



[Lee, D.](#) (2023) Identifying boundaries in spatially continuous risk surfaces from spatially aggregated disease count data. *Annals of Applied Statistics*, 17(4), pp. 3153-3172. (doi: [10.1214/23-AOAS1755](https://doi.org/10.1214/23-AOAS1755))

There may be differences between this version and the published version.  
You are advised to consult the published version if you wish to cite from it.

<https://eprints.gla.ac.uk/293873/>

Deposited on 7 March 2023

Enlighten – Research publications by members of the University of Glasgow

<http://eprints.gla.ac.uk>

# IDENTIFYING BOUNDARIES IN SPATIALLY CONTINUOUS RISK SURFACES FROM SPATIALLY AGGREGATED DISEASE COUNT DATA

BY DUNCAN LEE<sup>1</sup>,

<sup>1</sup>*School of Mathematics and Statistics, University of Glasgow, [Duncan.Lee@glasgow.ac.uk](mailto:Duncan.Lee@glasgow.ac.uk)*

Spatially aggregated disease count data relating to a set of non-overlapping areal units are often used to make inference on population-level disease risk. This includes the identification of risk boundaries, which are locations where there is a sizeable change in risk between geographically neighbouring areal units. Existing studies provide spatially discrete inference on the areal unit footprint, which forces the boundaries to coincide with the entire geographical border between neighbouring units. This paper is the first to relax these assumptions, by estimating disease risk and the locations of risk boundaries on a grid of square pixels covering the study region that can be made arbitrarily small to approximate a spatially continuous surface. We propose a two-stage approach that first fits a Bayesian spatio-temporal realignment model to estimate disease risk at the grid-level, and then identifies boundaries in this surface using edge detection algorithms from computer vision. This novel methodological fusion is motivated by a new study of respiratory hospitalisation risk in Glasgow, Scotland between 2008 and 2017, and we identify numerous risk boundaries across the city.

**1. Introduction.** Population-level disease risk varies between different communities, and modelling this spatio-temporal variation informs public health policy. Examples include identifying the locations of high-risk sub-regions ([Wakefield and Kim, 2013](#)), and quantifying the changing magnitude of health inequalities ([Mackenbach et al., 2018](#)). This research utilises disease count data that have been spatially aggregated to a set of non-overlapping areal units such as census tracts, because individual level disease incidence data are not publicly available. The models commonly applied to these data encourage neighbouring areal units and time periods to have similar disease risks, which is achieved using spatio-temporally autocorrelated random effects within a Bayesian hierarchical model (see for example [Knorr-Held, 2000](#) and [Rushworth, Lee and Mitchell, 2014](#)). However, disease risk surfaces are likely to contain areas of gradual change that are separated by locations exhibiting steep risk gradients, the latter being known as *risk boundaries* ([Lu and Carlin, 2005](#)). Such boundaries may reflect the border where two different communities meet, and can be caused by physical barriers such as rivers or railway lines that prevent two communities mixing ([Mitchell and Lee, 2014](#)). The estimation of risk boundaries allows researchers to identify the spatial extent of a high-risk region to aid the targeting of health resources, as well as hypothesis generation about putative risk factors that may influence disease risk and vary on each side of the boundary.

Identifying boundaries (or step-changes) in data surfaces has been investigated in time series analysis ([Fisher et al., 2022](#)) and spatial statistics ([Bradley, Wikle and Holan, 2017](#)), with the latter including point-level, pixel-level and areal unit level data. For example, [Banerjee and Gelfand \(2006\)](#) propose curvilinear boundary detection for point-level data, [Qu, Bradley and Niu \(2021\)](#) extend this to multivariate and multiscale data, while [Syring and Li \(2017\)](#) identify boundaries in an image of pixels. For areal unit level data [Ma, Carlin and Banerjee](#)

---

*Keywords and phrases:* Bayesian inference, Disease risk modelling, Edge detection algorithms, Spatially continuous inference.

(2010), Lee and Mitchell (2012), Berchuck, Mwanza and Warren (2019) and Lee, Meeks and Pettersson (2021) identify boundaries by jointly modelling disease risk and the underlying neighbourhood structure of the areal units utilising conditional autoregressive (CAR) priors.

This paper aims to estimate the spatial locations and temporal evolution of boundaries in the risk of hospitalisation due to respiratory disease across the city of Glasgow, Scotland, so that the spatial extents of high-risk areas can be identified allowing resources to be targeted at areas with the greatest need. Glasgow is chosen for our study because it exhibits large health inequalities (<http://www.healthscotland.scot/health-inequalities/>), which has led Lee and Mitchell (2013) and Lee, Meeks and Pettersson (2021) to identify numerous risk boundaries between geographically adjacent areal units. However, as these studies undertake inference on the areal unit level they implicitly assume that: (i) disease risk is constant within each areal unit; and (ii) the risk boundaries must coincide with the geographical borders separating neighbouring areal units. That is, boundaries are not allowed to run through the interior of an areal unit. The areal units used in this study and in the two previous studies listed above are Intermediate Zones (IZ), which have an average population of around 4,000 people. Assumptions (i) and (ii) may be unrealistic for areal units comprising around 4,000 people, and if they are then this would adversely affect the accuracy of the risks and boundaries estimated by Lee and Mitchell (2013) and Lee, Meeks and Pettersson (2021). Furthermore, the IZ borders were spatially redrawn in 2012 following the 2011 UK census, which means that areal unit level inference, such as that conducted by the above studies would not be comparable before and after this boundary change.

This paper proposes the first approach for solving these problems, by estimating disease risk on an approximately spatially continuous scale and identifying risk boundaries at approximately any location in the study region using only spatially aggregated data. It thus provides spatially consistent inference when the areal units have differing spatial supports over time, which addresses the limitations of the existing Glasgow respiratory hospitalisation studies described above. Our two-stage approach first uses Bayesian spatial realignment methodology to estimate disease risk on a temporally consistent grid of pixels covering the study region, and then identifies boundaries in these gridded risk surfaces using edge detection methods borrowed from computer vision (see Muntarina, Shorif and Uddin, 2022 for a review). The pixels can be made arbitrarily small to approximate a spatially continuous surface, which thus allows risk boundaries to be identified at approximately any location in the study region. This approach is the first to fuse spatial re-alignment count data models with edge detection methodology, and we are not aware of existing methodologies that take in areal unit level count data (with varying borders over time) as the input, and produce a spatially continuous risk surface with the locations of boundaries identified. We note though that spatial re-alignment count data models have a rich history, including areal interpolation (e.g. Flowerdew and Green, 1989) and model-based approaches (e.g. Mugglin and Carlin, 1998, Bradley, Wikle and Holan, 2016, Taylor, Andrade-Pacheco and Sturrock, 2018, and Gramatica, Congdon and Liverani, 2021). The motivating Glasgow study is described in Section 2, our novel methodological fusion is proposed in Section 3, while its efficacy is assessed in Section 4 via simulation. The results of the Glasgow study are presented in Section 5, while Section 6 concludes the paper.

## 2. Motivating study.

2.1. *Aims and questions of interest.* This paper is motivated by a new study of respiratory hospitalisation risk in Glasgow, Scotland, between 2008 and 2017, which extends Lee and Mitchell (2013) and Lee, Meeks and Pettersson (2021) by using data for a longer time period

that crosses the IZ border changes in 2012. The overarching aim is to estimate the spatio-temporal variation in disease risk and the locations of risk boundaries across the city, and we specifically focus on the following questions.

1. Where are the most prominent boundaries that persist over the entire study period?
2. Has the spatial variation in risk and the locations of risk boundaries changed over time, and how has this impacted health inequalities across the city?
3. How reasonable are the implicit assumptions made by existing studies that disease risk is constant within an IZ and risk boundaries coincide with the borders separating neighbouring Intermediate Zones?

Answering 1. and 2. allows us to identify the spatial extents of the main high-risk sub-regions, which informs public health policy as to where limited health resources should be targeted and the extent to which this has changed over time. Answering 3. allows us to examine how reasonable the assumptions made by existing IZ-level modelling approaches are. For example, if within IZ variation in disease risk is small and risk boundaries largely correspond to IZ borders, then IZ-level inference would be appropriate. However, if this is not the case then it suggests that IZ-level inference is not appropriate for these data. Finally, we note that existing methods such as [Lee, Meeks and Pettersson \(2021\)](#) are not able to answer questions 1 - 3 for our study, firstly because of the IZ border changes in 2012 causing spatial misalignment in the areal units over time, and secondly because it only delivers IZ-level inference.

*2.2. Respiratory hospitalisation data.* The disease data are yearly spatially aggregated counts of the numbers of individuals admitted to hospital due to respiratory disease between 2008 and 2017 for each Intermediate Zone in Glasgow. The boundaries of the IZs were spatially redrawn in 2012 following the 2011 UK census, meaning that the first half of the study period uses the 2001 definition of IZs, while the second half uses the 2011 definition of IZs. Thus for the 5-year period 2008 to 2012 the disease counts  $\{Y_t(\mathcal{A}_k^{(1)})\}$  relate to  $K^{(1)} = 133$  IZs denoted by  $\mathcal{S}^{(1)} = \{\mathcal{A}_1^{(1)}, \dots, \mathcal{A}_{K^{(1)}}^{(1)}\}$ , where  $k$  denotes the IZ and  $t$  denotes the year. Similarly, for the 5-year period 2013 to 2017 the disease counts  $\{Y_t(\mathcal{A}_k^{(2)})\}$  relate to  $K^{(2)} = 136$  IZs denoted by  $\mathcal{S}^{(2)} = \{\mathcal{A}_1^{(2)}, \dots, \mathcal{A}_{K^{(2)}}^{(2)}\}$ . These disease counts range from 17 to 282, with a median value of 91. The spatial footprints of the IZs are shown in panel (A) of Figure 1 as red (2001 IZs) and blue (2011 IZs) lines, which shows that the number of discordant boundaries is relatively small. The disease counts depend on the size and age-sex demographics of the population living in each IZ during each year, which we adjust for by computing the expected numbers of hospitalisations denoted by  $e_t(\mathcal{A}_k^{(1)})$  for  $t = 1, \dots, 5$  and  $e_t(\mathcal{A}_k^{(2)})$  for  $t = 6, \dots, 10$ . Specifically, the population in each IZ is partitioned into age-sex strata, and city-wide strata specific rates of respiratory hospitalisation are multiplied by the population totals and then summed over strata to compute the final expected count.

