

01 **Chapter 27**  
02 **Spatial Analytic Approaches to Explaining**  
03 **the Trends and Patterns of Drug**  
04 **Overdose Deaths**  
05  
06

07  
08 **Charlie DiMaggio, Angela Bucciarelli, Kenneth J. Tardiff, David Vlahov**  
09 **and Sandro Galea**  
10  
11

12  
13  
14  
15  
16  
17 **Abstract** To effectively utilize and interpret spatial analyses, substance use re-  
18 searchers, public health practitioners and policy makers should be familiar with  
19 some of the available data analytic techniques, each of which comes with advan-  
20 tages and drawbacks. In this chapter we first discuss three cluster detection tools  
21 and their associated software applications. We then present a Bayesian hierarchical  
22 approach, briefly reviewing its theoretical underpinnings, commonly used models,  
23 and how inferences may be drawn from a sample-based posterior distribution. We  
24 demonstrate the use of each approach on a set of substance abuse mortality data,  
25 comparing the results across the four tools. Our empiric illustration considers the  
26 role of neighborhood-level socioeconomic status (SES) in explaining opiate-related  
27 overdose deaths in New York City. We end with a discussion of the implications of  
28 the choice of technique and software on interpreting spatial analyses of substance  
29 abuse and conclude that the choice of a method will be driven by the question to  
30 be answered, data and software availability and the intended audience or context in  
31 which the research is being conducted.  
32

33  
34 **Introduction**  
35

36 Mapping techniques and spatial analysis have been used in a number of stud-  
37 ies seeking to describe and analyze substance abuse. Spatial analytic studies have  
38 demonstrated the correlation of drug use to deprivation indices (Squires, Bleeching,  
39 Schlecht, and Ruben 1995); the role social networks play in urban adolescent  
40 substance abuse (Mason, Cheung, and Walker 2004); the effect of ecologic level  
41

42  
43 S. Galea  
44 Department of Epidemiology, University of Michigan, 1214 S. University, Room 243, Ann Arbor,  
45 MI 48104, USA  
e-mail: sgalea@umich.edu

01 variables, such as legal prohibitions against alcohol sales (Schulte, Aultman-Hall,  
02 McCourt, and Stamatiadis 2003); and whether frequency and type of drug use are  
03 geographically located independent of neighborhood characteristics (Latkin, Glass,  
04 and Duncan 1998).

05 There are a number of spatial analytic tools available to epidemiologists, each  
06 having advantages as well as drawbacks. To effectively utilize and interpret spa-  
07 tial analyses of substance abuse, researchers, public health practitioners, and policy  
08 makers should be familiar with some of the available data analytic techniques.

09 In this chapter, we first discuss three cluster detection tools and their associ-  
10 ated software applications: nearest neighbor index (NNI; ESRI 2005), Ripley's  
11 K-function (Levine 2004), and a space-time and time permutation scan statistic  
12 (Kuldorf 2005). We briefly describe these techniques and then demonstrate their  
13 use on a set of substance abuse mortality data, comparing the results across the  
14 three tools. We then introduce hierarchical spatial modeling (Imperial College and  
15 Medical Research Council 2004). We will discuss the advantages and disadvan-  
16 tages of a Bayesian approach, some commonly used models, and how to draw  
17 inferences from the sampled posterior distribution. We will demonstrate this ap-  
18 proach on our data set and compare the results to those we obtained with cluster  
19 detection tools. As an empiric illustration, we consider the role of neighborhood-  
20 level socioeconomic status (SES) in explaining opiate-related overdose deaths in  
21 New York City (NYC). We conclude with a discussion of the implications of the  
22 choice of software and techniques on interpreting spatial analyses of substance  
23 abuse.

24

25

### 26 *Data and Variable Definitions*

27

28 We manually reviewed medical files at the Office of the Chief Medical Examiner  
29 (OCME) of NYC and identified all cases of fatal accidental drug overdose occur-  
30 ring in the city between 2000 and 2004. The OCME is responsible for assessing  
31 all deaths of persons believed to have occurred in an unnatural manner in NYC.  
32 Therefore, all overdose deaths in NYC would have been reviewed by the OCME  
33 and included in this chart abstraction.

34

35 The OCME medico-legal investigators use the decedent's medical history, the  
36 circumstances and environment of the fatality, autopsy findings, and laboratory data  
37 in attributing the cause of death and other criteria to each case being reviewed.  
38 The variables we abstracted for our analysis included the decedents' gender, age,  
39 address of residence, and location of injury. We geocoded residential and injury  
40 locations using ArcGIS, version 9.1 (ESRI 2005). For analysis purposes, place  
41 of injury (location of death) was used. In the analysis, only decedents who were  
42 successfully assigned an address of injury were included. Overdose deaths may  
43 include more than one drug being present; to increase the reliability of our mea-  
44 sures across the different analytic techniques, we restricted our analyses to cases  
45 in which opiates were the only (in the case of scan statistics) or primary cause of  
46 death.

**Table 27.1** Demographic characteristics of successfully geocoded opiate-related drug overdose deaths, New York city 2000–2004

	No.	%
Total	2426	
Sex		
Male	1883	77.6
Female	543	22.4
Age		
15–24 yrs	146	6.0
25–34 yrs	431	17.8
35–44 yrs	898	37.0
45–54 yrs	762	31.4
55–64 yrs	162	6.7
65–74 yrs	25	1.0
Over 74 yrs	1	0.0
Race/ethnicity		
White	1069	44.1
Black/not Hispanic	560	23.1
Hispanic	776	32.0
Asian	11	0.5
Other	10	0.4
Year of death		
2000	543	22.4
2001	455	18.8
2002	484	20.0
2003	506	20.9
2004	438	18.1

### *Descriptive Epidemiology*

From 2000 through 2004, the OCME reported 3982 fatal overdose deaths within NYC. Of these, 3777 occurred among NYC residents, in which 2516 were determined to have opiate toxicity as the primary cause of death. Together, 96.4% (2426 out of 2516) cases were successfully geocoded. These cases constituted the study base for our subsequent analyses. Their demographic characteristics are presented in Table 27.1.

