1

Introduction

1.1 Why spatial and spatio-temporal statistics?

In the last few decades, the availability of spatial and spatio-temporal data has increased substantially, mainly due to the advances in computational tools which allow us to collect real-time data coming from GPS, satellites, etc. This means that nowadays in a wide range of fields, from epidemiology to ecology, to climatology and social science, researchers have to deal with geo-referenced data, i.e., including information about space (and possibly also time).

As an example, we consider a typical epidemiological study, where the interest is to evaluate the incidence of a particular disease such as lung cancer across a given country. The data will usually be available as counts of diseases for small areas (e.g., administrative units) for several years. What types of models allow the researchers to take into account all the information available from the data? It is important to consider the potential geographical pattern of the disease: areas close to each others are more likely to share some geographical characteristics which are related to the disease, thus to have similar incidence. Also how is the incidence changing in time? Again it is reasonable to expect that if there is a temporal pattern, this is stronger for subsequent years than for years further apart.

As a different example, let us assume that we are now in the climatology field and observe daily amount of precipitation at particular locations of a sparse network: we want to predict the rain amount at unobserved locations and we need to take into account spatial correlation and temporal dependency.

Spatial and spatio-temporal models are now widely used: typing "statistical models for spatial data" in ™Google Scholar returns more than 3 million hits and "statistical models for spatio-temporal data" gives about 159,000. There are countless scientific papers in peer review journals which use more or less complex and

Spatial and Spatio-temporal Bayesian Models with R-INLA, First Edition.
Marta Blangiardo and Michela Cameletti.

innovative statistical models to deal with the spatial and/or the temporal structure of the data in hand, covering a wide range of applications; the following list only aims at providing a flavor of the main areas where these types of models are used: Haslett and Raftery (1989), Handcock and Wallis (1994) and Jonhansson and Glass (2008) work in the meteorology field; Shoesmith (2013) presents a model for crime rates and burglaries, while Pavia *et al*. (2008) used spatial models for predicting election results; in epidemiology Knorr-Held and Richardson (2003) worked on infectious disease, while Waller *et al*. (1997) and Elliott *et al*. (2001) presented models for chronic diseases. Finally, Szpiro *et al*. (2010) focused on air pollution estimates and prediction.

1.2 Why do we use Bayesian methods for modeling spatial and spatio-temporal structures?

Several types of models are used with spatial and spatio-temporal data, depending on the aim of the study. If we are interested in summarizing spatial and spatio-temporal variation between areas using risks or probabilities then we could use statistical methods like disease mapping to compare maps and identify clusters. Moran Index is extensively used to check for spatial autocorrelation (Moran, 1950), while the scan statistics, implemented in SaTScan (Killdorf, 1997), has been used for cluster detection and to perform geographical surveillance in a non-Bayesian approach. The same types of models can also be used in studies where there is an aetiological aim to assess the potential effect of risk factors on outcomes.

A different type of study considers the quantification of the risk of experiencing an outcome as the distance from a certain source increases. This is typically framed in an environmental context, so that the source could be a point (e.g., waste site, radio transmitter) or a line (e.g., power line, road). In this case, the methods typically used vary from nonparametric tests proposed by Stone (1988) to the parametric approach introduced by Diggle *et al*. (1998).

In a different context, when the interest lies in mapping continuous spatial (or spatio-temporal) variables, which are measured only at a finite set of specific points in a given region, and in predicting their values at unobserved locations, geostatistical methods – such as kriging – are employed (Cressie, 1991; Stein, 1991). This may play a significant role in environmental risk assessment in order to identify areas where the risk of exceeding potentially harmful thresholds is higher.

Bayesian methods to deal with spatial and spatio-temporal data started to appear around year 2000, with the development of Markov chain Monte Carlo (MCMC) simulative methods (Casella and George, 1992; Gilks *et al*., 1996). Before that the Bayesian approach was almost only used for theoretical models and found little applications in real case studies due to the lack of numerical/analytical or simulative tools to compute posterior distributions. The advent of MCMC has triggered the possibility for researchers to develop complex models on large datasets without the need of imposing simplified structures. Probably the main contribution to spatial and spatio-temporal statistics is the one of Besag *et al*. (1991), who developed the Besag–York–Mollié (BYM) method (see Chapter 6) which is commonly used for disease mapping, while Banerjee *et al*. (2004), Diggle and Ribeiro (2007) and Cressie and Wikle (2011) have concentrated on Bayesian geostatistical models. The main advantage of the Bayesian approach resides in its taking into account uncertainty in the estimates/predictions, and its flexibility and capability of dealing with issues like missing data. In the book, we follow this paradigm and introduce the Bayesian philosophy and inference in Chapter 3, while in Chapter 4 we review Bayesian computation tools, but the reader could also find interesting the following: Knorr-Held (2000) and Best *et al*. (2005) for disease mapping and Diggle *et al*. (1998) for a modeling approach for continuous spatial data and for prediction.

1.3 Why INLA?

MCMC methods are extensively used for Bayesian inference, but their limitation resides in their computational burden. This has become an important issue, considering the advances in data collection, leading to availability of *big datasets*, characterized by high spatial and temporal resolution as well as data from different sources. The model complexity of taking into account spatial and spatio-temporal structures with large datasets could lead to several days of computing time to perform Bayesian inference via MCMC.