From these data an exploratory measure of disease risk is the standardised morbidity ratio (SMR) defined as  $\text{SMR}_t(\mathcal{A}_k^{(1)}) = Y_t(\mathcal{A}_k^{(1)})/e_t(\mathcal{A}_k^{(1)})$  for  $t = 1, \dots, 5$  and  $\text{SMR}_t(\mathcal{A}_k^{(2)}) = Y_t(\mathcal{A}_k^{(2)})/e_t(\mathcal{A}_k^{(2)})$  for  $t = 6, \dots, 10$ . IZs with an SMR that is respectively greater / less than one indicate respectively higher / lower risks than the Glasgow average during the study period, where a value of 1.2 corresponds to a 20% increased risk. The spatial trend in the SMR for 2017 is displayed in panel (C) of Figure 1, which shows that a number of regions have sizeable differences in their risks compared to their geographical neighbours, suggesting the presence of boundaries in the risk surface. The temporal trend in the SMR is shown in panel (D) of Figure 1, where jittering has been added to the Year direction to improve the visibility

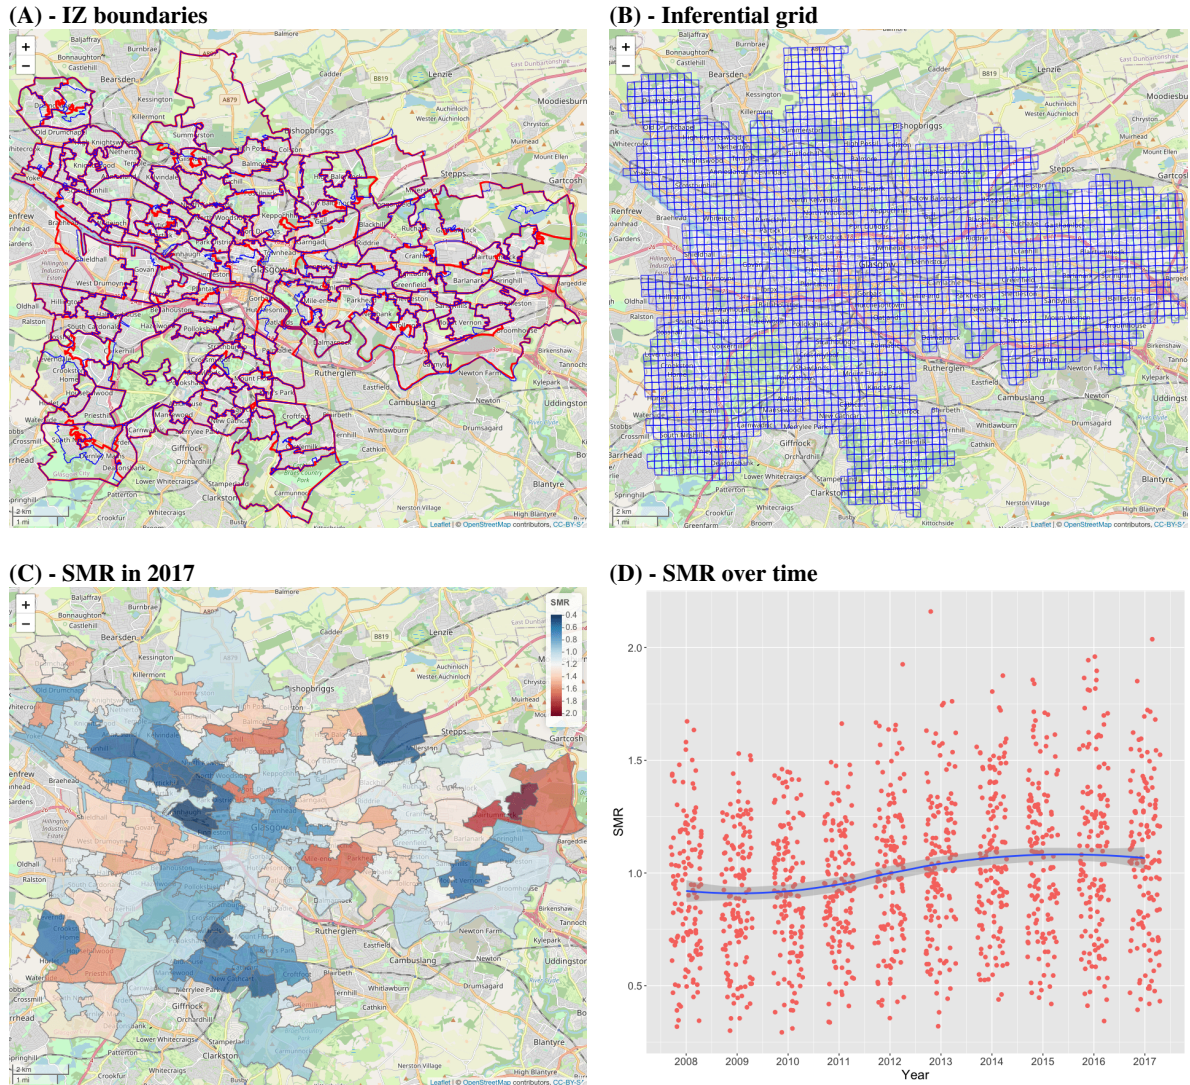


FIG 1. Summary of the study region and data: (A) - the IZ boundaries in 2001 (red) and 2011 (blue); (B) - the  $250m^2$  inferential grid; (C) - the spatial pattern in the SMR in 2017; and (D) - the temporal trend and variation in the SMR. In panel (D) the SMR values have been jittered in the horizontal (Year) direction to improve the presentation, and the blue line is a LOESS trend.

of the points, and a trend line has been estimated using LOESS smoothing. The SMR exhibits a small increasing trend in risk, with mean risks of 0.93 in 2008 compared to 1.06 in 2017.

**2.3. Spatial inferential grid.** Our modelling framework estimates disease risk for each year on a grid of  $M = 2,923$  square pixels  $\mathcal{S} = \{\mathcal{H}_{ij}\}$  covering the study region (shown in panel (B) of Figure 1), where  $\mathcal{H}_{ij}$  is the pixel in the  $i$ th row and  $j$ th column. These pixels can be made arbitrarily small to approximate a spatially continuous surface, thus allowing risk boundaries to be identified approximately anywhere in the study region. The pixels have sides of length  $250m$ , which is a trade-off between the granularity of the estimated risk surface and the computational overhead of fitting the model.

2.4. *Covariate information.* It would be desirable to have covariate information at the pixel level to aid in estimating the sub IZ variation in respiratory hospitalisation risk. A prime candidate in this regard would be measures of smoking prevalence or social-economic deprivation that typically correlate well with respiratory disease risk (see for example [Mack-enbach et al., 2018](#)), but in Scotland small-area statistics such as these are not available on a gridded footprint, precluding the inclusion of such data in the models. Environmental and topological data are often available on a gridded footprint and so could in theory be used to inform a grid-level model. However, they are unlikely to exhibit strong relationships with the health outcome in our study, so are not considered here. Additionally, even if covariates with strong links to respiratory hospitalisation were available, then the accuracy of covariate effect estimates when the response is being spatially re-scaled has not been well studied, meaning that any resulting effect estimates would need to be interpreted with caution.

**3. Methodology.** Let the set of spatially aggregated observed and expected disease counts for time  $t$  and areal unit  $k$  be denoted by  $\{Y_t(\mathcal{A}_k^{(t)}), e_t(\mathcal{A}_k^{(t)})\}$  for  $t = 1, \dots, N$ , where the set of  $K^{(t)}$  areal units at time  $t$  change over time. We use these data to estimate disease risk and identify risk boundaries on a spatially-continuous footprint approximated by the grid of pixels  $\mathcal{S} = \{\mathcal{H}_{ij}\}$ . Our two-stage approach consists of a Bayesian hierarchical spatial re-alignment model that estimates disease risk on this gridded footprint (stage 1), and the application of edge detection methodologies from computer vision to identify risk boundaries (stage 2).

### 3.1. Stage 1 - Estimating grid-level disease risk.

3.1.1. *Computing the grid-level expected disease counts.* We first estimate the expected numbers of hospitalisations at the grid-level  $\{e_t(\mathcal{H}_{ij})\}$  from the known areal unit level quantities  $\{e_t(\mathcal{A}_k^{(t)})\}$  using area-weighted interpolation ([Flowerdew and Green, 1993](#)). Letting  $a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})$  denote the area of intersection between areal unit  $\mathcal{A}_k^{(t)}$  and pixel  $\mathcal{H}_{ij}$ , the expected counts are estimated as

$$(1) \quad e_t(\mathcal{H}_{ij}) = \sum_{k=1}^{K^{(t)}} e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = \sum_{k=1}^{K^{(t)}} \frac{a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})}{\sum_{r,s} a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{rs})} e_t(\mathcal{A}_k^{(t)}).$$

Here  $e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})$  denotes the expected number of hospitalisations at time  $t$  from the population who live in the intersection area  $\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}$ , and the first equality sums these contributions over all  $K^{(t)}$  areal units. The second equality estimates  $e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})$  by allocating the expected hospitalisations from  $e_t(\mathcal{A}_k^{(t)})$  to each pixel  $\mathcal{H}_{ij}$  proportionally to its area of intersection with  $\mathcal{A}_k^{(t)}$ . Thus trivially  $\sum_{i,j} e_t(\mathcal{H}_{ij}) = \sum_{k=1}^{K^{(t)}} e_t(\mathcal{A}_k^{(t)})$  as would be desired, and most  $e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = 0$  because most  $a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = 0$ . Finally, as these are expected rather than observed hospitalisation counts computed using indirect standardisation, they do not need to be integer valued.

This approach assumes that the density of the expected hospitalisation counts is constant within each IZ, which may be an unrealistic assumption. Such an assumption could potentially be relaxed using gridded population data split by age-sex strata, such as that available from WorldPop (<https://hub.worldpop.org/>). However, such data are not used here because they are only estimates and hence likely to contain errors, which may induce errors into the estimated grid-level risks. We note that age-sex splits in the population data

are needed because respiratory hospitalisation rates vary dramatically by age and sex, and so only knowing the overall population total / density in each grid square would only be appropriate if the age-sex demographics of the population did not vary by area which is not the case in Glasgow.