### **Cluster Detection Techniques**

We conducted cluster analyses for all opiate-related deaths. We first described cross-sectional spatial distribution of all fatal opiate-related deaths in NYC that occurred between 2000 and 2004 using an average NNI statistic. We then utilized an L function transformation (derived from Ripley's K-function) to produce graphs to assess at what distance the clustering (if present) was observed to be the greatest for each year. We then used a space–time permutation model to assess both the spatial and temporal clustering opiate-related overdose deaths. This approach uses

count (case) data only, and assesses not only spatial clustering characteristics but also the role of time as a variable over the 5-year period. In this way, we determined if any clusters were statistically significant when adjusting for the year in which the deaths occurred. Finally, we applied a space–time scan statistic that differs from the permutation model in that it also adjusts for the underlying population at the census tract level.

### *First-order Clustering Technique: Average Nearest Neighbor Index*

#### **Description**

The global presence or absence of clustered overdose incidences can be assessed using the average NNI. This index is a measure of how similar the mean distance of all cases is to the expected mean distance for a hypothetical random distribution (Mitchell 2005). The equation for calculating the average NNI is expressed as (Mitchell 2005):

$$d = \left( \frac{\sum_i C_i}{n} \right) - \left( \frac{0.05}{\sqrt{n/A}} \right)$$

where the average NNI ( $d$ ) is equal to the summed distance to each feature nearest neighbor ( $\sum_i C_i$ ) divided by the number of features ( $n$ ) or the ‘observed distribution’ of mean features minus the product of 0.05 divided by the square root of the number of features ( $n$ ) divided by the study area ( $A$ ) or the ‘expected mean distance for a random distribution’ (Clark and Evans 1954). Clustering is suggested when the observed average distance is greater than the mean random distance ( $d < 1$ ). An index value close to 1 indicates randomness, while a value greater than 1 indicates dispersion of cases. Within ArcGIS, version 9.1, tests of significance (a  $z$ -score and  $p$ -value) are included with the NNI output. If the  $z$ -score is negative, this suggests the cases are clustered. Conversely, if the  $z$ -score is positive, this suggests the cases are dispersed, while a value close to zero indicates the random nature of cases.

#### **Application**

Table 27.2 displays the average NNI results of all opiate-related drug overdose deaths. The results suggest that the greatest clustering of any overdose occurring in NYC was in 2000 (NNI = 0.7636;  $z$ -score =  $-10.5474$ ). Given the statistically significant negative  $z$ -score values of all 5 years, individually, opiate-related drug overdose demonstrates significant clustering for all years between 2000 and 2005.

01 **Table 27.2** Average nearest neighbor analysis of opiate-related drug overdose deaths, New York  
 02 city, 2000–2004\*

03	Year of Death	Average	z-score	p-value
04		Nearest	(Standard	
05		Neighbor	Deviations)	
06		Ratio		
07	2000–2004	0.7023	–27.8189	<0.0001
08	2000	0.7636	–10.5474	<0.0001
09	2001	0.7216	–11.2869	<0.0001
10	2002	0.7659	–9.9835	<0.0001
11	2003	0.8235	–7.1310	<0.0001

12 \*Based on weighted counts of injury location and direct distance measurement.

13  
 14 Moreover, clustering became less dense toward the later years of the study period as  
 15 displayed by the gradually increasing z-scores.

16  
 17  
 18 ***Second-order Clustering Technique: Ripley K-Function***  
 19 ***(L-Transformation) Statistic***  
 20

21 **Description**

22 While the average NNI considers only the distance between one case and its nearest  
 23 other case, the Ripley’s K-function statistic is a second-order statistic that considers  
 24 the complete distribution of all distances in the point pattern of cases (Levine 2004).  
 25 It tests the cumulative distribution function of the entire set of inter-point distances  
 26 among the point data. When K statistics are transformed into a square root function,  
 27 the result is called a L function transformation (L(d)). The square-root transforma-  
 28 tion results in a linear function. This statistic can be very useful when exploring the  
 29 nature, in terms of distance, of the case clustering within the entire study area. L  
 30 function equation is expressed as (Levine 2004):  
 31

32  
 33  
 34 
$$L(d) = \sqrt{\frac{A \sum_{i \neq j} \sum I_{ij} d_{ij}}{\pi n(n-1)}}$$
  
 35  
 36  
 37

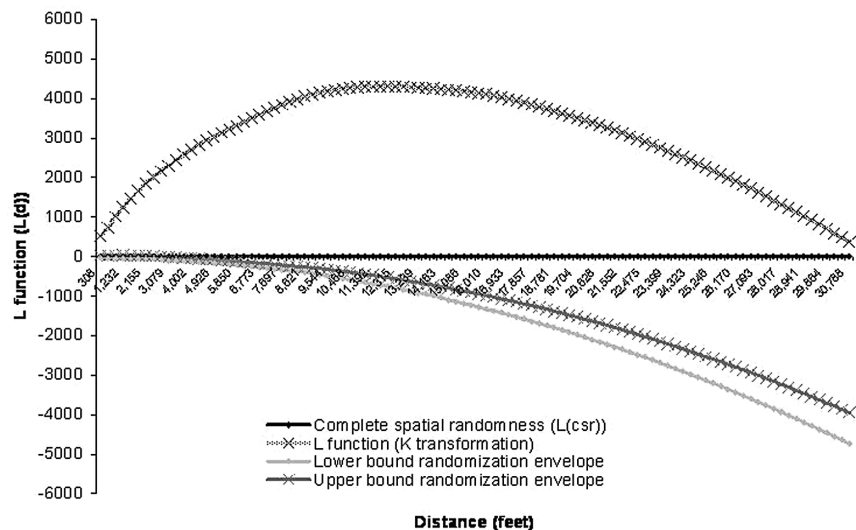
38 The numerator is the Ripley’s K-function, where the distance (d) is measured be-  
 39 tween case (i) and every other case (j); then each distance is multiplied by the  
 40 weight for the case paring (I<sub>ij</sub>), and all the values are summed (i ≠ j indicates  
 41 the distance between cases) are not included in the sum (Levine 2004). Finally, the  
 42 result is multiplied by the study area (A) and divided by the number of cases (n)  
 43 squared. The denominator is π multiplied by the number of possible case pairings  
 44 (represented as n–1). The square root of the product is then taken. At any given  
 45 distance (represented by the x-axis of the result graph), if the line of observed L

01 values is above that of the expected values [ $L = 0$  or complete spatial randomness  
 02 (CSR)], then the cases are more clustered than expected for a random distribution  
 03 (with the peak of the graph representing the greatest clustering detected at a speci-  
 04 fied distance) (Levine 2004). Once the curve falls below the CSR line, cases at that  
 05 point become dispersed at a given distance.