To overcome this issue, here comes the integrated nested Laplace approximations (INLA), a deterministic algorithm proposed by Rue *et al*. (2009) which has proven capable of providing accurate and fast results. It started as a stand-alone program but was then embedded into `R` (as a package called `R-INLA`), and since then it has become very popular amongst statisticians and applied researchers in a wide range of fields, with spatial and spatio-temporal models being possibly one of the main applications for it. The website `www.r-inla.org` provides a great resource of papers and tutorials and it contains a forum where users can post queries and requests of help. In this book we provide a detailed documentation of the `INLA` functions and options for modeling spatial and spatio-temporal data and use a series of examples drawn from epidemiology, social and environmental science.

1.4 Datasets

In this section, we briefly describe the datasets that we will use throughout the book. They are available for download from `R` packages or from the INLA website (`https://sites.google.com/a/r-inla.org/stbook/`), where we also provide the `R` code used to run all the examples.[1]

[1] From now onward, we use the `typewriter` font for computer code.

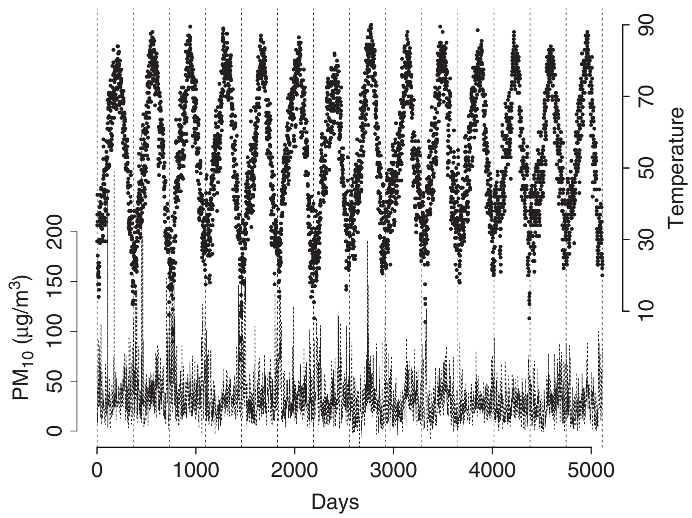

Figure 1.1 Daily temperature (points) and PM_{10} concentration (line) in Salt Lake City (1987–2000).

1.4.1 National Morbidity, Mortality, and Air Pollution Study

The National Morbidity, Mortality and Air Pollution Study (NMMAPS) is a large time series study to estimate the effect of air pollution on the health of individuals living in 108 US cities during the period 1987–2000. Several papers have been published on the data, methods, and results from this study (see, for instance, Samet *et al*. 2000). Detailed information about the database can be found on the Internet-based Health and Air Pollution Surveillance System (iHAPSS) website (`http://www.ihapss.jhsph.edu/`). Data on the daily concentration of particulates with an aerodynamic diameter of less than 10 (PM_{10}) and nitrogen dioxide (NO_2), both measured in $\mu g/m^3$, as well as daily temperature for Salt Lake City are contained in the file `NMMAPSraw.csv`.

We use this dataset to study the relationship between PM_{10} and temperature as an illustration of a linear regression model (Chapter 5). A plot which shows the trend of PM_{10} and temperature for the 14 years of available data is presented in Figure 1.1.

1.4.2 Average income in Swedish municipalities

Statistics Sweden (`http://www.scb.se/`) has created a population registry of Sweden, with detailed socioeconomic information at the individual and household level for all Swedish municipalities. This dataset was used by the EURAREA Consortium (EURAREA Consortium, 2004), a European research project funded by EUROSTAT, to investigate methods for small area estimation and their application. Gómez-Rubio *et al.* (unpublished) also used this dataset to illustrate how

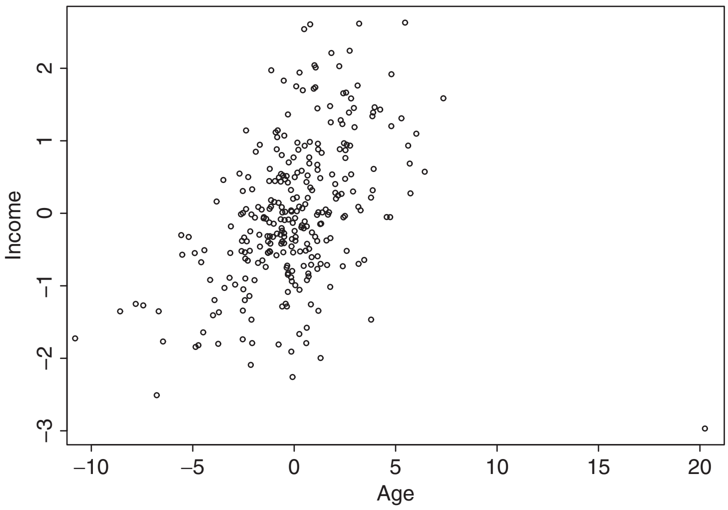

Figure 1.2 Relationship between the average age of the household and the average household income for 284 Swedish municipalities.

Bayesian hierarchical models can be implemented to provide good quality small area estimates and focused on the relationship between the average household income and the average age of the household heads. We are using a simulated version of the dataset[2] to show how to implement the Student *t* likelihood to deal with outliers in a linear regression model in Chapter 5.

The data are available in the file `income.csv` and contains the simulated average household income (`income`) and the average age for the head of the household (`age`) for 284 Swedish municipalities. Both variables are standardized. Figure 1.2 shows the relationship between the two variables.