3.1.2. *Grid-level risk model.* The grid-level risk model uses data  $\{Y_t(\mathcal{H}_{ij}), e_t(\mathcal{H}_{ij})\}$ , and while  $\{e_t(\mathcal{H}_{ij})\}$  are computed as above, the disease counts  $\{Y_t(\mathcal{H}_{ij})\}$  are unknown and treated as additional random variables in our Bayesian model fitting algorithm (see 3.1.3). The first level of the model has the following Poisson log-linear form:

$$(2) \quad \begin{aligned} Y_t(\mathcal{H}_{ij}) &\sim \text{Poisson}[e_t(\mathcal{H}_{ij})\theta_t(\mathcal{H}_{ij})] \\ \ln[\theta_t(\mathcal{H}_{ij})] &= \mathbf{x}_t(\mathcal{H}_{ij})^\top \boldsymbol{\beta} + \phi_t(\mathcal{H}_{ij}), \end{aligned}$$

where  $\{\theta_t(\mathcal{H}_{ij})\}$  are the grid-level model-based estimates of disease risk. The spatio-temporal variation in risk can be modelled by covariate  $(\mathbf{x}_t(\mathcal{H}_{ij})^\top \boldsymbol{\beta})$  and random effect  $(\phi_t(\mathcal{H}_{ij}))$  components, where  $\mathbf{x}_t(\mathcal{H}_{ij})$  denotes a  $p \times 1$  vector of covariates and  $\boldsymbol{\beta}$  are the associated regression parameters. Each regression parameter is assigned an independent weakly informative zero-mean Gaussian prior distribution with variance 100,000, which allows the data to speak for themselves when estimating its value. However, in the motivating case study we only include an intercept term in the model as grid-level covariates are not available.

The random effects  $\{\phi_t(\mathcal{H}_{ij})\}$  model the spatio-temporal variation in disease risk after covariate adjustment, and numerous prior structures have been proposed in the literature including spatial and temporal main effects and a spatio-temporal interaction (Knorr-Held, 2000), spatially correlated linear time trends (Bernardinelli et al., 1995), separate spatial surfaces at each time period (Waller et al., 1997), and a temporally evolving spatially correlated surface (Rushworth, Lee and Mitchell, 2014). Here we adopt the latter specification because the motivating case study pertains to mostly chronic and non-infectious respiratory disease cases such as COPD and lung cancer, and hence disease risk is likely to evolve slowly over time. Therefore partition the random effects  $\phi = \{\phi_t(\mathcal{H}_{ij})\}$  into  $\phi = (\phi_1, \dots, \phi_N)$ , where  $\phi_t = \{\phi_t(\mathcal{H}_{ij})\}$  denotes the vector of random effects relating to time period  $t$  for all pixels  $\{\mathcal{H}_{ij}\}$ . We consider the following first and second order autoregressive processes for  $\{\phi_t\}$ , the former being proposed by Rushworth, Lee and Mitchell (2014).

#### First order autoregressive process - AR(1)

$$(3) \quad \begin{aligned} \phi_t | \phi_{t-1} &\sim N(\alpha \phi_{t-1}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}) \quad \text{for } t = 2, \dots, N \\ \phi_1 &\sim N(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}). \end{aligned}$$

#### Second order autoregressive process - AR(2)

$$(4) \quad \begin{aligned} \phi_t | \phi_{t-1}, \phi_{t-2} &\sim N(\alpha_1 \phi_{t-1} + \alpha_2 \phi_{t-2}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}) \quad \text{for } t = 3, \dots, N \\ \phi_1, \phi_2 &\sim N(\mathbf{0}, \tau^2 \mathbf{Q}(\mathbf{W}, \rho)^{-1}). \end{aligned}$$

Here  $\mathbf{0}$  denotes an  $M \times 1$  vector of zeros, and the model specification is completed by the weakly informative priors  $\tau^2 \sim \text{inverse-gamma}(1, 0.01)$ ,  $\rho, \alpha \sim \text{Uniform}(0, 1)$  and  $f(\alpha_1, \alpha_2) \propto 1$ . Temporal autocorrelation is modelled by the mean  $\alpha \phi_{t-1}$  or  $\alpha_1 \phi_{t-1} + \alpha_2 \phi_{t-2}$ , where  $\alpha = 0$  and  $\alpha_1 = \alpha_2 = 0$  respectively correspond to temporal independence for the two specifications. Note, temporal stationarity of the autoregressive processes is not enforced in the models, i.e. by restricting the values of  $(\alpha, \alpha_1, \alpha_2)$ , because in a general case the data may exhibit a prominent temporal trend.

Spatial autocorrelation is induced by the precision matrix  $\mathbf{Q}(\mathbf{W}, \rho) = \rho[\text{diag}(\mathbf{W}\mathbf{1}) - \mathbf{W}] + (1 - \rho)\mathbf{I}$ , where  $(\mathbf{1}, \mathbf{I})$  are respectively an  $M \times 1$  vector of ones and an  $M \times M$  identity matrix. This precision matrix corresponds to the conditional autoregressive (CAR) prior for spatial autocorrelation proposed by [Leroux, Lei and Breslow \(2000\)](#), where the spatial autocorrelation structure assumed by the model is defined by a binary  $M \times M$  neighbourhood matrix  $\mathbf{W}$  that quantifies which pairs of pixels are close together. Here  $w_{ij,rs} = 1$  if pixels  $(\mathcal{H}_{ij}, \mathcal{H}_{rs})$  share a common border (but not a vertex), and  $w_{ij,rs} = 0$  otherwise. Based on this neighbourhood structure the level of spatial dependence is controlled globally for the entire study region by  $\rho$ , where  $\rho = 1$  corresponds to strong autocorrelation (the intrinsic CAR model proposed by [Besag, York and Mollié, 1991](#)), while  $\rho = 0$  corresponds to spatial independence. Assuming that  $\rho > 0$  then if  $w_{ij,rs} = 1$  then  $[\phi_t(\mathcal{H}_{ij}), \phi_t(\mathcal{H}_{rs})]$  are modelled as partially spatially autocorrelated, while if  $w_{ij,rs} = 0$  then  $[\phi_t(\mathcal{H}_{ij}), \phi_t(\mathcal{H}_{rs})]$  are modelled as conditionally independent given the remaining random effects. This partial autocorrelation structure assumes that all pairs of spatially neighbouring pixels have correlated random effects, despite the goal of the paper being to identify the locations of risk boundaries representing steep risk gradients.

An alternative would be to jointly estimate the spatial dependence structure as part of the model, for example by treating by the set of  $\{w_{ij,rs}\}$  corresponding to neighbouring pixels as random quantities to be estimated (e.g. [Rushworth, Lee and Sarran, 2017](#)). However, this adds an additional 5,620 parameters to be estimated in addition to both the grid-level disease risks and the grid-level disease counts, which need to be estimated using only the areal unit level data. Inference from such a model would likely be extremely sensitive to the prior distribution assumed for  $\{w_{ij,rs}\}$ , with little prior to posterior learning. There are more parsimonious alternatives for modelling  $\{w_{ij,rs}\}$  that depend on covariates (e.g. [Lee and Mitchell, 2012](#), and [Berchuck, Mwanza and Warren, 2019](#)), but as described above grid-level covariate information is not available. Additionally, the use of any of these methods could bias the boundary identification in stage 2 below, because assuming two neighbouring pixels have conditionally independent random effects (by estimating  $w_{ij,rs} = 0$ ) would make it more likely that a steep risk gradient is estimated than if the random effects were assumed to be partially autocorrelated (by estimating  $w_{ij,rs} = 1$ ).

**3.1.3. Inference.** Inference for this model uses the data augmented Markov chain Monte Carlo (MCMC) simulation algorithm proposed by [Taylor, Andrade-Pacheco and Sturrock \(2018\)](#), which jointly updates the grid-level model parameters  $\boldsymbol{\Omega} = \{\beta, \phi, \tau^2, \rho, \alpha\}$  (or  $\boldsymbol{\Omega} = \{\beta, \phi, \tau^2, \rho, \alpha_1, \alpha_2\}$  for the AR(2) model) and the grid-level disease counts  $\mathbf{Y}^{\mathcal{H}}(\mathcal{S}) = \{Y_t(\mathcal{H}_{ij})\}$  conditional on the disease and covariate data  $\mathcal{D}$  and the area of intersection data  $\mathcal{A} = \{a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})\}$ . The disease and covariate data  $\mathcal{D}$  comprise the IZ-level observed disease counts  $\{Y_t(\mathcal{A}_k^{(t)})\}$ , the spatially re-aligned grid-level expected counts  $\{e_t(\mathcal{H}_{ij})\}$ , and any grid-level covariate information  $\{\mathbf{x}_t(\mathcal{H}_{ij})\}$ . The algorithm iterates the following two steps.

1. Sample from  $f(\boldsymbol{\Omega} | \mathbf{Y}^{\mathcal{H}}(\mathcal{S}), \mathcal{D}, \mathcal{A})$ , the conditional distribution of the grid-level model parameters  $\boldsymbol{\Omega}$  given the current values of the grid-level disease counts  $\mathbf{Y}^{\mathcal{H}}(\mathcal{S})$  and the observed data  $\{\mathcal{D}, \mathcal{A}\}$ .
2. Sample from  $f(\mathbf{Y}^{\mathcal{H}}(\mathcal{S}) | \boldsymbol{\Omega}, \mathcal{D}, \mathcal{A})$ , the conditional distribution of the grid-level disease counts  $\mathbf{Y}^{\mathcal{H}}(\mathcal{S})$  given the grid-level model parameters  $\boldsymbol{\Omega}$  and the observed data  $\{\mathcal{D}, \mathcal{A}\}$ .

Further details of these two steps are given in Section 1.1 of the supplementary material, while software to fit this model is also available as supplementary material. The software is written in R, but is made computationally efficient by updating the random effects in C++ and utilising the triplet form representation of  $\mathbf{W}$ .



3.2. *Stage 2 - Boundary detection.* The second stage identifies boundaries in the gridded risk surfaces estimated in stage 1, using edge detection algorithms from computer vision. The aim of these algorithms is to identify locations of rapid change in the colours of an image, and a review of these algorithms by [Muntarina, Shorif and Uddin \(2022\)](#) classify them into 5 categories. Here we use gradient-based methods, which identifies a pixel as a boundary if it has a large risk gradient (rate of change). We choose this category because spatial smoothing has been used to estimate the gridded risk surfaces in stage 1 of our approach, meaning that boundaries will manifest themselves where a small number of neighbouring pixels exhibit a steep risk gradient, rather than there being a step change in risk between 2 neighbouring pixels. We consider two different types of gradient-based methods below, which respectively identify boundaries that are either thick (multiple pixels wide) or thin (one pixel wide). Thick boundaries allow one to identify sub-regions of rapid change in disease risk, while thin boundaries allow one to pinpoint the location of maximal change within this sub-region. While elucidating the algorithms we denote the disease risk surface generically by  $\theta = \{\dot{\theta}(\mathcal{H}_{ij})\}$ , and discuss the different options for this surface in Section 3.2.3 below.

