

OPEN

A real-time spatio-temporal syndromic surveillance system with application to small companion animals

Alison C. Hale^{1,7*}, Fernando Sánchez-Vizcaíno^{2,7*}, Barry Rowlingson¹, Alan D. Radford^{3,4}, Emanuele Giorgi¹, Sarah J. O'Brien^{5,6} & Peter J. Diggle¹

Lack of disease surveillance in small companion animals worldwide has contributed to a deficit in our ability to detect and respond to outbreaks. In this paper we describe the first real-time syndromic surveillance system that conducts integrated spatio-temporal analysis of data from a national network of veterinary premises for the early detection of disease outbreaks in small animals. We illustrate the system's performance using data relating to gastrointestinal disease in dogs and cats. The data consist of approximately one million electronic health records for dogs and cats, collected from 458 UK veterinary premises between March 2014 and 2016. For this illustration, the system predicts the relative reporting rate of gastrointestinal disease amongst all presentations, and updates its predictions as new data accrue. The system was able to detect simulated outbreaks of varying spatial geometry, extent and severity. The system is flexible: it generates outcomes that are easily interpretable; the user can set their own outbreak detection thresholds. The system provides the foundation for prompt detection and control of health threats in companion animals.

Surveillance systems have been developed globally for animal and/or public health purposes, facilitating the prevention and control of disease or infection nationally and regionally. During the past decade, the emergence of new diseases¹ and the increasing threat of bio-terrorism have motivated the development of syndromic surveillance systems in public health focused on the early detection of health threats that require effective public health action^{2,3}. Syndromic surveillance uses health-related data that precedes diagnosis. Although data of this kind are less specific than data from confirmed diagnoses they are typically more timely, which is an important consideration for real-time or near-real-time analysis and interpretation⁴. In veterinary medicine the development of systems for early health-event detection has followed a similar path to that previously taken in public health⁵. A recent inventory of current and planned European veterinary syndromic surveillance systems showed wide interest in European countries for syndromic surveillance, but also highlighted the novelty of this field⁶.

Small companion animal populations largely lack co-ordinated national and international disease surveillance. This has produced a deficit in our understanding of the dynamics and burden of the full range of endemic/emerging diseases in companion animals and leaves these populations susceptible to the emergence of health threats. Lack of disease surveillance also has implications for human health, as approximately 75 percent of new and emerging diseases are zoonotic⁷. However, as health records become digitised in veterinary practices they become more available for research⁸, providing an opportunity to improve companion animal syndromic surveillance in clinical settings and the possibility of linking this with human syndromic surveillance. Recently, electronic syndromic surveillance data on companion animals has become available in real-time on a national scale

¹Centre for Health Informatics, Computing, and Statistics (CHICAS), Lancaster Medical School, Lancaster University, Lancaster, LA1 4YW, UK. ²Bristol Veterinary School, Langford Campus, University of Bristol, Bristol, BS40 5DU, UK. ³NIHR Health Protection Research Unit in Emerging and Zoonotic Infections, University of Liverpool, Liverpool, UK. ⁴Department of Infection Biology, Institute of Infection and Global Health, Leahurst Campus, University of Liverpool, Neston, CH64 7TE, UK. ⁵Department of Public Health and Policy, Institute of Psychology Health and Society, The Farr Institute@HeRC, University of Liverpool, Liverpool, L69 3GL, UK. ⁶NIHR Health Protection Research Unit in Gastrointestinal Infections, University of Liverpool, Liverpool, UK. ⁷These authors contributed equally: Alison C. Hale and Fernando Sánchez-Vizcaíno. *email: a.c.hale@lancaster.ac.uk; f.s-vizcaino@bristol.ac.uk

in the UK through surveillance schemes such as the Small Animal Veterinary Surveillance Network (SAVSNET)⁹. SAVSNET harnesses the growing volume of patient electronic health records (EHRs) available from small animal practices and complementary data from diagnostic laboratories to improve animal and human health through rapid and actionable research and surveillance.

Here we propose a real-time syndromic surveillance system that uses a spatio-temporal model in conjunction with Bayesian inference for the early detection of health-event outbreaks. Specifically, we use a Markov Chain Monte Carlo (MCMC) algorithm to generate samples from the Bayesian predictive distribution of the underlying spatio-temporal surface. These samples are then used to compute predictive probabilities at given thresholds; a high predictive probability at a particular location and time gives an early warning of a possible disease outbreak. The system provides end-users (i.e. practising veterinary surgeons) decision-support tools for immediate analysis and easy interpretation of their data. As an example, we apply our model to small companion animal EHRs collected over two years by SAVSNET from a large network of UK veterinary premises. We illustrate the feasibility of our proposed surveillance system using gastrointestinal (GI) disease in dogs and cats as an example.