1.4.3 Stroke in Sheffield

Maheswaran *et al*. (2006) analyzed the effect of outdoor modeled nitrogen oxide (NOx) levels, classified into quintiles, on stroke mortality in Sheffield (UK) between 1994 and 1998, using a Bayesian hierarchical model with spatial random effects. An association was observed between higher levels of NOx (in μg/m^3) and stroke mortality at the small area level (1030 enumeration districts, including on average 150 households). We use this dataset as an illustration of the Binomial generalized linear model in Chapter 5 and then of the hierarchical models in the same chapter.

The numbers of observed and expected stroke cases in each enumeration district, together with the NOx exposure and a measure of social deprivation (Townsend

[2] It was not possible to use the real data for privacy issues.

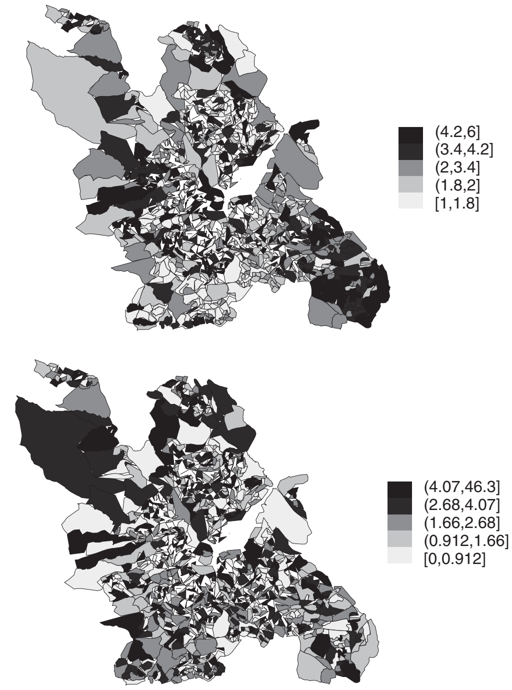

Figure 1.3 NOx concentration (top) and proportion of strokes registered per 1000 individuals (bottom) for enumeration districts in Sheffield (UK).

index), are available in the file `Stroke.csv`. Both variables are available in quintiles. Figure 1.3 shows the maps of NOx concentration (top) and the proportion of observed (O) strokes registered per 1000 individuals obtained as $\frac{\mathrm{O}_i}{\mathrm{Pop}_i} \times 1000$ (bottom) at the enumeration district level.

1.4.4 Ship accidents

McCullagh and Nelder (1989) used this dataset to study the rate of incidents in ships. The data are provided in the file `Ships.csv`. They include identification number (`id`), ship type (`type`), construction period (`built`), operation period (`oper`), and number of incidents (`y`). The natural logarithm of the number of months in operation is specified as the offset (`months`).

We use this dataset to illustrate the Poisson regression presented in Chapter 5. Figure 1.4 shows boxplots of the relationship between each predictor and the number of incidents (outcome).

1.4.5 CD4 in HIV patients

We consider simulated data from a clinical trial comparing two alternative treatments for HIV-infected individuals. 80 patients with HIV infection were randomly

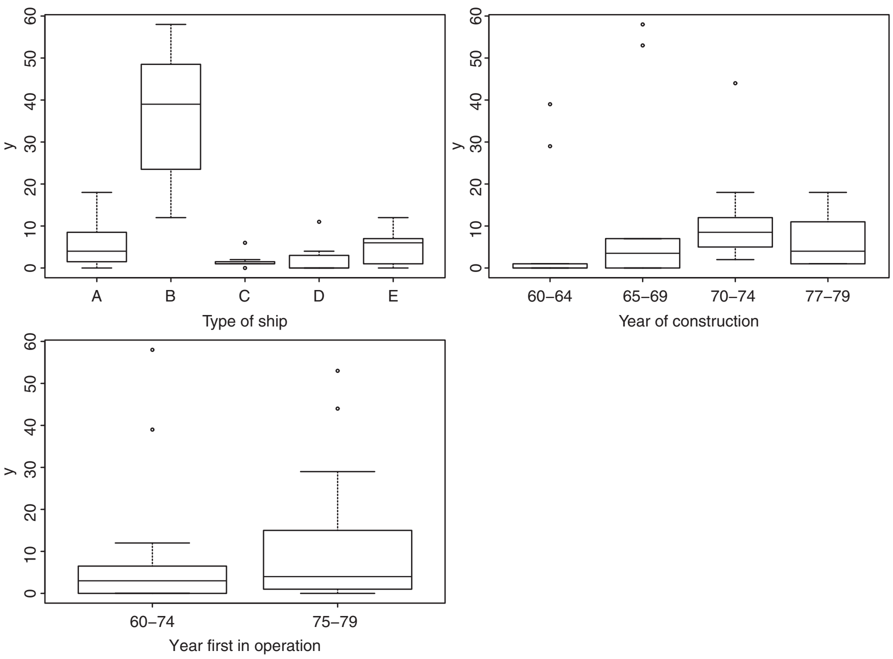

Figure 1.4 Boxplots of the relationship between each predictor and the number of accidents: type of ship (top left), year of construction (top right), and year when first in operation (bottom left).

assigned to one of two treatment groups: drug = 0 (didanosine, ddI) and drug = 1 (zalcitabine, ddC). Counts of CD4, cells commonly used in HIV positive patients as they are part of the immune system, were recorded at study entry (time $t = 0$) and again at 2, 6, and 12 months. An indicator of whether the patient had already been diagnosed with AIDS at study entry was also recorded (AIDS = 1 if patient diagnosed with AIDS, and 0 otherwise). The data can be found in the file `CD4.csv`.