06 To test the null hypothesis of global spatial randomness of opiate-related over-  
 07 dose for the period of interest, we computed a 95% CI (referred to as the envelope)  
 08 of the L-function ( $L(d)$ ) using a Monte Carlo method of specifying 100 simulated  
 09 random patterns (Levine 2004). At a given distance (represented on the  $x$ -axis),  
 10 a value of  $L(d)$  (represented on the  $y$ -axis) outside the confidence interval (CI)  
 11 envelope is interpreted as a significant departure from CSR toward clustering or  
 12 dispersion. When the function peaks at the largest, most positive value and remains  
 13 outside the CI envelope, this is considered to be the distance at which cases tend to  
 14 be the most clustered.

16 **Application**

17  
 18 Figure 27.1 presents a graph of the L-function statistic for the entire 2000–2004 time  
 19 period, and suggests that most clusters of opiate-related overdose are fairly compact  
 20 and that the greatest clustering occurs at distance of approximately 12,007.34 feet  
 21 or 2.27 miles [distance at which  $L(d)$  peaks in the output]. This suggests that most  
 22 clusters occur with a radius of approximately 2.25 miles. Thereafter, the cases be-  
 23 come more dispersed. The  $L(d)$  curve also remains outside the 95% confidence en-  
 24 velope and, therefore, remains statistically significant. Additional curves for individual  
 25



44 **Fig. 27.1** Graph of L-function statistic. New York City opiate-related overdose deaths, 2000–2004

01 years (not presented) are also clustered at approximately 2–3 miles, with the tightest  
 02 clusters observed to be in 2000 (10,775.82 feet or 2.27 miles).

03

04

05

06

07

08

09

### *Space–Time Modeling Techniques: Space–Time Permutation Statistic and Space–Time Scan Statistic (SaTScan)*

#### **Description**

10

11

12

13

14

15

16

17

18

19

20

21

22

23

24

25

26

27

28

29

30

31

32

33

34

35

36

37

38

39

40

41

42

43

44

45

Pure spatial analyses, e.g., NNI and Ripley’s K, are useful when exploring cross-sectional health outcomes. When the variable of ‘time’ (in units of hours, days, months, years, etc.) is of interest, we will need a model that assesses the trend of the outcome over both space and time. We are interested in whether the same areas experience clustering year after year, asking: Are the cases clustered and, if so, do they continue to cluster over time given the nature of the study area?

The space–time permutation scan statistic model uses only case data. There is no requirement for specifying the underlying population data. It makes minimal assumptions about the time, geographic location, or size of the potential case clustering. The model adjusts for what is termed as ‘purely’ spatial and temporal variation in the case data for a given area (Kulldorff, Heffernan, Hartman, Assuncao, and Mostashari 2005). Using a Poisson-based probability model, a series of overlapping scanning windows (cylindrical in form) move across the spatial plane (the base of the cylinder) while also scanning the point data for temporal clusters (the height of the cylinder). The circular base represents the geographical area or the study area while the height of the cylinder scans for time (in days, months, or years) clustering. For each location, the scanning window calculates the number of observed and expected cases. The statistical significance of an observed ‘cluster’ is then evaluated taking into account the multiple testing methods (0, 9, or 999 Monte Carlo replications). For each center and radius of the cylinder base, the method iterates over all possible temporal cylinder lengths. Cylinders can be geographically large and temporally short (forming a flat disc), or can be geographically small and temporally long (forming a pole), or any combination in between. The number of observed cases is divided by the calculated expected number of cases for each cylinder to the power of the observed inside the cylinder, and then multiplied by the observed, divided by the expected to the power of the observed outside the cylinder. The approximation, a Poisson generalized likelihood ratio, is expressed as (Kulldorff et al. 2005):

$$\left(\frac{C_A}{\mu_A}\right)^{C_A} \left(\frac{C - C_A}{C - \mu_A}\right)^{(C - C_A)}$$

where  $C$  is the total number of observed cases,  $\mu_A$  represents the mean number of expected number of cases within the cylinder, and  $C_A$  represents the observed number of cases within the cylinder.

01 The space–time scan statistic is also based on Poisson modeling, just as the  
02 space–time permutation model, but allows for scanning of purely spatial, purely  
03 temporal, and special temporal clusters.

04 These models are most readily applied using the SaTScan software package,  
05 which is available for download (Kuldorf 2005) after registration and can be trans-  
06 lated in ArcGIS for viewing of the cluster statistics.

## 08 **Application**

09  
10 The space–time permutation scan statistic of opiate-related overdose was mapped  
11 to give a visual display of the model output (see Fig. 27.2a). In terms of spatiotem-  
12 poral clustering for opiate-related overdose citywide, seven clusters were detected.  
13 The primary cluster was detected in the northwestern portion of Manhattan and  
14 the southern region of Bronx in 2001, while the other six secondary clusters were  
15 located in various parts of the city. It is notable that none of the clusters reached a  
16 level of statistical significance.

17 The majority of opiate-related drug clusters seem to have occurred in 2000,  
18 although the primary cluster in northern Manhattan/South Bronx was detected in  
19 2001. It is suggested that opiate-related fatal overdose is not only changing inci-  
20 dence pattern, but also such cases are decreasing.

21 In our comparison space–time scan statistic model, we attempted to see how  
22 clusters may change when adjusting for the underlying population counts. In this  
23 analysis, we used population counts at the census tract level for the entire NYC area  
24 for a finer and more exact population adjustment. Using a small neighborhood unit  
25 of population count adjustments allows for a finer resolution of cluster detection.