Gastrointestinal (GI) disease is one of the four syndromes for which SAVSNET currently gathers information for every consultation it receives. GI disease affects animal welfare, can be expensive to manage and may be transmissible to other pets¹⁰ or, more rarely, to people¹¹. Current approaches to preventing and controlling GI disease in companion animals have focussed on individuals or small groups of animals. This seems to have had little impact on GI disease, which remains one of the commonest reasons for presenting for veterinary care in the UK^{9,10,12–15}, although precise data to confirm this has been lacking. A more coordinated population-scale approach to GI disease surveillance in companion animals is needed.

This paper focuses on the early detection of a GI disease *outbreak*, which we define as an unexplained, spatially and temporally localised increase in the fraction of GI consultations amongst all consultations. We illustrate the performance of our proposed surveillance system on simulated GI disease outbreaks of varying spatial extent and severity. This is, to our knowledge, the first surveillance system that conducts integrated spatio-temporal analysis of data from a national network of veterinary practices so as to enable real-time detection of spatially and temporally localised changes in reporting patterns across the network.

The paper is structured as follows. First, we give details of the SAVSNET and socioeconomic data used in this paper. We then give the rationale for our methodological approach, describe the spatio-temporal stochastic model that is the foundation of our surveillance system, and report the results of fitting our model to our SAVSNET-acquired data. We then simulate spatio-temporal GI outbreaks by perturbing the actual SAVSNET data in various ways to demonstrate the ability of the surveillance system to achieve timely outbreak-detection. Finally, we discuss the similarities and differences between our proposed system and other approaches in the literature, and also extensions for joint human and veterinary surveillance.

Data Sources

Savsnet. *Data collection.* Data were collected electronically in near-real-time from volunteer veterinary premises or sites using a compatible version of the practice management system (PMS) namely RoboVet (VetSolutions, Edinburgh) and Teleos Systems Ltd (Birmingham). This study used data for dogs and cats collected over the period between 1st March 2014 and 29th February 2016. In our analysis we included data from an increasing number of premises as they enrolled in the RoboVet and Teleos systems. By 29th February 2016 we had data from 216 practices (amounting to a total of 458 distinct premises) located in England, Wales and Scotland. The data were extracted from consultations where a booked appointment was made to see a veterinary surgeon or nurse, including out-of-hours consultations. Through the SAVSNET system a compulsory, single-question questionnaire is appended at the end of each consultation allowing the attending veterinary surgeon or nurse to categorise the main reason for the animal's presentation into syndromes (currently GI disease, respiratory disease, pruritus and renal disease) or other routine veterinary interventions (i.e., trauma, neoplasia, 'other sick', vaccination, 'other healthy' or post-operative check-up). Specifically, the definition provided to participating veterinary surgeons to categorise the animal presentation as GI disease is that the main reason for the animal's presentation are signs including but not limited to diarrhoea, vomiting, weight loss and poor appetite. A full description of the SAVSNET data collection protocol has been described by Sánchez-Vizcaíno *et al.*⁹. The data for this study were gathered on a consultation-by-consultation basis, and include the date the animal was seen, unique identifiers for practice, premise and animal, the animal description (including species, breed, sex and date of birth), the syndromic level classification and the full postcode of each veterinary premise and pet owner.

Data were only gathered if the owner had not opted out of study participation. The collection and use of these data were approved by the University of Liverpool's Research Ethics Committee (RETH00964); as such all collection and use of these data were performed in accordance with the relevant guidelines and regulations.

Data management. Text-based data for species and breed were cleaned to deal with misspellings or the use of non-standard terms by mapping to standard terms. A full description of this cleaning procedure has been described elsewhere¹⁶. Many breeds were present in the data set, some represented by only a few individuals, limiting the scope for analysis by breed. Thus, for the purposes of this study, only the animal's classification as purebred or crossbred was used.

To identify localised outbreaks we needed to geocode all postcodes. The text-based data for each owner's full postcode were automatically cleaned by applying mapping rules of typical misspellings (e.g. letter 'O' instead of zero). Any remaining records containing erroneous postcodes were discarded from our outbreak prediction as they could not be geocoded. Similarly, if the age of the animal was recorded outside the range 0 to 25 years then the record was excluded. SAVSNET records with missing data were removed before the analysis. If an animal attended a veterinary premise on more than one occasion during the study period we included all attendances

Species	Sex	Breed-status	Number of animal consultations by age category		
			<1 year	1 < 8 years	>=8 years
Dog	Female	Crossbred	429	1089	957
Dog	Female	Purebred	2266	6411	4969
Dog	Male	Crossbred	448	1151	916
Dog	Male	Purebred	2777	6876	4874
Cat	Female	Crossbred	488	1242	2295
Cat	Female	Purebred	123	233	403
Cat	Male	Crossbred	514	1319	1989
Cat	Male	Purebred	142	354	403

Table 1. Age profile of dogs and cats attending SAVSNET veterinary premises for a gastrointestinal disease consultation stratified by sex and breed-status. The number of dog and cat consultations shown included only animals with a mapped breed-status, sex and date of birth within the range of 0 to 25 years recorded.

without adjustment, on the grounds that multiple visits occurring within a short time period (e.g. within a few days) would likely indicate a more serious illness episode.