3.2.1. *Thick boundaries.* Two commonly used methods for identifying thick boundaries were proposed by [Sobel, I and Feldman, G \(1968\)](#) and [Prewitt \(1970\)](#), which begin by estimating the gradient at each pixel in both east-west (EW) and north-south (NS) directions. They do this by applying the following masks to each pixel in turn:

- **Sobel masks** -  $\nabla_{EW}^S = \frac{1}{4} \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}$  and  $\nabla_{NS}^S = \frac{1}{4} \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}$ .
- **Prewitt masks** -  $\nabla_{EW}^P = \frac{1}{3} \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}$  and  $\nabla_{NS}^P = \frac{1}{3} \begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}$ .

Each directional mask is placed on the gridded risk surface centered on the pixel of interest  $\mathcal{H}_{ij}$ , and the set of 9 pixel level risks in the enclosing  $3 \times 3$  square are multiplied by the corresponding elements in the mask before summing to obtain the estimated derivative. For example, the east-west derivative of  $\dot{\theta}(\mathcal{H}_{ij})$  based on the Sobel mask is estimated by:

$$\nabla_{EW}^S[\dot{\theta}(\mathcal{H}_{ij})] = \frac{1}{4} \left\{ [\dot{\theta}(\mathcal{H}_{i-1,j+1}) + 2\dot{\theta}(\mathcal{H}_{i,j+1}) + \dot{\theta}(\mathcal{H}_{i+1,j+1})] - [\dot{\theta}(\mathcal{H}_{i-1,j-1}) + 2\dot{\theta}(\mathcal{H}_{i,j-1}) + \dot{\theta}(\mathcal{H}_{i+1,j-1})] \right\}.$$

This estimator smoothes the directional derivative  $\dot{\theta}(\mathcal{H}_{i,j+1}) - \dot{\theta}(\mathcal{H}_{i,j-1})$  over the rows above and below row  $i$ , which reduces the impact of random noise in the surface on the estimated derivative. The Sobel mask gives the row / column in question double the weight of its neighbours, while the Prewitt mask gives it the same weight as its neighbours. The magnitude of the risk gradient at  $\mathcal{H}_{ij}$  based on the Sobel masks is computed by

$$|\nabla^S[\dot{\theta}(\mathcal{H}_{ij})]| = \sqrt{\nabla_{EW}^S[\dot{\theta}(\mathcal{H}_{ij})]^2 + \nabla_{NS}^S[\dot{\theta}(\mathcal{H}_{ij})]^2},$$

where an analogous formula holds for the Prewitt gradient. Pixel  $\mathcal{H}_{i,j}$  is then classified as a boundary in the risk surface if  $|\nabla^S[\dot{\theta}(\mathcal{H}_{ij})]| > \xi$  for a chosen threshold  $\xi$ , and is classified as a non-boundary otherwise. The threshold  $\xi$  is chosen to be an upper quantile of the distribution of  $\{|\nabla^S[\dot{\theta}(\mathcal{H}_{ij})]|\}$  over all pixels, and boundaries identified using higher thresholds represent larger boundaries with steeper gradients. Note, that derivatives and hence boundaries cannot be computed for pixels on the very edge of the study region, because the above formulae require neighbouring pixels to be present on all sides.

3.2.2. *Thin boundaries.* The two most commonly used gradient-based algorithms that detect thin boundaries are the Laplacian of Gaussian (LoG) detector proposed by Marr and Hildreth (1980) and the Canny detector proposed by Canny (1986), and a brief description is given below. A full exposition is given in Section 1.2 of the supplementary material.

### Laplacian of Gaussian detector

The LoG detector has three stages, with the first applying a Gaussian smoothing mask with variance  $\sigma^2$  to the estimated risk surface to reduce its noise. We consider  $\sigma^2 = 0.2, 2$  in the simulation studies that follow, which show that this choice has very little impact on the results. The second stage computes the Laplacian second derivative of this smoothed surface, and identifies locations where this changes sign from negative to positive or vice-versa between spatially adjacent pixels, which are known as *zero crossings*. Pixels with zero second derivatives are either a local maximum (positive gradient) or a local minimum (negative gradient) in the gradient surface, which thus could correspond to boundary locations. However, zero crossings can occur in the Laplacian surface between two spatially adjacent pixels that exhibit a small rather than a large change in risk, meaning that a zero crossing is not a sufficient condition for identifying a risk boundary. Therefore the third stage identifies pixel  $\mathcal{H}_{ij}$  as a risk boundary if: (i) it is a zero crossing; and (ii) its gradient is above a threshold  $\xi$ , which is an upper quantile of the gradient distribution taken across all zero crossing locations.

### Canny detector

The LoG detector is known to be badly affected by noise in the image, which can result in: (i) isolated pixels being identified as boundaries that do not form part of a connected boundary segment; and (ii) holes in a connected boundary segment, which are pixels that should have been identified as part of a boundary but were not. Canny (1986) proposed a four-stage method for overcoming these limitations, which is still the most commonly used gradient-based method. The first stage applies a Gaussian smoothing mask with variance  $\sigma^2$  to the estimated risk surface to reduce its noise, which is analogous to the first stage in the LoG detector. The second stage estimates the overall magnitude and direction (with respect to the horizontal  $x$ -axis) of the gradient at each pixel. Stage three applies a non-maximal suppression procedure to the estimated gradient surface, which restricts the boundaries identified to be a single pixel wide. Finally, stage four applies a two-level threshold procedure with higher and lower thresholds  $(\xi, \xi/2)$  to the gradient surface obtained from stage 3, which removes isolated boundary points generated by noise and ensures that connected boundary segments are identified.

3.2.3. *Implementation.* Stage 1 produces  $r = 1, \dots, R$  grid-level posterior risk samples  $\{\theta_t^{(r)}(\mathcal{H}_{ij})\}$  for pixel  $\mathcal{H}_{ij}$  and time period  $t$ , and we propose the following approach for identifying *time-varying* boundaries based on a chosen threshold  $\xi$ .

- **Point estimation** - Classify pixel  $\mathcal{H}_{ij}$  as a boundary or not at time  $t$  by applying a boundary detection algorithm to the posterior median risk surface, that is set  $\hat{\theta}(\mathcal{H}_{ij}) = \text{Median}_r\{\theta_t^{(r)}(\mathcal{H}_{ij})\}$  in the previous subsection. We use the median rather than the mean risk because the posterior risk distributions exhibit right skew.
- **Posterior uncertainty** - Classify pixel  $\mathcal{H}_{ij}$  as a boundary or not at time  $t$  separately for each of the  $R$  posterior risk samples. That is set  $\hat{\theta}(\mathcal{H}_{ij}) = \theta_t^{(r)}(\mathcal{H}_{ij})$  separately for each sample  $r$  in the previous subsection. Then compute the posterior probability that  $\mathcal{H}_{ij}$  is a risk boundary as the proportion of the  $R$  samples where  $\mathcal{H}_{ij}$  is classified as a boundary.

Here  $\xi$  quantifies the *magnitude* of the boundary identified, with pixels classified as a boundary under higher thresholds of  $\xi$  corresponding to steeper risk gradients. Then given

$\xi$  the posterior probability quantifies the likelihood that a pixel is a risk boundary of that magnitude. Obviously there is no single *true* value of  $\xi$  for transforming a set of continuous gradients into a binary set of *{boundary, non-boundary}* values, as one can have small boundaries that exhibit moderate gradients as well as large boundaries that exhibit steep gradients. In practice we follow the pragmatic solution of [Lu and Carlin \(2005\)](#), which identifies boundaries at a small number of different thresholds  $\xi$ .

An alternative approach to point estimation would be to classify pixel  $\mathcal{H}_{ij}$  as a risk boundary if its posterior probability of being a boundary was above a cutoff, say  $\nu$ . However, this is inferior to using the posterior median risk for boundary detection both conceptually and in practice. Conceptually, the construction of the Canny detector includes steps to ensure the boundaries identified are a single pixel wide and form connected boundary segments, and these desirable properties are lost when combining the results of applying the Canny detector to each posterior risk sample (see the uncertainty maps in Section 5 that illustrate this). In practice, a small simulation study in Section 2.3 of the supplementary material shows that regardless of the cutoff  $\nu$  chosen for classifying boundaries, this approach performs worse in terms of correct boundary identification than applying the algorithms to the posterior median risk surface. This inferior performance is likely to be because the posterior median provides a better estimate of the disease risk surface than the individual samples from its posterior distribution do, leading to improved boundary identification.

The above implementation identifies *time-varying* boundaries, which allows an assessment of how much risk boundaries change over time. However, the identification of *time-invariant* boundaries that persist over the entire study period is also of interest, and point estimation in this case is achieved by applying the boundary detection algorithms to the posterior median of the time averaged risks, i.e. setting  $\hat{\theta}(\mathcal{H}_{ij}) = \text{Median}_r \{ \frac{1}{N} \sum_{t=1}^N \theta_t^{(r)}(\mathcal{H}_{ij}) \}$  in the previous subsection. In this case posterior uncertainty is quantified by applying the boundary detection algorithms to  $\hat{\theta}(\mathcal{H}_{ij}) = \frac{1}{N} \sum_{t=1}^N \theta_t^{(r)}(\mathcal{H}_{ij})$  separately for each posterior sample  $r$ .

**4. Simulation study - time-varying boundaries.** This study quantifies the ability of our approach to accurately estimate the locations of grid-level risk boundaries that change over time. A second study focusing on time-invariant boundaries is presented in Section 2.2 of the supplementary material, while a study comparing posterior median classification with the alternative posterior uncertainty classification is presented in Section 2.3 of the supplementary material. The template for all these studies matches the Glasgow application, with disease counts available at  $K^{(1)} = 133$  IZs (2001 definition) for the first five years and for  $K^{(2)} = 136$  IZs (2011 definition) for the second five year period. Our methodology is then used to estimate disease risk and identify risk boundaries on the  $250m^2$  grid presented in Figure 1.

4.1. *Data generation.* Grid-level observed disease counts  $\{Y_t(\mathcal{H}_{ij})\}$  are generated from (2), where  $\{e_t(\mathcal{H}_{ij})\}$  are based on the real data but are varied in size in our simulation design. The true risk surfaces are generated as  $\{\theta_t(\mathcal{H}_{ij}) = \exp[\phi_t(\mathcal{H}_{ij}) + \delta_t(\mathcal{H}_{ij})]\}$ , where  $\{\phi_t(\mathcal{H}_{ij})\}$  is generated from an AR(1) process similar to (3). Here, the spatial autocorrelation matrix  $\mathbf{Q}(\mathbf{W}, \rho)^{-1}$  is replaced with a geostatistical style Gaussian autocorrelation matrix, whose autocorrelation between pixels  $1km$  apart is 0.9. Strong temporal autocorrelation is induced by setting  $\alpha = 0.9$  while  $\tau^2 = 0.03$  (both from (3)), and both were chosen to match the motivating study results.