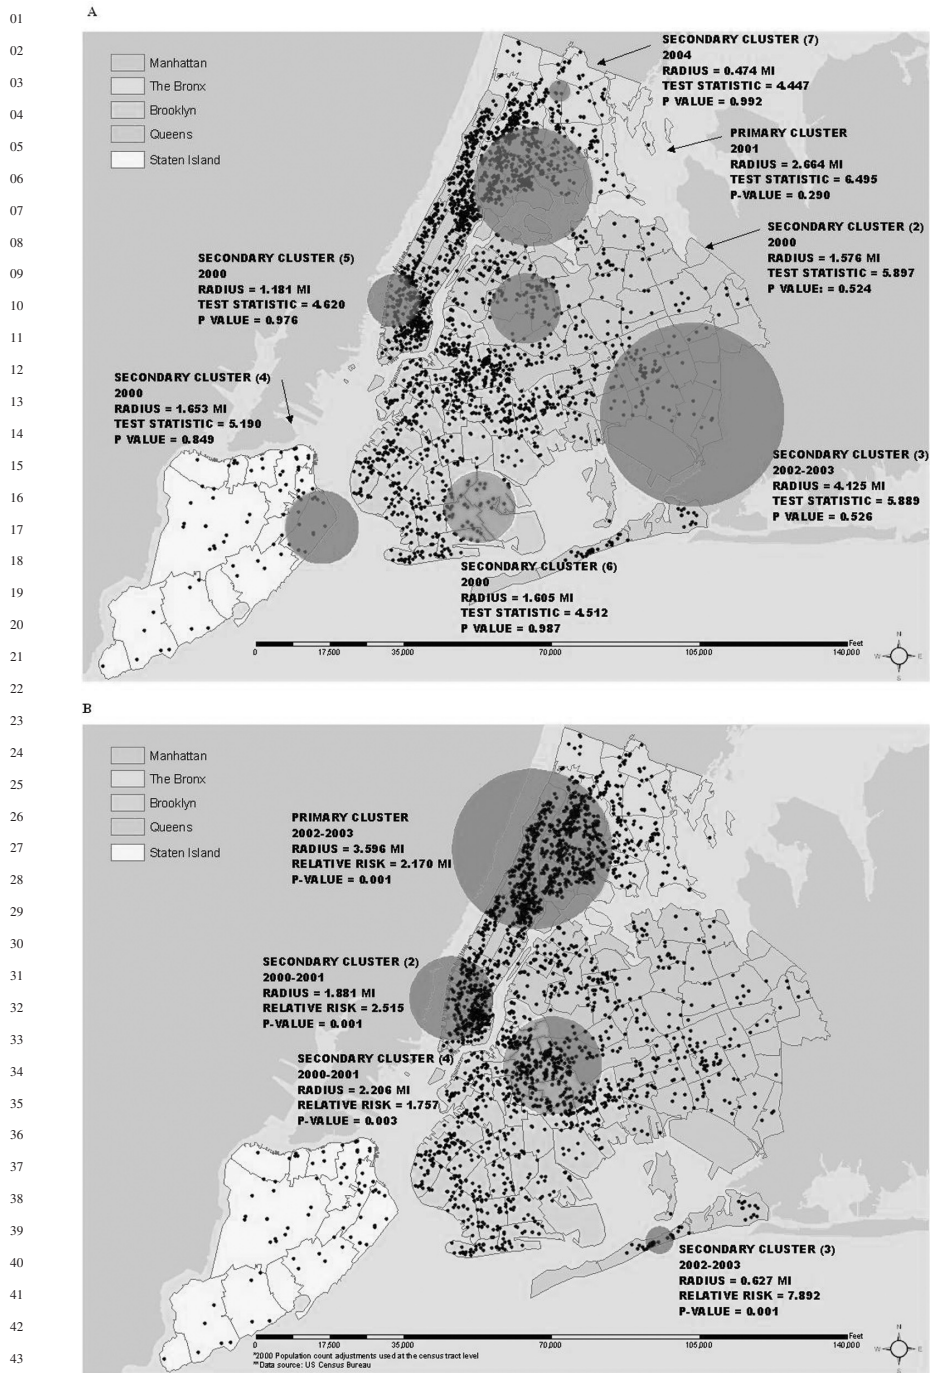
26 The space–time scan statistic detected fewer clusters, all of which were statis-  
27 tically significant (Fig. 27.2b). What remains as the primary cluster (just as with  
28 the space–time permutation model cluster map) is the one located in the northern  
29 region of Manhattan and the southern region of Bronx. Compared to the space–time  
30 permutation cluster map, the central radius point of the cluster is shifted slightly  
31 west (encompassing more of Manhattan than Bronx), the radius of the cluster is  
32 approximately one mile larger (radius = 3.596 miles as compared to 2.664 miles),  
33 and the year in which cases were significantly detected was 2002–2003 as compared  
34 to 2001 for the space–time permutation model.

35 There are additional differences between the two approaches. A cluster in lower  
36 Manhattan detected by both methods is somewhat larger in the space–time scan,  
37 which additionally detected higher than expected cases in both 2000 and 2001 as  
38 compared to only 2000 in the space–time permutation scan.

## 41 **Bayesian Hierarchical Models**

42  
43 Bayesian hierarchical modeling is frequently used in spatial epidemiological anal-  
44 yses, but may be unfamiliar to some substance abuse researchers. Multilevel spa-  
45 tial modeling, though, has been used to capture context in community studies





**Fig. 27.2** Opiate-related deaths, New York City, 200–2004. (a) Space-time permutation scan indicating 7 clusters (b) Controlling for underlying population count (See also Plate 48 in the Colour Plate Section)

01 of substance abuse (Luke 2005) and in studies of drug-related crime (Law and  
02 Bayesian 2004). In this section, we review Bayesian methods, consider how they  
03 may address certain difficulties encountered in other spatial analytic techniques,  
04 and present the results of their use in our sample data set.

05

06

## 07 *Description*

08

### 09 **Mapping issues and the Bayesian approach**

10

11 In the classical maximum likelihood approach to risk measures, such as standardized  
12 mortality ratios, the risk estimate for each area  $j$  is given by the observed  $j$ /expected  
13  $j * 100$  with the standard error under an assumption of a Poisson distribution for  
14 each area given by the square root of the observed number divided by the expected  
15 number.