We use this dataset to illustrate the hierarchical structure in regression models in Chapter 5. Figure 1.5 shows the distribution of CD4 for each patient, stratified by the use of drugs and the AIDS diagnosis.

1.4.6 Lip cancer in Scotland

Clayton and Kaldor (1987) analyzed the lip cancer rates in Scotland in the years 1975–1980 at the county level in order to evaluate the presence of an association between sun exposure and lip cancer. The example is part of the `GeoBUGS` Manual (Spiegelhalter *et al.* 1996; see `http://mathstat.helsinki.fi/openbugs/Manuals/GeoBUGS/Manual.html`) and is used here to illustrate hierarchical models (Chapter 5). The dataset is available as an `R` workspace (see Section 3.3), named `LipCancer.RData`, containing the number of counties

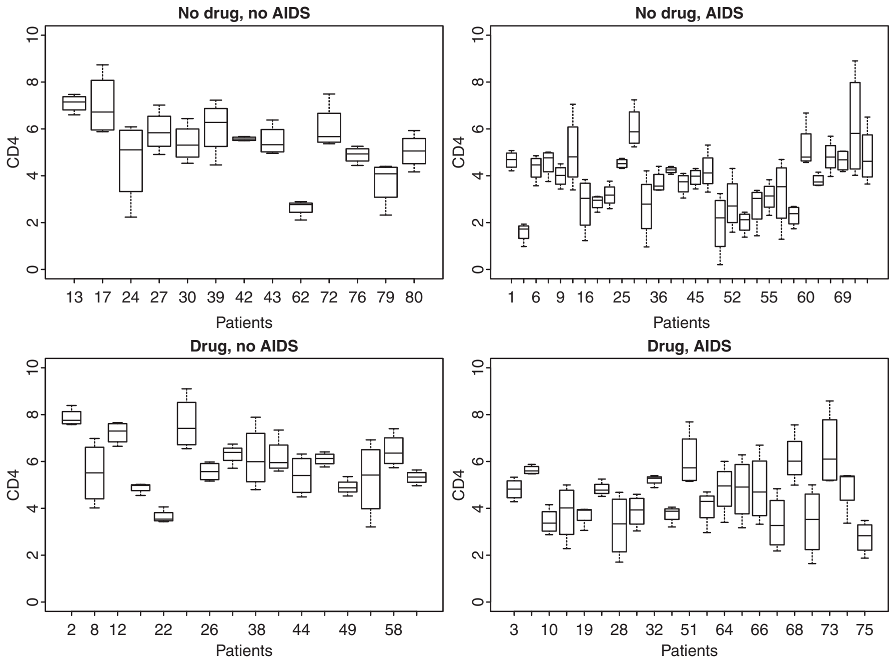

Figure 1.5 Boxplots for the CD4 counts of 80 HIV positive patients (at different time points), stratified by the drug use and the AIDS diagnosis.

($N = 56$), the observed cases of lip cancer in each county (O), the expected number of cases adjusted by the age and sex distribution of the population in the county (E), and the exposure variable (X), which measures the percentage of the population working in an outdoor environment (agriculture, fishing, or forestry), thus highly exposed to the sun. Figure 1.6 shows the map of the exposure variable (top) and of the standardized morbidity ratio ($\text{SMR}_i = \text{O}_i/\text{E}_i$, bottom) for the 56 counties.

1.4.7 Suicides in London

Congdon (2007) studied suicide mortality in 32 London boroughs (excluding the City of London) in the period 1989–1993 for male and female combined, using a disease mapping model and an ecological regression model.

We use this example to illustrate the intrinsic conditional autoregressive (iCAR) structure described in Chapter 6. The dataset is available as an R workspace, named `LondonSuicides.RData`, which contains the number of boroughs (N), the number of observed suicides in the period under study (O), the number of expected cases of suicides (E), an index of social deprivation (X1), and an index of social fragmentation (X2), which represents the lack of social connections and of sense of community.

Figure 1.6 Percentage of people working in an outdoor environment (top) and SMR for lip cancer (bottom) in the 56 Scottish counties.

Figure 1.7 shows the distribution of the social deprivation index (top), the social fragmentation index (center) and the SMR for suicides (bottom) in the 32 London boroughs.

1.4.8 Brain cancer in Navarra, Spain

Gómez-Rubio and Lopez-Quilez (2010) developed a statistical method to perform cluster detection on rare diseases and applied it to the study of the brain cancer incidence in Navarra (Spain), following the previous work of Ugarte *et al*. (2006). The data are available as an `R` workspace named `Navarre.RData`: the object `brainnav` contains the observed cases (`OBSERVED`), expected (`EXPECTED`)

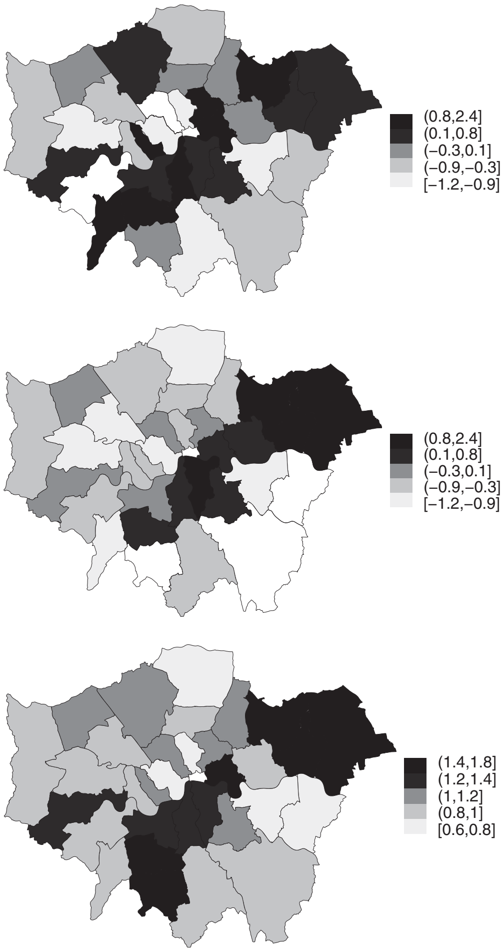

Figure 1.7 Distribution of social deprivation index (top), social fragmentation index (center), and SMR of suicide (bottom) in the 32 London boroughs.