Boundaries are induced into the grid-level risk surfaces via a piece-wise-constant spatial surface  $\{\delta_t(\mathcal{H}_{ij})\}$  that varies over time. Boundary template 1 is used for  $t = 1, 2, 3$ ; boundary template 2 is used for  $t = 4, 5, 6, 7$ ; and boundary template 3 is used for  $t = 8, 9, 10$ , and maps of these templates together with example simulated risk surfaces are displayed in Section 2.1

of the supplementary material. To ensure the boundaries are realistic, boundary template 1 is generated by clustering the time-averaged estimated risks from the real data in Section 5 using the k-means algorithm with 2 clusters. Pixels in the high risk cluster have  $\delta_t(\mathcal{H}_{ij}) = \lambda$  while those in the low risk cluster have  $\delta_t(\mathcal{H}_{ij}) = -\lambda$ . Thus risk boundaries occur where two spatially adjacent pixels have  $\delta_t(\mathcal{H}_{ij})$  values of  $\lambda$  and  $-\lambda$ . However, to create a single pixel at the centre of each risk boundary we set  $\delta(\mathcal{H}_{ij}) = 0$  for one of these two pixels, which means that each risk boundary spans three adjacent pixels where  $\delta(\mathcal{H}_{ij})$  moves from  $(-\lambda, 0, \lambda)$ . The value of  $\lambda$  controls the size of the risk boundary as larger values of  $\lambda$  lead to steeper risk gradients, and its value is varied in our simulation design. Boundary templates 2 and 3 represent a slow evolution in the boundary locations compared to template 1, with around 7% of the pixels having changed risk category between boundary templates 1 and 2, while the change between templates 2 and 3 corresponds to 10% of the pixels.

Based on  $\{\theta_t(\mathcal{H}_{ij})\}$  and the area intersection data  $\{a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})\}$ , we generate observed and expected disease counts for each intersection area  $\{\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}\}$ . The majority have  $a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = 0$ , in which case  $Y_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = 0$ . Conversely, if  $a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) > 0$  then we compute  $e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) = \gamma \frac{a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})}{\sum_{r,s} a(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{rs})} e_t(\mathcal{A}_k^{(t)})$  using area weighted interpolation, where  $e_t(\mathcal{A}_k^{(t)})$  comes from the real data and  $\gamma$  controls disease prevalence and is varied in our simulation design. Finally, we generate  $Y_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij}) \sim \text{Poisson}[e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})\theta(\mathcal{H}_{ij})]$ , which has a similar mean model to (2). Then IZ-level data are computed via  $Y_t(\mathcal{A}_k^{(t)}) = \sum_{i,j} Y_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})$  and  $e_t(\mathcal{A}_k^{(t)}) = \sum_{i,j} e_t(\mathcal{A}_k^{(t)} \cap \mathcal{H}_{ij})$  respectively, which are the only quantities used when estimating risk boundaries.

*4.2. simulation design.* We generate 100 data sets under each of 6 scenarios, which comprise all pairwise combinations of the following 2 factors.

- **Boundary size** - The magnitude of the risk boundaries (i.e. the risk gradient at a boundary) is either small ( $\lambda = 0.125$ ) or large ( $\lambda = 0.25$ ), and example realisations of both cases are presented in Section 2.1 of the supplementary material.
- **Disease prevalence** - The disease is either rare ( $\gamma = 0.5$ , around 33 disease cases on average), moderately prevalent ( $\gamma = 1$ , around 66 disease cases on average) or common ( $\gamma = 2$ , around 132 disease cases on average), with the middle case matching the size of the counts observed in the motivating study.

*4.3. Results.* The AR(1) random effects model was used in this study rather than the AR(2) model due to its simpler form, and the real data results show the estimated risk surfaces from both models are almost identical. Inference for each data set is based on 6,000 MCMC samples from a single Markov chain, which was run for 350,000 samples that were thinned by 50 to reduce their autocorrelation and the first 50,000 were removed as the burnin period. The data augmentation step was implemented every  $G = 30$ th iteration of the algorithm, and pilot runs showed this was sufficient for MCMC convergence. The generation of these samples took around 2 hours on an iMac with 32 GB of memory and a 3.8 GHz i7 processor, and the mean effective sample size across the set of grid-level risk parameters was around 700.

The four edge detection algorithms Canny, LoG, Prewitt and Sobel were then applied to the posterior median grid-level risk surfaces estimated from each simulated data set separately for each year, and both the Canny and LoG detectors were applied with two levels of smoothing defined by  $\sigma^2 = 0.2, 2$ . Each method classifies every pixel as either a boundary or a non-boundary based on a threshold gradient  $\xi$ , and this binary classification is computed for 27

Boundary Size	Disease Prevalence	Boundary detection method					
		Canny (0.2)	Canny (2)	LoG (0.2)	LoG (2)	Prewitt	Sobel
Small	Low	0.85	0.83	0.78	0.76	0.84	0.85
Small	Medium	0.87	0.86	0.80	0.78	0.87	0.87
Small	High	0.89	0.88	0.81	0.80	0.90	0.90
Large	Low	0.90	0.90	0.83	0.82	0.92	0.92
Large	Medium	0.92	0.92	0.85	0.84	0.94	0.94
Large	High	0.93	0.93	0.86	0.85	0.95	0.95

TABLE 1

Area under the curve results for each boundary detection method and scenario.

different  $\xi$  values ranging between the 1st and 99th percentile of the estimated gradients across the study region. For each value of  $\xi$  we compute: (i) True positive rate - percentage of the true boundaries that are correctly identified; and (ii) False positive rate - percentage of the true non-boundaries that are falsely identified as boundaries. As previously discussed each true risk boundary is 3 pixels wide, so we define an estimated boundary as correctly identifying a true boundary if it coincides with any of the 3 pixels that make up the boundary. From these metrics we compute receiver operating characteristic (ROC) curves presented in Figure 2, and the area under the curve (AUC) statistics presented in Table 1 to quantify the accuracy of boundary identification, and these results are averaged over all time periods and simulated data sets.

The figure and table show that all methods perform well in all scenarios, with AUC values ranging between 0.76 and 0.95. Better performance is achieved when the disease is more prevalent (compare the top and bottom rows of Figure 2), which is because there are more disease cases and hence better estimates of disease risk. For example, the differences in the AUC between the low and high prevalence cases is between 0.02 and 0.06. Secondly, better performance is achieved for larger boundaries that are easier to identify (compare the left and right columns in Figure 2), with the AUC increasing by between 0.04 and 0.08. The amount of smoothing applied by the Canny and Laplacian of Gaussian detectors has minimal effects on the results, as the ROC curves for  $\sigma^2 = 0.2$  and  $\sigma^2 = 2$  are almost identical in all scenarios.

In comparing the different methods the Prewitt and Sobel detectors (thick boundaries) deliver almost identical results, while Canny always outperforms the Laplacian of Gaussian method (thin boundaries). The differences in AUC between the Canny and Prewitt / Sobel detectors are minimal, and as discussed above, they produce different boundary types (*thin* and *thick*). These different types explain the differences in the true and false positive rates obtained by these detectors, with Prewitt / Sobel identifying more pixels as boundaries and thus having higher true and false positive rates compared to the Canny detector.

## 5. Results from the Glasgow respiratory hospitalisation study.

5.1. *Stage 1 - Grid-level risk estimation.* The AR(1) and AR(2) grid-level disease risk models were fitted to the IZ-level respiratory hospitalisation data with no covariates, where the grid-level expected numbers of respiratory hospitalisations were estimated by (1). Inference is based on 9,000 MCMC samples generated from 3 independent Markov chains, which were each burnt in for 100,000 samples before being run for a further 300,000 samples that were then thinned by 100 to reduce their autocorrelation. Convergence of the Markov chains for  $(\beta, \rho, \alpha, \alpha_1, \alpha_2, \tau^2)$  and a sample of the grid-level risks  $\{\theta_t(\mathcal{H}_i)\}$  was assessed by trace-plots and the Gelman-Rubin diagnostic (Gelman et al., 2013), and both indicate that convergence is likely with the potential scale reduction factors being less than 1.1 in all cases. The