16

17 There are problems with this approach for spatial analyses. The map may be  
18 dominated by extreme values based on a few cases in small populations (Devine,  
19 Louis, and Halloran 1994). These rare events contribute to more heterogeneity than  
20 is assumed by a Poisson model (where  $\mu$  is expected to be close to 1 and equal to  $S$ ).  
21 A simple maximum likelihood approach also does not account for spatial correla-  
22 tion. Influential covariates of an outcome, which may be unmeasured, are likely to be  
23 similar in adjacent areas resulting in risk estimates that are also spatially correlated  
24 and similar. In situations when there are a small number of correlated cases relative  
25 to those at risk and Poisson ‘noise’ obscures the ‘signal’ of the spatial pattern in the  
26 data, hierarchical Bayesian modeling can be useful (Richardson, Abellan, and Best  
2006).

27

28 In a Bayesian approach, our two main sources of information about the risk es-  
29 timate for an area ( $\theta$ ) are our prior beliefs about  $\theta$ , called the *prior distribution* or  
30  $p(\theta)$ , and the *likelihood* of observing our data given  $\theta$  or  $L(y|\theta)$ . We thus specify a  
31 probability distribution of risk estimates ( $\theta$ ) that vary across the areas of the map in  
32 some defined fashion, e.g., they may be normally distributed or Poisson distributed.  
33 This prior distribution may be based on previous studies, literature reviews, or expert  
34 opinions, and informs about  $\theta$  through our beliefs or assumptions. The likelihood  
35 informs about  $\theta$  via the data itself. When we have lots of data, the likelihood pre-  
36 dominates our analysis, and our results will essentially be the maximum likelihood  
37 estimate. When we have less data, the prior has greater influence (Greenland 2006;  
38 Lawson, Browne, and Vidal 2003). The result of combining the prior distribution  
39 and the likelihood is called our posterior distribution.

39

40 Choice of our prior distribution is critical as it essentially indicates how we  
41 believe the parameter would behave if we had no data from which to make our  
42 decision. What, for example, might we expect is the probability that someone living  
43 within 3 miles of a certain location would die from an opiate-related overdose?  
44 Our best guess might be, for example, 1 in 20 or about 5%, and that this prob-  
45 ability varies around this point estimate in a normal fashion with a variance of,  
say, 0.01 or 1%. This estimate may be based on previous studies, law enforcement

01 data, clinical experience, or a combination of sources. What if we conduct a study  
 02 that indicates the risk of an opiate overdose within 3 miles of the location is 45%?  
 03 How likely is our observed data given our postulated prior probability? Our pos-  
 04 terior distribution combines our expectation with what we actually observe. In a  
 05 very common sense way, it tells us, for example, that if the results of our study  
 06 differ markedly from our best existing information, we should perhaps be somewhat  
 07 skeptical.

08 The approach is hierarchical or mixed because we specify a distribution of  
 09 *hyperparameters* ( $\lambda$ ) for our risk parameter,  $\theta$  allowing it to vary across each area.  
 10 One could, for example, say that  $y_i$  is the empirical (observed) rate of substance  
 11 abuse-related deaths in zip code  $i$ ,  $\theta$  is the true underlying rate, and  $\lambda$  how that true  
 12 rate varies (Banerjee, Carlin, and Gelfand 2004).

13 As noted, the posterior distribution ( $\Pr[\theta|y]$ ) is based on our prior assumptions  
 14 and our observed data. It follows Bayes' theorem and is proportional to the likeli-  
 15 hood times the prior: (Greenland 2006)

$$16 \Pr[\theta|y] \propto \Pr[y|\theta] * \Pr[\theta],$$

17  
 18  
 19 As described by Richardson, Abellon, and Best (2006), hierarchical Bayesian  
 20 spatial models describe observed cases in an area as Poisson distributed with a mean  
 21 equal to the expected number of cases times the risk for that area:

$$22 O_i \sim \text{Poisson}(\rho_i E_i).$$

23  
 24  
 25 At the second level of the model, the risk for each area ( $\rho$ ) is transformed to a  
 26 log scale (making relationships additive rather than multiplicative) and is described  
 27 as an intercept term ( $a$ ) and two random effects, one spatial ( $\theta$ ), the other non-  
 28 spatial ( $\Psi$ ):

$$29 \log \rho_i = a_i + \theta_i + \psi_i$$

30  
 31  
 32 The spatially structured component is described as a conditional autoregressive  
 33 (CAR) Gaussian process [ $\theta \sim \text{CAR normal}(W, \tau_\theta)$ ] where the conditional dis-  
 34 tribution of each  $\theta_i$ , given all the other  $\theta_i$ 's, is normal with  $\mu$  = the average  $\theta$   
 35 of its neighbors and a precision ( $\tau_\theta$ ) proportional to the number of neighbors.  $W$   
 36 represents the matrix of neighbors that defines the neighborhood structure. The  
 37 non-spatial component of the model ( $\psi_i$ ) is defined as normally distributed with  
 38  $\mu = 0$  and precision ( $\tau_\psi$ ). The model is completed by assigning hyper-priors to the  
 39 precision terms  $\tau_\theta$  and  $\tau_\psi$ .  
 40  
 41

#### 42 **The Poisson–gamma model**

43  
 44 This hierarchical Bayesian approach most frequently described in the mapping liter-  
 45 ature is the Poisson–gamma model. In this formulation, the risk ( $\theta$ ) is described as

01 a set of parameters that may include any number of explanatory variables (Law-  
02 son et al. 2003). The prior distribution of the observed outcome  $y$  is described  
03 as  $y|\theta \sim \text{Po}(\theta E)$ , and the hyper-prior distribution of risk is  $\theta|\alpha, \beta \sim \text{gamma}$   
04  $(\alpha, \beta)$ , with  $\mu = \alpha/\beta$  and  $\sigma^2 = \alpha/\beta^2$  (Banerjee et al. 2004; Lawson et al.  
05 2003). We could further specify  $\alpha$  and  $\beta$ , but we assume that beyond a cer-  
06 tain point, further model specification will have little practical effect on our  
07 results.

08 We commonly choose a non-informative (proper) or arbitrarily vague-prior that  
09 is uniform or ‘flat’ to allow the data to predominate and lead us to a posterior  
10 distribution that is dominated by the likelihood. A gamma (0.5, 0.0005) has been  
11 suggested as reasonable (Law and Bayesian 2004).