ones, and SMR (`SMR`) of brain cancer in 1988–1994 for the 40 health districts in the Navarra region of Spain.

As the data contains a large proportion of zeros (32.5%), the standard Poisson model used for disease mapping is not appropriate, so we use this example to illustrate the zero inflated Poisson models (ZIP) in Chapter 6.

Figure 1.8 shows the SMR for brain cancer in the 40 health districts.

1.4.9 Respiratory hospital admission in Turin province

Atmospheric pollution is known to be associated with respiratory hospital admissions and mortality in small area studies (see, for instance, Sunyer *et al.*, 1997, 2003). We use an example on data for PM_{10} and hospital admissions for respiratory

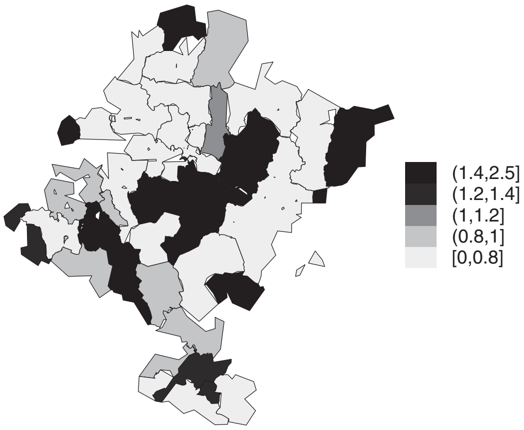

Figure 1.8 SMR of brain cancer for the 40 health districts in Navarra (Spain).

causes in the Turin province (Italy) in 2004 to introduce a zero inflated binomial regression model (ZIB) in Chapter 6, as the outcome variable is very sparse and the risk factor is spatially structured. The number of observed hospitalizations for respiratory causes at municipality level, the population for the same spatial units, and the average annual concentration of PM_{10} are available in the file `dataResp.csv`.

The map in Figure 1.9 (top) shows the distribution of the percentage of respiratory hospital admissions (over the total population) in the 315 municipalities considered in this example. The map in Figure 1.9 (bottom) shows the distribution of the average PM_{10} for the same period and the same areas.

1.4.10 Malaria in the Gambia

Diggle *et al*. (2002) studied the prevalence of malaria in children sampled from a village in the Gambia using generalized linear models. The dataset `gambia` (available as dataframe in the `geoR` package) contains data on eight variables and 2035 children living in 65 villages. The response variable (`pos`) is a binary indicator of the presence of malarial parasites in a blood sample. Other child level covariates are: age (`age`, in days), usage of bed nets (`netuse`), and information about whether the bed nets are treated with insecticide (`treated`). Village level covariates regard the vegetation index (`green`) and the inclusion or not of the village in the primary health care system (`phc`).

We use this example to illustrate the Bayesian kriging in Chapter 6 through the stochastic partial differential equation (SPDE) approach of Lindgren *et al*. (2011). The map in Figure 1.10 shows the Gambia region with the location of the villages.

1.4.11 Swiss rainfall data

In 1997 a statistical exercise named *The Spatial Interpolation Comparison 97 project* was organized by the Radioactivity Environmental Monitoring (Joint

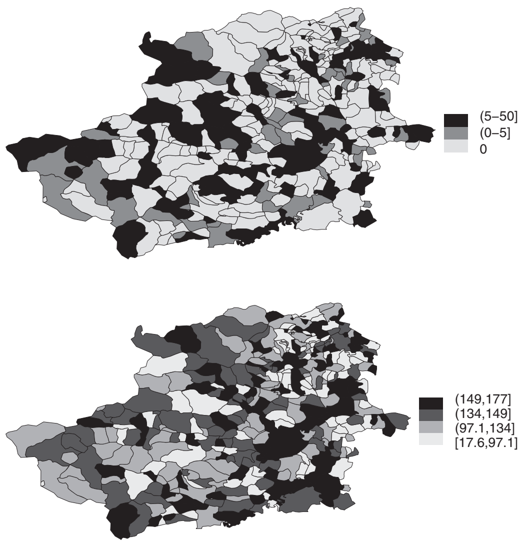

Figure 1.9 Distribution of the percentage of hospital admissions for respiratory causes in 2004 (over the total population) in the 315 municipalities in the Turin province, Italy (top). Map of the average PM_{10} concentration in 2004 for the 315 municipalities (bottom).

