominate the map (Plate 3).

4. DISCUSSION

Most areas of the state which were identified as having a higher than expected number of LBWs using the spatial filter and the spatial scan statistic techniques are located in the areas which have low household incomes and large proportions of African-Americans. This is consistent with the existing literature [17], which has shown an association between LBW, low socio-economic

Table II. Results of spatial scan statistic test for most likely low birth weight clusters.

Geographic area	LBW births	Total births	SIR	*P*-value
South Bronx & Harlem	4061	40906	1.6	< 0.001
Bedford-Stuyvesant & Crown Heights (Brooklyn)	4283	44882	1.5	< 0.001
Jamaica (Queens)	1355	14675	1.5	< 0.001
Buffalo	668	7057	1.5	< 0.001
Rochester	460	4829	1.5	< 0.001
Coney Island (Brooklyn)	144	1408	1.6	< 0.001
Syracuse	168	1807	1.5	0.005

status and race. Such a relationship can be looked at more formally, using for example the Kafadar–Tukey urbanization index [18], but that is beyond the scope of this methodological paper.

Spatial filter methods based on circles with a fixed or nearly fixed population size are more appropriate for mapping disease in geographic areas with widely varying population density, than methods which use a fixed geographical circle size. The major advantage is that the random noise in rural areas is smoothed to the same extent as in urban areas, while at the same time enhancing the geographical resolution of the urban area estimates. A second advantage is that all estimates throughout the map are approximately equally reliable. This is because the variances of the estimates depend on the population size at risk.

The major disadvantage of the fixed population size technique is that the map does not have the same resolution throughout. In our example the grid point rates are calculated based on populations within circles as small as 0.75 km in urban areas while in the rural areas the filter can exceed 50 km. A map showing a high rate of disease in a rural area may give the mistaken impression that the elevation is local in nature, when in fact the rate is based on a much larger area. This requires some care when interpreting the maps. A map which shades each grid by the geographical circle size needed to capture the desired population provides added information about the geographic resolution of the smoothed map. Plate 11 provides an example of such a map using a filter size set at 250 births.

In areas with a more homogeneous population density, it may be better to use a fixed geographical filter size. The mapping of birth defects in Des Moines is a case in point. Rushton and Lolonis [4] could easily describe their map by saying that the birth defect prevalence at any grid point was based on the births located within a 0.4-mile radius. The reader easily understands this. If the population density is constant throughout the map, then the methods are identical, whether one uses a fixed population or a fixed geographical area.

While the two methods with circle sizes based on population are more appropriate than methods with fixed circle sizes in areas with heterogeneous population densities, the question remains of whether it is better to use the constant or the nearly constant population size filter. In terms of their estimates, the differences are minimal, and, for example, there is no way of claiming that the map in Plate 5 is more appropriate than the map in Plate 8, or vice versa. The nearly constant method is easier to calculate, and would be easier to explain to a lay audience. For the constant method though, the variances would be more equal across the map. All of these criteria are too marginal though for us to clearly advocate one method over the other.

Whichever technique is used, there remains the issue of choosing the appropriate population filter size. There is no intrinsic size that can be used to smooth all maps. The choice depends to some extent on the aetiology of the disease. If the filter is too large, it may remove true variation, which could point to aetiological associations, while if it is too small it will not remove all the random noise in the data. Since smoothed maps are primarily a descriptive tool, it is useful to simultaneously provide several maps with different amounts of smoothing in order to be able to see different types of spatial patterns (Plates 5–7).

The analysis of the geographic distribution of disease using smoothed maps should always be complemented with a statistical test for spatial randomness to determine whether the observed pattern is likely due to chance or not, and to locate elevated areas that are statistically significant. We applied one such test, the spatial scan statistic, and were able to determine which of the high areas observed by the spatial smoothing technique were likely chance occurrences, and which were statistically significant after adjusting for the multiple testing of many possible cluster locations and sizes.

The use of a test for spatial randomness does not mean that the smoothed maps of rates should be discarded, since the general pattern of the disease rate is of interest for generating hypotheses about the cause of the spatially variable rates. To provide an example of how the two methods can work in tandem, note that the statistically significant areas are in poor urban inner city areas in New York City, Syracuse, Rochester and Buffalo. The fact that the smoothed map also shows an elevation in the inner city of Albany might then have a natural explanation, even though this particular area was not statistically significant based on the spatial scan statistic. While the elevated rate in the Albany area is not statistically significant it does warrant further investigation from a public health perspective since it shares socio-demographic characteristics similar to the other areas of high incidence.

The smoothed maps all use a circular spatial filter to calculate LBW rates. However, we do not necessarily get circular areas when we shade the grid points which have elevated LBW rates or low *p*-values. For example, a long and narrow shaped area could comprise a number of grid points which represent a series of centroids of overlapping clusters which may follow a river or transportation corridor. It is important to realize that the smoothed maps do not show the areas with elevated rates *per se*, but the location of the centroids for which the surrounding circles of a specified size contain an elevated rate.