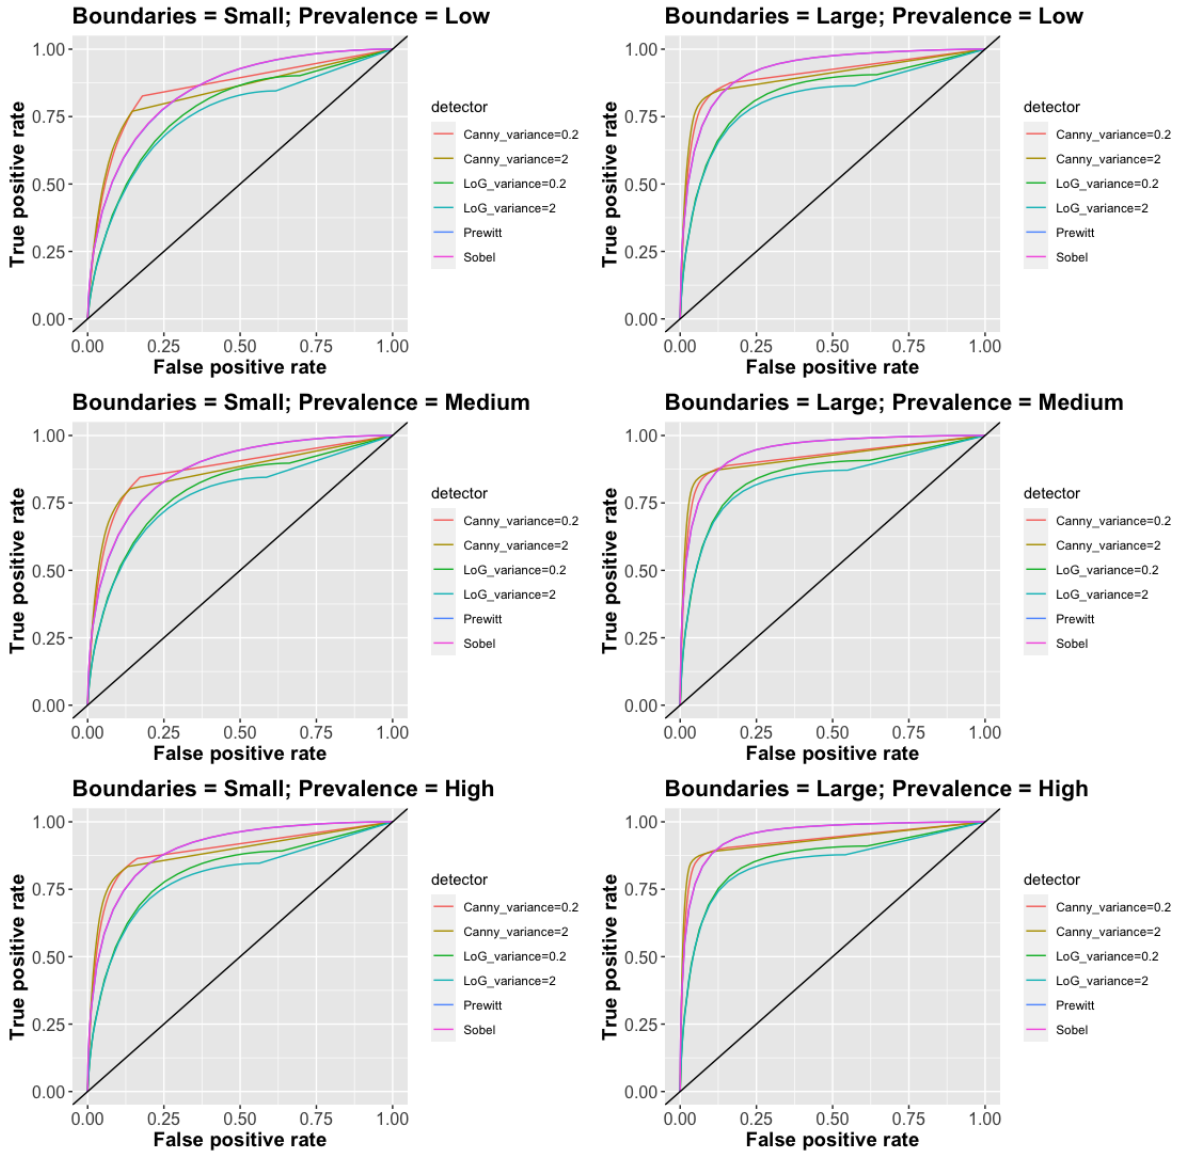


FIG 2. Receiver operating characteristic curves for each boundary detection method and simulation scenario. The black lines are the  $y = x$  line and represent random guessing.

generation of these samples took around 6 hours on an iMac with 32 GB of memory and a 3.8 GHz i7 processor, and the mean effective sample size across the grid-level risks was 1,854 for the AR(1) model and 2,927 for the AR(2) model.

The disease risk surfaces exhibit strong spatio-temporal dependence, which is illustrated by the posterior median and 95% credible intervals for the dependence parameters  $\rho$  (spatial) and  $(\alpha, \alpha_1, \alpha_2)$  (temporal). For the AR(1) model  $\rho$ : 0.91 (0.87, 0.94) and  $\alpha$ : 0.94 (0.91, 0.97), while for the AR(2) model  $\rho$ : 0.91 (0.86, 0.94),  $\alpha_1$ : 0.58 (0.49, 0.68) and  $\alpha_2$ : 0.39 (0.29, 0.49). The two models also produce very similar grid-level disease risk estimates, with their posterior medians having a Pearson's correlation coefficient of 0.99 and a mean absolute difference of 0.035. Furthermore, their widely applicable information criterion (WAIC, [Watanabe, 2010](#)) and effective number of independent parameters (in brackets) are very similar, being AR(1): 132,594 (15,340); and AR(2): 132,613 (15,240). Therefore we use the

results from the AR(1) model as its WAIC is slightly lower, but this choice should not affect the risk boundaries identified below.

## 5.2. Stage 2 - Boundary identification.

5.2.1. *Time-invariant boundaries.* We first identify the locations of the most prominent risk boundaries that persisted over the 10-year study period (motivating question 1.), by applying the boundary detection algorithms to the time-averaged risk surfaces as described in Section 3.2.3. We use the Sobel and Canny detectors here, because the simulation study suggests they can accurately identify thick and thin boundaries respectively. In contrast the LoG detector performed poorly compared to the Canny detector, while the Prewitt detector gave almost identical results to the Sobel detector, and hence both are omitted. Results for the Canny detector that presents the locations of maximal change are displayed in Figure 3, while results for the Sobel detector that shows sub-regions with relatively large gradients are shown in Section 3 of the supplementary material. The Canny detector is applied with a smoothing variance of  $\sigma^2 = 2$ , and the boundaries using  $\sigma^2 = 0.2$  are almost identical and are not shown for brevity. Maps (A) and (C) present the posterior median of the time-averaged risk surfaces, with boundaries identified by dots at threshold levels of  $\xi = 0.7$  and  $\xi = 0.9$  respectively using the posterior median approach. The posterior boundary probabilities associated with these thresholds are presented in panels (B) and (D). In both cases the grey lines correspond to the IZ-level borders.

The boundaries identified in panels (A) and (C) visually correspond to regions of rapid change in disease risk, being present where low risk areas (shaded yellow) neighbour high risk areas (shaded red). They also mainly comprise connected boundary segments of multiple spatially adjacent pixels, with almost no isolated boundary points. More boundaries are identified when  $\xi = 0.7$  (panel A) compared to when  $\xi = 0.9$  (panel C), which is because the latter is a subset of the former by the construction of the algorithm. The posterior probabilities that each pixel is a risk boundary at both thresholds  $\xi$  are presented in panels (B) and (D), with the darker shaded pixels having a higher probability of being a boundary. The highest probability pixels largely correspond to the boundaries identified in panels (A) and (C), but these elevated probability areas are not a single pixel wide, suggesting that there is some uncertainty as to the exact location of the maximal risk gradient between spatially neighbouring pixels.

One of the prominent boundary segments separates the low risk city centre and west-end (between Park district and Kelvindale) parts of the city from the higher risk suburbs of Possilpark to the north and Govan to the south. The latter boundary largely follows the river Clyde, and physical barriers such as this have previously been shown to lead to social barriers that separate communities and prevent them mixing (Mitchell and Lee, 2014). One of the other prominent features is the long curved boundary segment in the east-end of the city separating the low risk Springhill area with its more high-risk northerly neighbour Blairtummock, and a sizeable part of this boundary also coincides with a physical barrier, namely the main railway line that connects Glasgow with the capital city of Edinburgh.

Finally for comparison purposes with the existing literature, panel (E) presents estimated risks and boundaries (dots) on an IZ-level footprint. The risks are estimated from an IZ-level equivalent of the Poisson log-linear spatio-temporal model with AR(1) temporal autocorrelation presented in Section 3, while the boundaries are identified using the recent graph-theoretic approach proposed by Lee, Meeks and Pettersson (2021). Note, this model and the resulting risk map are only based on data from 2013-2017, because the changes in the IZ borders in 2012 preclude it from being applied to the entire study period as the IZs change over time. The results show that while the boundaries agree with those identified above in

a sizeable number of cases, there are numerous examples where the boundaries identified appear to be false positives and separate two IZs that have very similar risks. This may be because this approach is not designed to form extended boundary segments, resulting in a number of short disconnected boundaries that do not appear to reflect large changes in disease risk. Finally, this approach only provides a spatially discrete approximation to the risk surface, because the risk is assumed to be constant within each IZ.

*5.2.2. Time-varying risk boundaries.* To observe the extent to which the boundaries have changed over time (motivating question 2.), we apply the Canny detector with  $\sigma^2 = 2$  to the posterior median risk surface for each year separately. The resulting boundaries are displayed in Figure 4 for: (A) 2008, (B) 2011, (C) 2014, and (D) 2017, where the boundaries correspond to a threshold  $\xi$  of the 70th percentile of the set of possible boundaries across the study region. The spatial pattern in disease risk is similar for all time periods, with Pearson's correlation coefficients between any two spatial risk surfaces ranging between 0.78 (2008 and 2017) and 0.90 (2014 and 2017). Additionally, a substantial number of the boundaries are consistently observed for all four years, such as the boundary separating the high-risk east-end areas of Blairtummock and Easterhouse from their southerly neighbours, and the demarcation between the low-risk west-end and city centre parts of the city from the higher risk suburbs of Possilpark to the north and Govan to the south. In fact, between 41% (2008 and 2017) and 59% (2014 and 2017) of the boundaries are exactly the same between any of the two years presented in the figure, while a sizeable proportion of the boundaries that are discordant have only changed slightly in space.

The main change over time is that the high-risk areas have increased in risk but shrunk in size, which is evident from the smaller but darker shaded red areas present in the 2017 map compared to the 2008 map. An example is Port Dundas in the middle of the city, which is completely surrounded by a boundary in 2017 where as in 2008 it is part of a much larger but less extreme high-risk area. These changes have increased the magnitude of the inequality (variation) in the risk of respiratory disease hospitalisation over the study period, with for example the spatial standard deviations being 0.33 and 0.44 in 2008 and 2017 respectively. The reduction of health inequalities is a major target for the Scottish Government ([NHS Health Scotland, 2016](#)), and these results suggest that this is a complex problem that cannot be easily solved.

*5.2.3. Appropriateness of IZ-level inference.* The majority of existing studies such as [Lee, Meeks and Pettersson \(2021\)](#) make inference on disease risk and boundary locations at the areal unit level at which they have data, which in this study would imply that: (i) disease risk is constant within an IZ; and (ii) the risk boundaries identified coincide with the geographical borders between neighbouring IZs. We examine the appropriateness of these assumptions (motivating question 3.) by first computing the magnitude of the within IZ variation in disease risk separately for each year (assumption i). Specifically, for each IZ  $\mathcal{A}_k^{(t)}$  in year  $t$  we consider all pixels  $\mathcal{H}_{ij}$  that overlap with  $\mathcal{A}_k^{(t)}$ , and compute a weighted standard deviation of their posterior median risks. Here the weights are proportional to the areas of overlap, so that if a pixel has a very small overlap with an IZ then it does not heavily influence its estimated standard deviation. In 2008 the average (over all IZs) within IZ standard deviation in risk was 0.12, which was 37% of the magnitude of the standard deviation across all pixels in the study region. By 2017 these values had increased to a within IZ standard deviation of 0.20, which was 46% of the region-wide standard deviation. Thus there is some