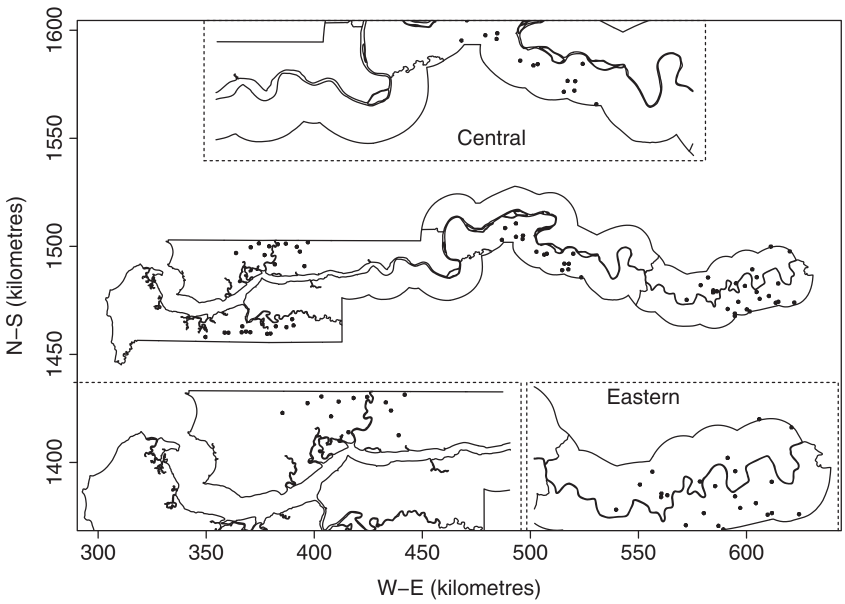

Figure 1.10 Map of the Gambia, Africa: the dots identifies villages where malaria prevalence has been recorded in children.

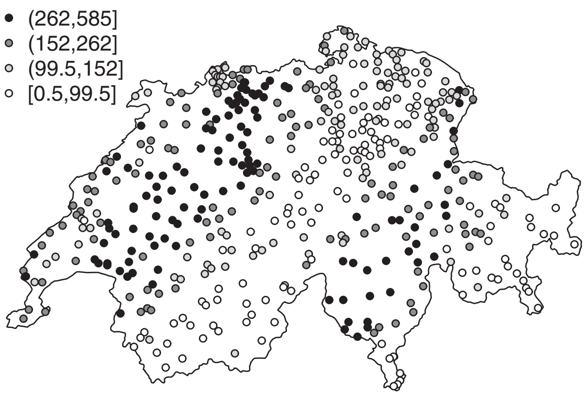

Figure 1.11 Rainfall data (in 10th of mm) collected on May 8, 1986, at 467 locations in Switzerland.

Research Centre, European Commission) to explore the impact of human factors in selecting and using spatial interpolation algorithms for mapping purposes (Dubois, 1998). The participants were asked to estimate daily rainfall values at 367 sites in Switzerland using 100 observed measurements (in 10th of mm) taken on May 8, 1986.

The data are included in the `geoR` library as an object named `SIC` which is formed by four `geodata` objects (a `geodata` object is a list with two obligatory arguments given by `coords` and `data`) denoted by `sic.all`, `sic.100`, `sic.367`, and `sic.some`, which differ in the number of spatial locations. Each object contains the following variables: location coordinates (`coords`), rainfall measurements (`data`), and elevation values (`altitude`). Additionally, a matrix named `sic.borders` with Switzerland borders is included.

The spatial distribution of the rainfall data measured at the 467 spatial locations is displayed in Figure 1.11. In Chapter 6, we use the rainfall data to illustrate spatial prediction (i.e., kriging) for a continuous spatial process.

1.4.12 Lung cancer mortality in Ohio

Lawson (2009) presented a space–time disease mapping model on lung cancer mortality in the Ohio counties (USA) for the years 1968–1988. We use the same dataset here to illustrate the parametric spatio-temporal disease mapping approach in Chapter 7.

The data are stored in the `OhioRespMort.csv` file, which consists of a matrix of (88×21) rows (counties $\times$ years) and six columns with the name and the ID of the county (`NAME` and `county`), the year (`year`), the number of deaths (`y`), the number of exposed individuals (`n`), and the expected number of deaths (`E`).

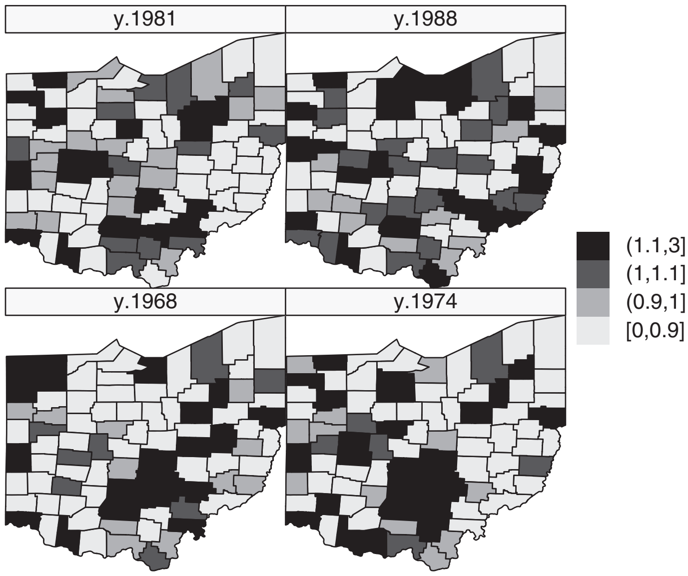

Figure 1.12 Distribution of standardized mortality rates of lung cancer in 88 counties in Ohio (USA) during 1968, 1974, 1981, and 1988.

Figure 1.12 displays the distribution of standardized mortality ratios ($\text{SMR}_i = y_{it}/E_{it}$) of respiratory cancer deaths for four years.