The exact boundaries of different areas with different rates cannot be determined from the smoothed maps, or by other statistical methods. These techniques only indicate the general location of areas with higher and lower rates. Moreover, these techniques do not determine the cause of the geographical variability. Further analysis is often needed through detailed geocoding and investigation of socio-demographic, environmental and behavioural risk factors, which may be associated with LBW in these populations.

In general, it is preferable to have actual co-ordinates of the cases and the underlying population rather than relying on centroids of ZIP codes or census areas. The maps produced will have better resolution since they are not dependent on only the location of the ZIP code centroid. In addition, if the data is geocoded to street address the approximate and exact methods of calculating the rates at the grid points will be virtually identical, since we add only a few individuals at a time to the circle as it captures the population needed to calculate the rate. In most situations, however, we do not have the data geocoded to the street address. It is very costly and time consuming to accurately geocode large health outcome data sets, particularly in the rural areas where exact street address information is unavailable. Thus aggregated data must often be used to produce regional maps.

The smoothing method based on a fixed geographic size [4] is a special case of the more general class of methods proposed by Kafadar [5]. Rather than treating all births and LBWs within a specific radius of a grid point equally, births closer to the centre are weighted more than those further away. While not presented in this paper, such a weighting mechanism could easily be combined with a population-defined radius.

What is a reasonable approach for producing maps for surveillance purposes? If the population density is heterogeneous, our recommendation is to produce a series of maps, using fixed or nearly fixed population size spatial filters (Plates 5–8). These should be studied in conjunction with statistical tests for spatial randomness, in the form of smoothed probability maps for filter sizes of interest (Talbot *et al*., manuscript in preparation) or a map showing the statistically significant clusters.

REFERENCES

1. Clayton D, Bernardinelli L. Bayesian methods for mapping disease risk. In *Geographical and Environmental Epidemiology: Methods for Small-Area Studies*, Elliott P, Cuzick J, English D, Stern R (eds). Oxford University Press: Oxford, 1992.
2. Lawson A, Biggeri A, Böhning D, Lesaffre E, Viel J-F, Bertollini R. (eds). *Disease Mapping and Risk Assessment for Public Health*. Wiley: Chichester, 1999.
3. Devine OJ, Louis TA, Halloran ME. Empirical Bayes methods for stabilizing incidence rates before mapping. *Epidemiology* 1994; **5**:622–630.
4. Rushton G, Lolonis P. Exploratory spatial analysis of birth defect rates in an urban population. *Statistics in Medicine* 1996; **15**:717–726.
5. Kafadar K. Smoothing geographical data, particularly rates of disease. *Statistics in Medicine* 1996; **15**:2539–2560.
6. Breiman L, Meisel W, Purcell E. Variable kernel estimates of multivariate densities. *Technometrics* 1997; **19**:135–144.
7. Abramson IS. On bandwidth variation in kernel estimates - a square root law. *Annals of Statistics* 1982; **10**:1217–1223.
8. Silverman BW. *Density Estimation for Statistics and Data Analysis*. Chapman and Hall: London, 1986.
9. Kulldorff M. A spatial scan statistic. *Communications in Statistics: Theory and Methods* 1997; **26**:1481–1496.
10. Kulldorff M, Nagarwalla N. Spatial disease clusters: detection and inference. *Statistics in Medicine* 1995; **14**:799–810.
11. United States Bureau of the Census. 1990 census of population and housing, summary tape file 1B CD-ROM, US Department of Commerce, 1991.
12. Geographic Data Technology Inc. GDT ZIP Code Files, New Hampshire, 1995.
13. Turnbull BW, Iwano E, Burnett W, Howe H, Clark L. Monitoring for clusters of disease: application to leukemia incidence in upstate New York. *American Journal of Epidemiology* 1990; **132**:S136–S143.
14. MapInfo Corp. MapInfo Professional v5.0, Troy, NY, 1998.
15. Northwood Geoscience. Vertical Mapper contour and display software v2.0, Ontario, Canada, 1998.
16. Kulldorff M, Rand R, Gherman G, Williams G, DeFrancesco D. SaTScan v2.1: Software for calculating the spatial and space-time scan statistics, Bethesda, MD, National Cancer Institute, 1998.
17. Committee to Study the Prevention of Low Birthweight, Division of Health Promotion and Disease Prevention, Institute of Medicine. *Preventing Low Birthweight*. National Academy Press, Washington DC, 1985.
18. Kafadar K, Tukey JW. U.S cancer death rates: a simple adjustment for urbanization. *International Statistical Review* 1993; **61**:257–281.