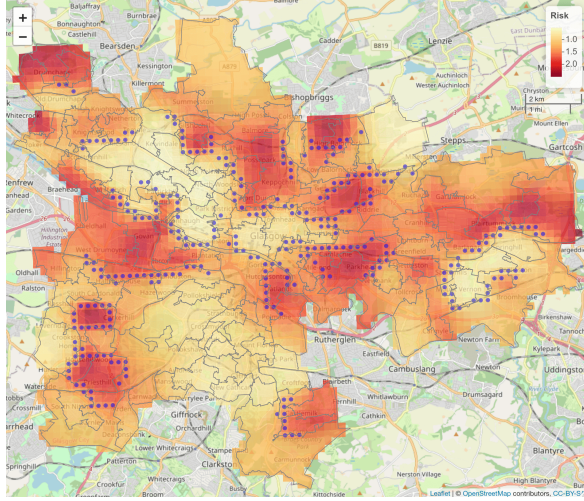
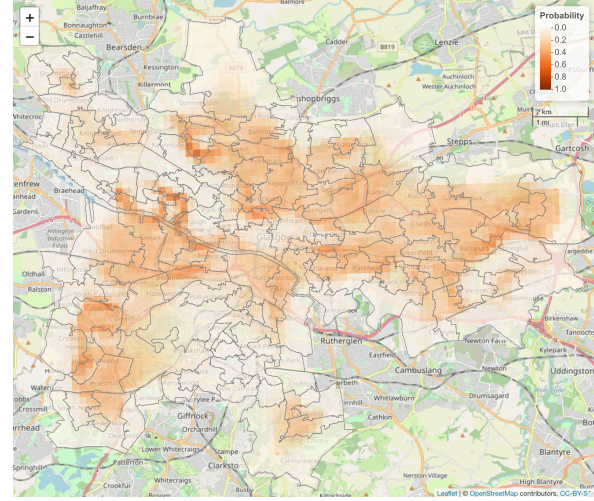
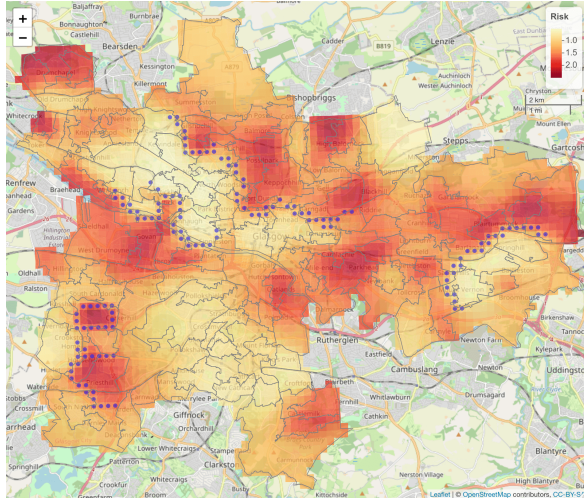
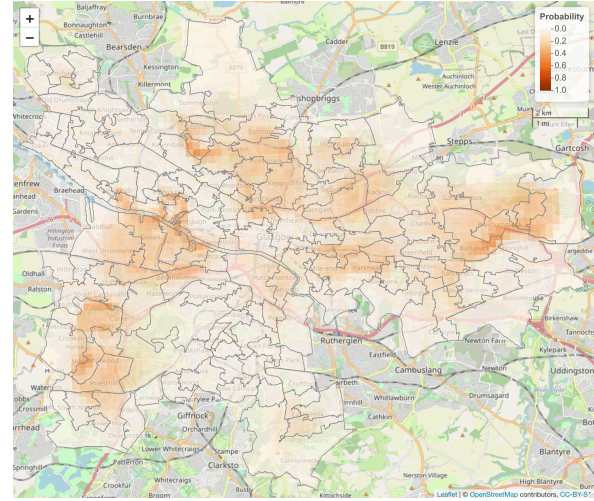
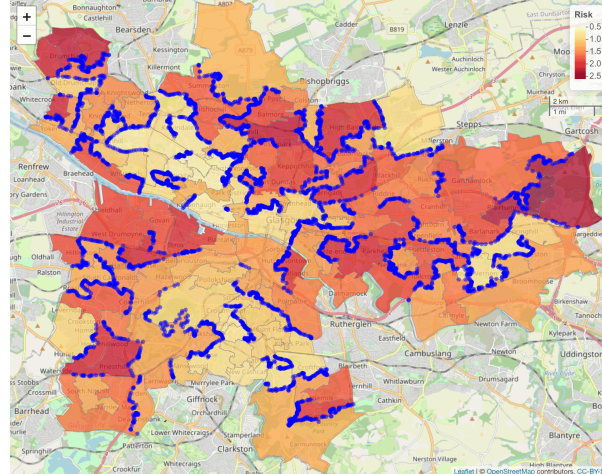
**(A) - Canny risks -  $\xi = 0.7$** **(B) - Canny posterior probabilities -  $\xi = 0.7$** **(C) - Canny risks -  $\xi = 0.9$** **(D) - Canny posterior probabilities -  $\xi = 0.9$** **(E) - Areal unit based boundaries**

FIG 3. The left maps (A) and (C) display the posterior median time-averaged risk surfaces with the locations of boundaries (dots) identified using the Canny detector ( $\sigma^2 = 2$ ) with  $\xi = 0.7$  (top) and  $\xi = 0.9$  (middle). The right maps (B) and (D) display the posterior probabilities that each pixel is a boundary based on thresholds of  $\xi = 0.7$  (top) and  $\xi = 0.9$  (middle). Panel (E) presents boundaries identified (blue dots) at the areal unit level using the graph-theoretic approach proposed by [Lee, Meeks and Pettersson \(2021\)](#).

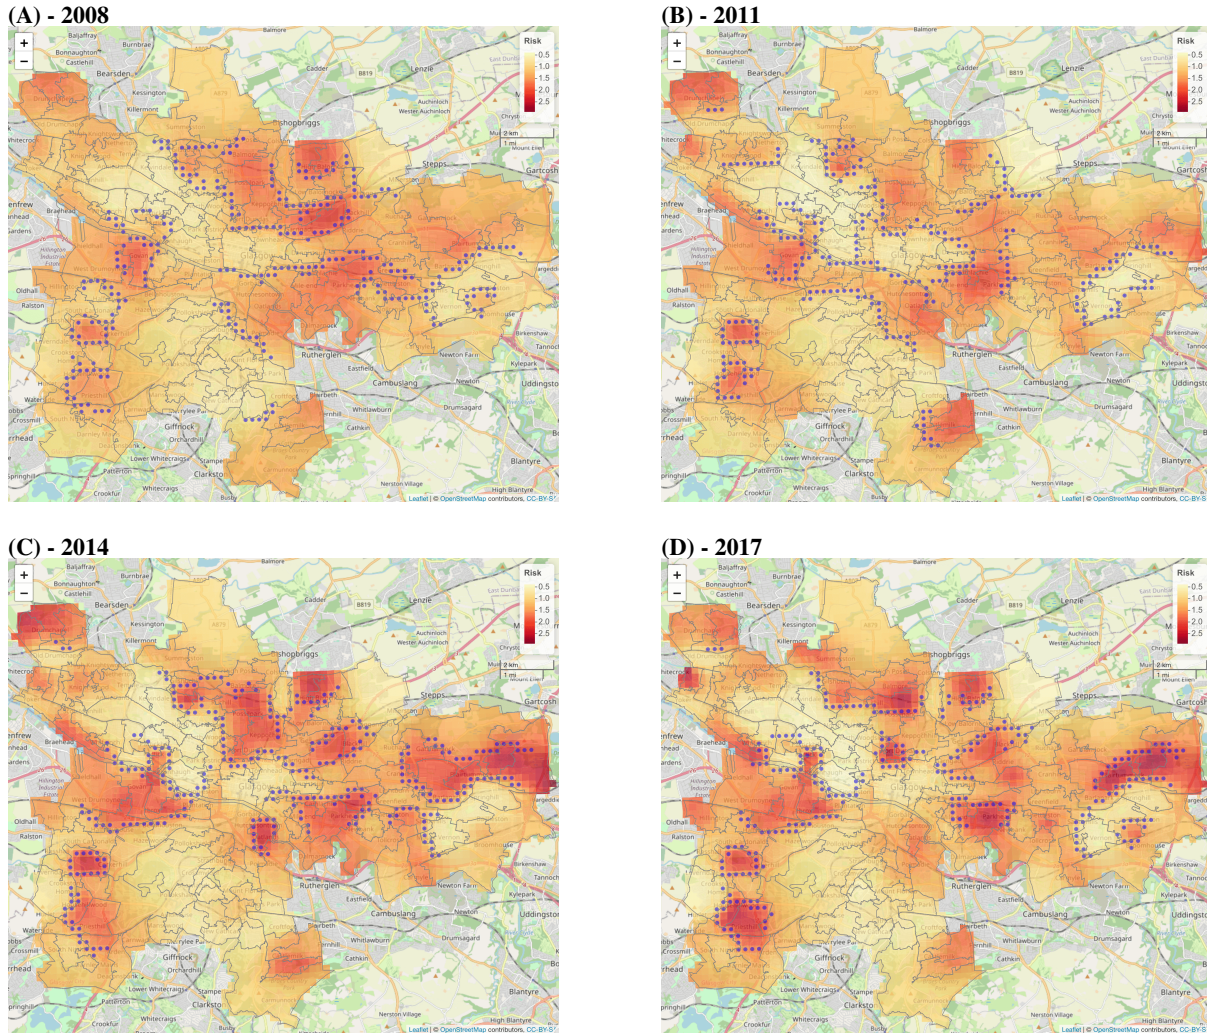


FIG 4. Maps of the estimated risk surfaces and the boundaries identified by the Canny detector for: (A) 2008, (B) 2011, (C) 2014, and (D) 2017. The blue dots represent boundary locations based on a threshold  $\xi$  corresponding to the 70th percentile of the set of possible boundaries across the study region.

within IZ variation in disease risk, which while less than the variation across the entire city, nonetheless suggests that the assumption of a constant risk within each IZ is inappropriate.

To examine the appropriateness of assumption (ii), we computed the percentage of the pixels identified as boundaries in Figure 4 that coincide with a border between two neighbouring IZs. Specifically, we computed the percentages of the boundary pixels that had an IZ border running through them, which were 78% (2008), 79% (2011), 78% (2014) and 77% (2017) respectively for the four years presented. This shows that the locations exhibiting the steepest risk gradients generally coincide with IZ borders, suggesting that on the whole the most prominent community boundaries are demarked by an IZ border, and that IZs typically represent either a single community or multiple communities whose disease risks are not too dissimilar.

**6. Discussion.** This paper has proposed the first approach for estimating the locations of disease risk boundaries at approximately any location in the study region using only spa-

tially aggregated areal unit level data whose borders change over time, using a novel fusion of Bayesian spatial re-alignment modelling and edge detection algorithms. This approach improves on the existing literature such as [Lu and Carlin \(2005\)](#), [Lee and Mitchell \(2012\)](#), [Ma, Carlin and Banerjee \(2010\)](#) and [Lee, Meeks and Pettersson \(2021\)](#), which all make the restrictive assumptions that: (i) the disease risk surface is piece-wise constant; and (ii) the boundaries identified must coincide with the border separating two geographically adjacent areal units. Our simulation studies provide convincing evidence of accurate boundary detection via ROC curves and AUC statistics across a range of scenarios, with the Sobel / Prewitt algorithms being appropriate if one wants to identify regions exhibiting steep risk gradients (thick boundaries), while the Canny algorithm is appropriate if one wishes to identify the exact location of the maximal risk gradient (thin boundaries). Our two-stage approach can be utilised by other researchers using the software provided in the supplementary material, which contains R functions to implement the methodology as well as simulated data sets to make the simulation study reproducible.