1.4.13 Low birth weight births in Georgia

Lawson (2009) considered counts of very low birth weight (<1500 g) in the counties of Georgia (USA) for the years 1994–2004 to perform spatio-temporal disease mapping. Here we consider counts of low birth weight (<2500 g) for the 159 counties of Georgia during 2000–2010 in order to illustrate the spatio-temporal Poisson nonparametric approach in Chapter 7.

The data were obtained from the Georgia Department of Public Health website through the OASIS web query system (`http://oasis.state.ga.us/`) and are stored in the `Lowbirthweight_births.csv` and `Total_births.csv` files, which contain for each county and year the number of low birth weight births and the total number of births, respectively.

Figure 1.13 displays the distribution of standardized incidence ratios of low birth weight for the 11 considered years.

1.4.14 Air pollution in Piemonte

Cameletti *et al*. (2011) analyzed PM_{10} concentration measured in the Piemonte region (Northern Italy) during October 2005–March 2006. The data come from a monitoring network composed of 24 stations.

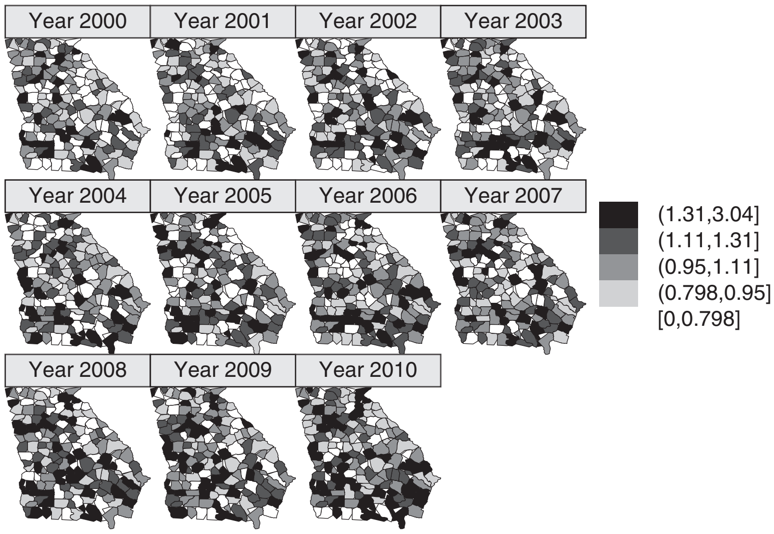

Figure 1.13 Distribution of standardized incidence rate of low birth weight births in 159 counties in Georgia (USA) during 2000–2010.

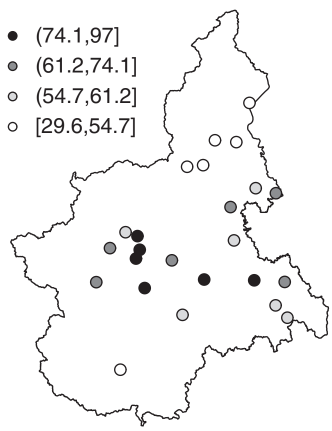

Figure 1.14 Average PM_{10} ($\mu g/m^3$) concentration during October 2005–March 2006 for the 24 monitoring stations in the Piemonte region (Northern Italy).

The data are stored in the `Piemonte_data_byday.csv` file which contains daily PM_{10} concentration (`PM10`, in $\mu g/m^3$) and some covariates: daily maximum mixing height (`HMIX`, in m), daily total precipitation (`PREC`, in mm), daily mean wind speed (`WS`, in m/s), daily mean temperature (`TEMP`, in K), daily emission rates of primary aerosols (`EMI`, in g/s), altitude (`A`, in m), and spatial coordinates (`UTMX` and `UTMY`, in km).

We use this dataset in Chapter 7 for implementing a spatio-temporal model with covariates to predict PM_{10} concentration all over the Piemonte region. Moreover, we deal with the so-called change of support problem in order to get concentration predictions at a lower scale given by health districts where mortality data are available. Figure 1.14 shows the average PM_{10} concentration computed over the period October 2005–March 2006 for the 24 monitoring stations in the Piemonte region.

References

Banerjee, S., Carlin, B., and Gelfand, A. (2004). *Hierarchical Modeling and Analysis for Spatial Data*. CRC.

Besag, J., York, J., and Mollie, A. (1991). Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, **43**, 1–59.

Best, N., Richardson, S., and Thompson, A. (2005). A comparison of Bayesian spatial models for disease mapping. *Statistical Methods in Medical Research*, **14**(1), 35–59.

Cameletti, M., Ignaccolo, R., and Bande, S. (2011). Comparing spatio-temporal models for particulate matter in Piemonte. *Environmetrics*, **22**, 985–996.

Casella, G. and George, E. (1992). Explaining the Gibbs sampler. *American Statistician*, **46**, 167–174.

Clayton, D. and Kaldor, J. (1987). Empirical Bayes estimates of age-standardized relative risks for use in disease mapping. *Biometrics*, **43**, 671–681.

Congdon, P. (2007). *Bayesian Statistical Modelling*. John Wiley and Sons, Ltd.

Cressie, N. (1991). *Statistics for Spatial Data*. Wiley.

Cressie, N. and Wikle, C. (2011). *Statistics for Spatio-Temporal Data*. Wiley.

Diggle, P. and Ribeiro, J. P. (2007). *Model-based Geostatistics*. Springer.

Diggle, P., Moyeed, R., and Tawn, J. (1998). Model-based geostatistics. *Journal of the Royal Statistical Society, Series C*, **47**, 299–350.