12 For simple models for which there is a closed form (i.e., they behave as true  
13 distributions and integrate to 1), we can estimate the posterior distribution directly  
14 via the maximum likelihood estimate, and a Bayesian approach is unnecessary. But  
15 for most reasonably realistic models, we will not be able to find a closed form and  
16 will need sample-based approaches.

17 Empirical Bayes methods approximate the posterior distribution (Devine et al.  
18 1994). Full Bayes methods base inferences on a sample of the full posterior distribu-  
19 tion. The results from such a sample are not as informative as the closed form itself,  
20 but are usually sufficient for inference. We increase the precision of our estimates  
21 by increasing the sample size (Banerjee et al. 2004).

22 One way to construct a sample from the posterior distribution is through Markov  
23 Chain Monte Carlo (MCMC) methods. Like a ‘random walk’ seen in time series  
24 analysis, the resulting series has no ‘memory’. Subsequent values depend only on  
25 the current value, and the series converges to a stationary distribution assumed to  
26 be the posterior. Unlike traditional Monte Carlo methods, MCMC methods produce  
27 correlated samples, because they base subsequent values on current values. Meth-  
28 ods such as thinning every other value may help decrease this correlation (Imperial  
29 College and Medical Research Council 2004).

30 Transition probabilities for selection into the series are typically determined  
31 through the use of the so-called Gibb’s sampler. A special case of the Metropolis–  
32 Hastings algorithm, the Gibb’s sampler generates conditional probability distribu-  
33 tions of a parameter given all other parameters, and transition probabilities are  
34 generated that result in a proposal value that accepts or rejects the value with a  
35 probability of 1 or 0 (Lawson et al. 2003). The algorithm is useful in the context of  
36 Markov random fields where the joint posterior distribution is complicated but the  
37 full posterior prior distributions have simple forms.

38 As noted, the spatial Poisson process consists of two components: uncorre-  
39 lated global heterogeneity ( $\Psi$ ), usually due to unmeasured confounders or effects  
40 throughout the data, and correlated or specific heterogeneity, due to spatial correla-  
41 tion or local effects ( $\theta$ ) (Lawson et al. 2003; Richardson et al. 2006. To capture  
42 both spatial variation and non-spatial random effects in an additive fashion, we  
43 model the natural log of risk as the sum of these two components (da Silva and  
44 Melo 2004).

### 01 **Setting the correlational structure**

02 Since we want to model the spatial components so that geographically close areas  
 03 present similar risks, we use information from other areas in the region to reduce  
 04 random variation unrelated to the risk represented by our risk estimate. This should  
 05 take spatial correlation into account and result in smoother informative maps. Here,  
 06 we see the advantage of a Bayesian approach. Modeling  $\theta$  as a random variable  
 07 rather than a fixed variable allows us to set a spatial correlational structure.

08 We can describe this structure via Markov random fields where each  $\theta$ , given all  
 09 the other  $\theta$ 's, depends only on its neighborhood. A Markov random field is a locally  
 10 specified joint distribution that can be determined by its full conditionals. Given a  
 11 joint distribution,  $\Pr[y_1 \dots y_n]$ , the set of full conditional distributions,  $\Pr[y_i | y_j]$ , that  
 12 we can create from it are uniquely determined. Brook's lemma tells us that we can  
 13 go in the opposite direction. If we have a set of full conditional distributions, we can  
 14 get the unique joint distribution from which they arose (Banerjee et al. 2004).

15 A locally determined, weighted structure can be represented by a CAR Gaussian  
 16 model where the conditional distribution of each  $\theta$  is given by:

$$17 \theta_i | \theta_j \sim N_1(\sum w_{ij} \theta_j / \sum w_{ij}, 1 / \tau \sum w_{ij})$$

18 where  $j$  is not equal to  $i$  and is an element of  $d$ , the set of the neighbors of  $i$  (da  
 19 Silva and Melo 2004).

20 The simplest and most commonly used definition of a neighborhood is the existence  
 21 of a common border between areas. In this case, the weights are specified as  
 22  $w_{ij} = 1$  if  $j$  is in  $d$ , and  $w_{ij} = 0$  if  $j$  is not in  $d$ . In this case, the  $\sum w_{ij}$  is simply the  
 23 number of neighbors of area  $i$ . So, the conditional prior mean of  $\theta$  is given by the  
 24 arithmetic average of the spatial effects of its neighbors, and the conditional prior  
 25 variance is proportional to the number of neighbors.

26 This structure has been used in a number of disease mapping studies (da Silva and  
 27 Melo 2004), but other approaches are possible. We could set up a proximity matrix  
 28 of weights,  $w_{ij}$ 's, based on the distance between or other relationships between the  
 29 spatial units, e.g., a set of first order  $w_{ij}$ 's ( $w_{ij}[1]$ ), if an areal unit is, say, less than  
 30 some predefined distance (Banerjee et al. 2004). We could also adopt more compli-  
 31 cated relationships between spatial units to represent, for example, the movement of  
 32 commuters from one area to another. The more the neighbors we include, the more  
 33 the smoothing we achieve. The limit would, of course, be averaging over all the  
 34 spatial units, which would be the overall average and would not be very informative  
 35 at the local level.

### 36 **Interpreting results**

37 An important consideration in MCMC methods is diagnosing convergence to the  
 38 stationary Markov Chain. A commonly accepted approach is to run and dynam-  
 39 ically monitor a specific number, e.g. 3, parallel chains and examine the trace plots  
 40 for when they start to overlap as an indication of convergence. We then discard the  
 41

01 burn-in period samples and base inference on the stationary Markov Chain. The  
02 Gellman Rubin statistic is useful in diagnosing convergence. It compares variation  
03 within chains to those between chains for evidence of scale reduction. When the  
04 scale reduction factor reaches 1, there is evidence of convergence. Other conver-  
05 gence statistics are based on examining individual chains (Lawson et al. 2003).

06 Once the posterior distribution has been sampled, a Bayesian CI has a straightfor-  
07 ward interpretation. In a 95% CI, we are 95% certain that the true value lies within  
08 it. It is most easily obtained by chopping off the  $\alpha/2$  tails of the posterior probability  
09 distribution.