The Glasgow study shows clear evidence of sizeable spatial variation in the risk of respiratory hospitalisation across the city, with risks ranging between less than half the size of the citywide average risk and over two and a half times greater than it. This spatial variation naturally results in risk boundaries being present, and both the Canny and Sobel algorithms identify a number of connected boundary segments that mainly persist over the entire 10 year study period. These boundaries, such as the demarcation of the city centre and west end areas, are similar to boundaries identified in 2005 by [Lee and Mitchell \(2013\)](#), even though in this study the boundaries were restricted to an areal unit footprint. The high-risk areas are mainly those that are socio-economically deprived, and reducing the link between poverty and elevated disease risk is a long-standing priority for Public Health Scotland (<http://www.healthscotland.scot/health-inequalities/measuring-health-inequalities>). Additionally, the size of within IZ risk variation is around 40% of the size of the between IZ risk variation, while just under 80% of the risk boundaries coincide with IZ borders. These figures suggest that in this study while making inference at the areal unit level should capture the main spatial patterns in disease risk, there will be some finer grain risk variation that will be missed by this commonly used inferential paradigm.

A potential limitation of our approach is the implicit assumption in the spatial realignment model that the random effects in all neighbouring pixels are partially autocorrelated, when the goal of the analysis is to find the locations of risk boundaries. The reasons for this choice were outlined in Section 3, and the simulation study in Section 4 showed that true risk boundaries are well identified by our methodology, suggesting that it is not detrimental to model performance. However, a natural avenue for future work would be to relax this assumption, perhaps by allowing the spatial dependence parameter  $\rho$  to vary across the study region. A further area for methodological development is that here we have focused exclusively on gradient-based boundary detection algorithms, because the aim was to identify locations exhibiting steep risk gradients. However, there are many non-gradient-based edge detection algorithms that one could apply to the grid-level risk surface, and future work will therefore compare the approaches used here to popular alternative methods such as SUSAN ([Smith and Brady, 1997](#)) and deep learning based approaches ([Zhiding et al., 2017](#)).

In future the application will be scaled up from the city of Glasgow to the whole of Scotland, which would require approximately 12 million  $250m^2$  pixels across all 10 years of the study compared to the 29,230 used for Glasgow. This ambition highlights a limitation of the current approach in that it is relatively computationally expensive, and is unlikely therefore to scale well to such a massive spatial study region. Therefore a focus in this future work will be the development of a simpler and hence more computationally efficient spatial realignment method for stage one of our approach. Finally, the existence of boundaries in areal

unit level data are not just of interest in an epidemiological context, because social scientists are interested in the identification of boundaries in other socio-demographic factors and their resulting effects on issues such as crime rates and educational attainment (e.g. [Dean et al., 2019](#)). Therefore future work will investigate the existence of boundaries in features such as social class, religious beliefs and ethnicity and their effects on society, which will require the spatial realignment models used in stage 1 to be adjusted to deal with continuous and binary areal unit level measurements.

## SUPPLEMENTARY MATERIAL

### Supplementary document

Additional description of the methods and the simulation and Glasgow studies.

### Supplementary code and data

Software in the form of R functions to fit the model together with simulated data (real data are the property of Public Health Scotland) to allow others to utilise the methods developed here.

## REFERENCES

- BANERJEE, S. and GELFAND, A. E. (2006). Bayesian Wombling. *Journal of the American Statistical Association* **101** 1487-1501.
- BERCHUCK, S., MWANZA, J. and WARREN, J. (2019). Diagnosing glaucoma progression with visual field data using a spatiotemporal boundary detection method. *Journal of the American Statistical Association* **114** 1063-1074.
- BERNARDINELLI, L., CLAYTON, D., PASCUTTO, C., MONTOMOLI, C., GHISLANDI, M. and SONGINI, M. (1995). Bayesian analysis of space-time variation in disease risk. *Statistics in Medicine* **14** 2433-2443.
- BESAG, J., YORK, J. and MOLLIE, A. (1991). Bayesian image restoration with two applications in spatial statistics. *Annals of the Institute of Statistics and Mathematics* **43** 1-59.
- BRADLEY, J., WIKLE, C. and HOLAN, S. (2016). Bayesian spatial change of support for count-valued survey data with application to the American community survey. *Journal of the American Statistical Association* **514** 472-487.
- BRADLEY, J., WIKLE, C. and HOLAN, S. (2017). Regionalization of multiscale spatial processes by using a criterion for spatial aggregation error. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **79** 815-832.
- CANNY, J. (1986). A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **8** 679-698.
- DEAN, N., DONG, G., PIEKUT, A. and PRYCE, G. (2019). Frontiers in residential segregation: understanding neighbourhood boundaries and their impacts. *Tijdschrift voor Economische en Sociale Geografie* **110** 271-288.
- FISHER, T., ZHANG, J., COLEGATE, S. and VANNI, M. (2022). Detecting and modeling changes in a time series of proportions. *The Annals of Applied Statistics* **16** 477 – 494.
- FLOWERDEW, R. and GREEN, M. (1989). *Statistical methods for inference between incompatible zonal systems* In *Accuracy of spatial databases* 239-247. Taylor and Francis.
- FLOWERDEW, R. and GREEN, M. (1993). *Developments in areal interpolation methods and GIS* In *Geographic information systems, spatial modelling and policy evaluation* 73–84. Springer Berlin Heidelberg.
- GELMAN, A., CARLIN, J., STERN, H., DUNSON, D., VEHTARI, A. and RUBIN, D. (2013). *Bayesian data analysis*, 3rd ed. Chapman and Hall / CRC.
- GRAMATICA, M., CONGDON, P. and LIVERANI, S. (2021). Bayesian modelling for spatially misaligned health areal data: A multiple membership approach. *Journal of the Royal Statistical Society: Series C* **70** 645-666.
- KNORR-HELD, L. (2000). Bayesian modelling of inseparable space-time variation in disease risk. *Statistics in Medicine* **19** 2555-2567.
- LEE, D., MEEKS, K. and PETERSSON, W. (2021). Improved inference for areal unit count data using graph-based optimisation. *Statistics and Computing* **31** 51.
- LEE, D. and MITCHELL, R. (2012). Boundary detection in disease mapping studies. *Biostatistics* **13** 415-426.
- LEE, D. and MITCHELL, R. (2013). Locally adaptive spatial smoothing using conditional auto-regressive models. *Journal of the Royal Statistical Society Series C* **62** 593-608.

- LEROUX, B., LEI, X. and BRESLOW, N. (2000). *Estimation of disease rates in small areas: A new mixed model for spatial dependence* In *Statistical models in epidemiology, the environment and clinical trials*, Halloran, M and Berry, D (eds) 135-178. Springer-Verlag, New York.
- LU, H. and CARLIN, B. (2005). Bayesian areal Wombling for geographical boundary analysis. *Geographical Analysis* **37** 265-285.
- MA, H., CARLIN, B. and BANERJEE, S. (2010). Hierarchical and joint site-edge methods for medicare hospice service region boundary analysis. *Biometrics* **66** 355-364.
- MACKENBACH, J., VALVERDE, J., ARTNIK, B. and ET AL. (2018). Trends in health inequalities in 27 European countries. *Proceedings of the National Academy of Sciences* **115** 6440–6445.
- MARR, D. and HILDRETH, E. (1980). Theory of edge detection. *Proceedings of the Royal Society of London. Series B, Biological Sciences* **207** 187-217.
- MITCHELL, R. and LEE, D. (2014). Is there really a ‘wrong side of the tracks’ in urban areas and does it matter for spatial analysis? *Annals of the Association of American Geographers* **104** 432-443.
- MUGGLIN, A. and CARLIN, B. (1998). Hierarchical modeling in geographic information systems: Population interpolation over incompatible zones. *Journal of Agricultural, Biological, and Environmental Statistics* **3** 111–130.
- MUNTARINA, K., SHORIF, S. and UDDIN, M. (2022). Notes on edge detection approaches. *Evolving Systems* **13** 169-182.
- PREWITT, J. (1970). *Object enhancement and extraction* In *Picture processing and psychopictorics*, Lipkin, B and Rosenfeld, A (eds) 75-149. New York: Academic Press.
- QU, K., BRADLEY, J. and NIU, X. (2021). Boundary detection using a Bayesian hierarchical model for multiscale spatial data. *Technometrics* **63** 64-76.
- RUSHWORTH, A., LEE, D. and MITCHELL, R. (2014). A Spatio-temporal model for estimating the long-term effects of air pollution on respiratory hospital admissions in Greater London. *Spatial and Spatio-temporal Epidemiology* **10** 29-38.
- RUSHWORTH, A., LEE, D. and SARRAN, C. (2017). An adaptive spatiotemporal smoothing model for estimating trends and step changes in disease risk. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **66** 141–157.
- NHS HEALTH SCOTLAND (2016). Health inequalities - what are they and how do we reduce them? <http://www.healthscotland.scot/media/1086/health-inequalities-what-are-they-how-do-we-reduce-them-mar16.pdf>.
- SMITH, S. and BRADY, M. (1997). SUSAN — a new approach to low level image processing. *International Journal of Computer Vision* **23** 45-78.
- SOBEL, I AND FELDMAN, G (1968). A 3x3 isotropic gradient operator for image processing. Presentation at Stanford A.I. project 1968.
- SYRING, N. and LI, M. (2017). BayesBD: An R package for Bayesian inference on image boundaries. *The R Journal* **9** 149–162.
- TAYLOR, B., ANDRADE-PACHECO, R. and STURROCK, H. (2018). Continuous inference for aggregated point process data. *Journal of the Royal Statistical Society: Series A* **181** 1125-1150.
- WAKEFIELD, J. and KIM, A. (2013). A Bayesian model for cluster detection. *Biostatistics* **14** 752-765.
- WALLER, L., CARLIN, B., XIA, H. and GELFAND, E. (1997). Hierarchical spatio-temporal mapping of disease rates. *Journal of the American Statistical Association* **92** 607-617.
- WATANABE, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research* **11** 3571-3594.
- ZHIDING, Y., CHEN, F., MING-YU, L. and SRIKUMAR, R. (2017). CASNet: deep category-aware semantic edge detection. *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)* 5964–5973.