Diggle, P., Moyeed, R., Rowlingson, B., and Thomson, M. (2002). Childhood Malaria in the Gambia: A case-study in model-based geostatistics. *Journal of the Royal Statistical Society, Series C (Applied Statistics)*, **51**(4), 493–506.

Dubois, G. (1998). Spatial interpolation comparison 97: Foreword and introduction. *Journal of Geographic Information and Decision Analysis*, **2**, 1–10.

Elliott, N., Wakefield, J., Best, N., and Briggs, D., editors (2001). *Spatial Epidemiology*. Oxford University Press.

EURAREA Consortium (2004). *Project reference volume*. EURAREA Consortium. `http://www.ons.gov.uk/ons/guide-method/method-quality/general-methodology/spatial-analysis-and-modelling/eurarea/index.html`.

Gilks, W., Richardson, S., and Spiegelhalter, D. (1996). *Markov Chain Monte Carlo in Practice*. Chapman & Hall/CRC.

Gómez-Rubio, V. and Lopez-Quilez, A. (2010). Statistical methods for the geographical analysis of rare diseases. *Advances in Experimental Medicine and Biology*, **686**, 151–171.

Handcock, M. and Wallis, J. (1994). An approach to statistical spatial-temporal modelling of meteorological fields. *Journal of the American Statistical Association*, **86**, 368–390.

Haslett, J. and Raftery, A. (1989). Space time modelling with long memory dependence: Assessing Ireland's wind power resource. *Journal of Applied Statistics*, **38**(1), 1–50.

Jonhansson, M. and Glass, G. (2008). High-resolution spatio-temporal weather models for climate studies. *International Journal of Health Geographics*, **7**(52), 1–8.

Killdorf, M. (1997). A spatial scan statistics. *Communications in Statistics: Theory and Methods*, **26**, 1481–1496.

Knorr-Held, L. (2000). Bayesian modelling of inseparable space-time variation in disease risk. *Statistics in Medicine*, **19** (17–18), 2555–2567.

Knorr-Held, L. and Richardson, S. (2003). A hierarchical model for space-time surveillance data on meningococcal disease incidence. *Journal of the Royal Statistical Society, Series C*, **52**(2), 169–183.

Lawson, A. (2009). *Bayesian Disease Mapping. Hierarchical Modeling in Spatial Epidemiology*. CRC Press.

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach (with discussion). *Journal of Royal Statistical Society, Series B*, **73**(4), 423–498.

Maheswaran, R., Haining, R., Pearson, T., Law, J., Brindley, P., and Best, N. (2006). Outdoor NOx and stroke mortality: Adjusting for small area level smoking prevalence using a Bayesian approach. *Statistical Methods in Medical Research*, **15**(5), 499–516.

McCullagh, P. and Nelder, J. (1989). *Generalized Linear Models, Second Edition*. Chapman & Hall.

Moran, P. (1950). Notes on continuous stochastic phenomena. *Biometrika*, **37**(1), 17–23.

Pavia, J., Larraz, B., and Montero, J. (2008). Election forecasts using spatio-temporal models. *Journal of the American Statistical Association*, **103**(483), 1050–1059.

Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian model by using integrated nested Laplace approximations (with discussion). *Journal of Royal Statistical Society, Series B*, **71**, 319–392.

Samet, J., Dominici, F., Zeger, S., Schwartz, J., and Dockery, D. (2000). The National Morbidity, Mortality, and Air Pollution Study, Part I: Methods and Methodologic Issues. *Research Report of the Health Effects Institute*, **94**.

Shoesmith, G. (2013). Space time autoregressive models and forecasting national, regional and state crime rates. *International Journal of Forecasting*, **29**(1), 191–201.

Spiegelhalter, D., Thomas, A., Best, N., and Gilks, W. (1996). BUGS: Bayesian Inference Using Gibbs Sampling, Version 0.5, (version ii). MRC Biostatistics Unit.

Stein, M. (1991). *Interpolation of Spatial Data: Some Theory of Kriging*. Springer.

Stone, R. (1988). Investigations of excess environmental risks around putative sources: Statistical problems and a proposed test. *Statistics in Medicine*, **7**, 649–660.

Sunyer, J., Spix, C., Quenel, P., Ponce-de Leon, A., Ponka, A., Barumandzadeh, T., Touloumi, G., Bacharova, L., Wojtyniak, B., Vonk, J., Bisanti, L., Schwartz, J., and Katsouyanni, K. (1997). Urban air pollution and emergency admissions for asthma in four European cities: the APHEA Project. *Thorax*, **52**(9), 760–765.

Sunyer, J. Atkinson, R., Ballester, F., Le Tertre, A., Ayres, J. G., Forastiere, F., Forsberg, B., Vonk, J. M., Bisanti, L., Anderson, R. H., Schwartz, J., Katsouyanni, K. (2003). Respiratory effects of sulphur dioxide: a hierarchical multicity analysis in the APHEA 2 study. *Occupation Environmental Medicine*, **60**(8), e2.

Szpiro, A., Sampson, P., Sheppard, L., Lumley, T., Adar, S., and Kaufman, J. (2010). Predicting intra-urban variation in air pollution concentrations with complex spatio-temporal dependencies. *Environmetrics*, **21**(6), 606–631.

Ugarte, M., Ibanez, B., and Militino, F. (2006). Modelling risks in disease mapping. *Statistical Methods in Medical Research*, **15**, 21–35.

Waller, L., Carlin, B., Xia, H., and Gelfand, A. (1997). Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association*, **92**(438), 607–617.