10 Being a sample-based approach, in MCMC methods, no two analysts will end  
11 up with exactly the same results. This makes variance assessment crucial. Recall  
12 that MCMC produces correlated samples. This will lead to underestimates of the  
13 variance. One may ‘thin’ the samples to decrease correlation, but this would result  
14 in throwing out information. Variance estimates based on an effective sample size  
15 (ESS), though, are available (Banerjee et al. 2004).

## 17 *Application*

18  
19  
20 We calculated standardized mortality ratios (SMR) for opiate-related deaths in NYC  
21 for the years 2000–2004 using the expected number of overdose deaths in a zip  
22 code tabulation area based on the mean number of such deaths in NYC throughout  
23 the 5-year study period. We were interested, in this example, in drawing inference  
24 about the potential role of SES as an explanatory variable for opiate-related overdose  
25 deaths in NYC. We used zip code level median household income (MHI) as a proxy  
26 for neighborhood-level SES throughout this example.

27 In the model, the likelihood of the observed values in the standardized morbidity  
28 ratio is modeled as a Poisson distribution. The log of the observed value is a function  
29 of the log of the expected value, an intercept, and the coefficient for a normally  
30 transformed median household income measure. Random effects are represented  
31 by a Gaussian intrinsic CAR model with the weights for adjacent neighbors set  
32 at 1. Non-informative prior distributions are placed on the intercept, the coeffi-  
33 cients, and on tau, the precision term for the CAR prior distribution for random  
34 effects.

35 Our interest is in mapping the zip code level risk estimates while accounting  
36 for the potential instability and autocorrelation of those rates and controlling for  
37 MHI.

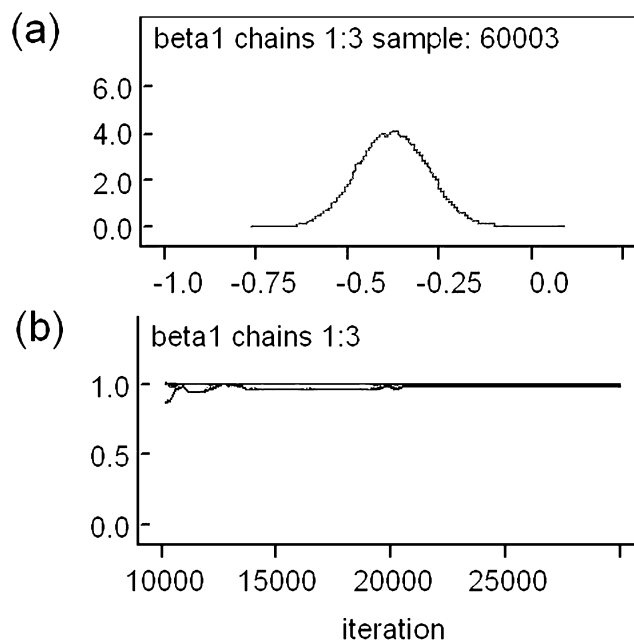
38 For this empiric illustration, we used WinBUGS (Imperial College and Medical  
39 Research Council 2004) software to run three parallel MCMSs with over-dispersed  
40 initial values for 120,000 iterations. The first 60,000 iterations were discarded as a  
41 burn-in, and our inferences were based on the second 60,000 iterations. We assessed  
42 convergence by examining trace histories for parallel chains, and we used R (20)  
43 software to conduct the Brooks, Gelman and Rubin, and the Geweke convergence  
44 diagnostics as well as the Heidelberger and Welch stationarity tests. We present our  
45 results as median values for the coefficients with their associated 95% equal-tailed

01 Bayesian CIs, histograms (kernel density graphs) of the sampled distributions, and  
 02 maps comparing raw and smooth risk estimates.

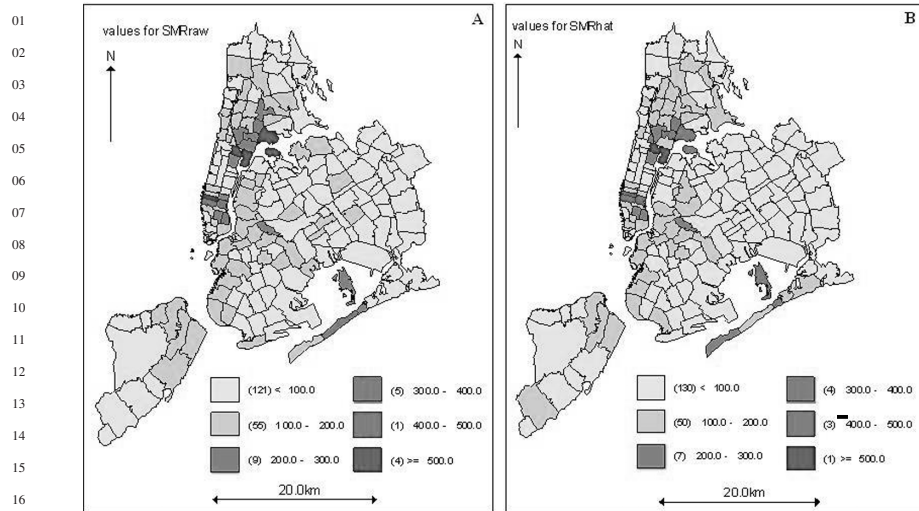
03 As indicated in Fig. 27.3a, the histogram for  $\beta_1$  (MHI coefficient) is smooth and  
 04 normally distributed. The Gellman Rubin statistic was calculated to be 1 for most  
 05 of the runs (Fig. 27.3b), and the trace history for  $\beta_1$  appears to reasonably overlap.  
 06 Taken together, this information gives us confidence that the model appropriately  
 07 converged to the posterior distribution and that our inferences based on this posterior  
 08 distribution are valid.

09 The median value for the MHI coefficient was  $-0.3782$  (95% CI 0.5681,  
 10  $-0.1855$ ). The interpretation is not straightforward because the dependent variable  
 11 in the model is the natural log of the SMR, and the MHI variable itself has been  
 12 normalized to achieve appropriate convergence in WinBUGS. It does, though, indicate  
 13 that SES, as measured by MHI, is strongly and significantly inversely related  
 14 to the number of heroin overdose deaths in a zip code area. Essentially, the number  
 15 of opiate-related deaths increases in a linear fashion as MHI declines. That this is as  
 16 expected indicates, to a certain extent, the validity of the model. Also, and perhaps  
 17 more importantly, the subsequent fitted SMR values now control for this important  
 18 potentially confounding variable.

19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41



42 **Fig. 27.3** Results from Markov Chain Monte Carlo run, association of median household income  
 43 with opiate overdose standardized mortality ratios, New York City zip code tabulation areas, 2000–  
 44 2004. (a) Histogram for Median Household Income beta coefficient. (b) Trace history for Median  
 45 Household Income beta coefficient



**Fig. 27.4** Opiate-related standardized mortality ratios, New York City zip code tabulation areas, 2000–2004. (a) Unadjusted (b) Smoothed estimated adjusted for median household income and autocorrelation with Bayesian Hierarchical Modeling

Figure 27.4 presents the raw and fitted zip code level SMRs. It appears that if we had simply mapped the raw SMRs, we would infer greater than expected rates in such areas as South Bronx and Northern Queens. Looking at the fitted values, these potential clusters seem to become less severe when we take the underlying distribution of the population and its SES characteristics into account. The most evident area of continuing concern is Harlem in northern Manhattan which, despite controlling for MHI, continues to display much greater than expected numbers of heroin overdose deaths.

## Discussion and Conclusions

The methods presented in this chapter have much to offer the substance use researcher. They can be viewed as offering incrementally more information and detail as one progresses from first-order cluster detection methods, such as the NNI through scan statistics, to more explanatory analytic techniques, such as hierarchical modeling.

While we did not spend much time on it, an essential first step in any spatial analysis is to describe the data in terms of summary statistics and simple plots. Not only does this provide key descriptive information, it also allows the researcher to assess whether the data meet the assumptions underlying subsequent tests, e.g., Poisson distributions of the outcome of interest.

First-order clustering methods offer a relatively straightforward and easily interpretable global assessment of whether clustering exists and how tight the clustering



01 appears to have been. These are important considerations, but most analysts will  
02 want to know exactly where and when the clustering occurred and whether any ap-  
03 parent clustering was simply due to chance. Scan statistics are an important public  
04 health tool in this regard. They are fairly straightforward in their application and  
05 interpretation, and allow the incorporation of a time variable.

06 When a population denominator is not available or appropriate, the space–time  
07 permutation model is very useful. Although none of the clusters detected through  
08 the use of space–time permutation model (Fig. 27.2a) were statistically significant,  
09 the results suggest that certain areas had higher counts than expected and that there  
10 were temporal changes in opiate-related drug overdose. Having a measure to simply  
11 detect higher than expected case counts is particularly worthwhile in public health  
12 research. In a setting where timely results based on possibly incomplete count data  
13 is a prime consideration, such as in syndromic surveillance (Heffernan et al. 2004),  
14 the scan statistic may be one’s first choice.

15 When population estimates are available, adjusting for areas that are more highly  
16 populated is appropriate, and the space-time scan statistic is a better, more precise  
17 measurement of cluster points. In our example (Fig. 27.2b), we were able to more  
18 accurately describe the location and statistical importance of clusters detected by the  
19 space–time permutation model. Again, when public health concerns are uppermost,  
20 this method may be particularly useful.

21 Bayesian methods may be most appropriate when potential explanatory variables  
22 are available and one’s interest is in assessing the *determinants* of health outcomes  
23 on a spatial level. It is important to appreciate that this is a smoothing method. When  
24 cluster detection is of uppermost concern, caution must be exercised that potential  
25 clusters are not smoothed away. While its appropriate utilization requires knowledge  
26 of MCMC and sample-based methods, full Bayesian analysis, as presented in our  
27 example, is not always necessary. Good empirical approximations are available and  
28 obviate the need to learn and use new complex statistical software (Devine et al.  
29 1994; Greenland 2006). But, when data are sparse and highly correlated and there  
30 is concern over noise obscuring spatial signals, a full Bayesian approach can help  
31 describe both the determinants and the patterns of the outcome of interest at a finer  
32 level.

33 It should be noted that we do not, in this discussion, dwell extensively on the  
34 implications of the results of the specific example we have used here to illustrate  
35 the material being described in this chapter. However, we considered the role of  
36 neighborhood-level SES in explaining rates of heroin-related overdose in the largest  
37 US urban area. Conceptual frameworks that consider the complex etiology of sub-  
38 stance use and its consequences (Galea, Rudenstine, and Vlahov 2005) have long  
39 suggested that a full consideration of the determination of substance use requires  
40 that we consider a range of individual- and group-level factors to understand popu-  
41 lation patterns of substance use. The hierarchical approach introduced here, suitably  
42 expanded, can be applied to test specific hypotheses and to pursue spatial etiologic  
43 questions, incorporating determinants at group and individual levels as necessary.

44 Ultimately, we note that the limitations of available data, including, for example,  
45 the use of zip codes as a neighborhood proxy, have been well discussed elsewhere in

01 the literature about the multilevel determination of population health (Osypuk and  
02 Galea 2007), and pertain just as much to spatial analyses as they do to all other types  
03 of epidemiologic analyses. Future work that makes use of the methods introduced  
04 here to address specific substance use-related etiologic hypotheses may benefit from  
05 application of these methods at different group levels of inference.

06 In conclusion, an appreciation for where and when health outcomes occur adds  
07 much to substance abuse epidemiology. A number of spatial and temporal methods  
08 are available, each with its own advantages and disadvantages in terms of complex-  
09 ity, data requirements, and underlying assumptions. The choice of a method will  
10 be driven by the question to be answered, data and software availability, and the  
11 intended audience or context in which the research is being conducted.

12  
13 **Acknowledgment** Dr. DiMaggio's efforts were supported by the US National Institute on Drug  
14 Abuse, grant number 1R03DA023431-01 and Centers for Disease Control and Prevention Health  
15 Protection Research Initiative Grant Number 1 K01CE000494-01. The efforts of Drs. Tardiff,  
16 Vlahov and Galea were supported in part by US National Institutes on Drug Abuse grants numbers  
17 DA06534 and DA017642.  
18  
19  
20  
21  
22  
23  
24  
25  
26  
27  
28  
29  
30  
31  
32  
33  
34  
35  
36  
37  
38  
39  
40  
41  
42  
43  
44  
45