Data summary. Of the 1,211,326 consultations collected between 1st March 2014 and 29th February 2016, 72.3% were for dogs and 27.7% for cats. In 80.7% of all consultations a valid age, breed-status (purebred or crossbred) and owner's full postcode were recorded; this subset of data is used for model selection and the basis for simulations. Gastrointestinal disease accounted for 4.0% of all presentations, amongst which 91.5% were recorded between Monday and Friday. Amongst animals presenting for GI disease, there was not a notable gender bias; 48.5% of dog consultations and 50.6% of cat consultations with a recorded sex were female. Where the breed-status was identified, 84.9% of dog GI disease consultations and 17.2% of cat GI disease consultations were purebreds. In animals with a date of birth recorded within the range 0 to 25 years, 65.4% of dog GI disease consultations and 47.4% of cat GI disease consultations were under eight years. The age profile of dogs and cats presenting for GI disease at SAVSNET veterinary premises stratified by sex and breed-status is shown in Table 1. Data for the two species were analysed separately.

Measure of deprivation. We used the pet owner's home postcode to assign a measure of deprivation to each owner using the most recent English¹⁷, Scottish¹⁸ and Welsh¹⁹ Indices of Multiple Deprivation (IMD) produced by their respective governments. A detailed description of how each government has developed their own measure of deprivation can be found elsewhere^{20–22}. The three country-specific IMD measures are not directly comparable. We therefore included *country* as a three-level factor and rescaled the ranks of each country's set of IMD scores to the range 0 to 1. For example, if for England the maximum rank was 32,000 and a location had rank 100 then the owner IMD explanatory variable would be assigned a value of 100/32,000.

Outbreak Detection Modelling

Rationale. As noted earlier, we define an *outbreak* as an unexplained spatially and temporally localised increase in the fraction of GI consultations amongst all consultations. The term “unexplained” refers to the fact that, for reasons that are well understood, some areas or times of year will experience higher fractions of GI consultations than others because of spatial variation in the local population susceptibility or temporal variation in the region-wide susceptibility to GI. We adjust for these known effects using measured explanatory variables, as described below in the section on explanatory variable selection. We then equate “unexplained” to “stochastic” and include this in our model as a latent, spatially and temporally correlated process $S_{i,t}$, where i denotes premise and t denotes time, in days. By definition, the expected value of each $S_{i,t}$ is zero, and our goal is to determine where and when its actual value is materially greater than zero. Note that the natural pattern of GI consultations will always be subject to fluctuations in time and space that cannot be explained fully by measured variables. It follows that outbreak detection is not a statistical hypothesis-testing problem. Our approach acknowledges this by the fact that the actual value of $S_{i,t}$ will never be exactly zero. Our formal solution is therefore to calculate, for each premise i and day t , the predictive probability q (i.e. the probability conditional on all available data up to and including day t) that $S_{i,t} > l$, where l is a user-specified threshold representing an effect large enough to be of practical concern. We then declare an outbreak affecting premise i if this probability exceeds q_0 , the required positive predictive value per premise, say $q_0 = 0.95$ or 0.99 . As with any prediction problem using observational data, it is not possible simultaneously to control both the positive and negative predictive probabilities.

Prediction model. To accommodate the spatial and temporal correlations that would characterise an outbreak of GI disease, we use a spatio-temporal mixed effects regression model, and fit the model using Bayesian inference. We define our binary response variable $Y_{j,it}$ to take the value 1 if the j^{th} consultation at the i^{th} premise on day t is a GI disease presentation and 0 otherwise. Conditionally on an unobserved, spatio-temporally structured random effect $S_{i,t}$, the $Y_{j,it}$ are distributed as mutually independent Bernoulli variables with probabilities $p_{j,i,t}$ defined by

$$\Phi^{-1}(p_{j,i,t}) = d_{j,i,t}^T \theta + S_{i,t} \quad (1)$$

where $\Phi^{-1}(\cdot)$ is the quantile function of the standard Normal distribution. The vector $d_{j,i,t}$ denotes the set of explanatory variables and θ their associated regression parameters. We discuss selection of explanatory variables, $d_{j,i,t}$, below.

The spatio-temporally structured collection of random effects for all premises and days is written as

$$S = (S_{(1)}^T, \dots, S_{(\tau)}^T)^T \quad (2)$$

where $S_{(t)} = (S_{1,t}, \dots, S_{n,t})^T$ and we denote by τ and n , respectively, the total numbers of days and premises contained in the data-set. The complete vector S follows a multivariate Normal distribution with mean zero and covariance matrix that incorporates the spatio-temporal context of the data. Specifically, we assume that, conditionally on its past, $S_{(t)}$ follows a multivariate Gaussian distribution with mean vector $\varphi S_{(t-1)}$ and spatial covariance matrix Ω , which we construct as follows. Firstly, we associate with premise i a polygon consisting of all points closer to premise i than to any other premise; the resulting polygons, V_i are called Voronoi polygons. Secondly, we define the neighbours of i to be the set $N(i)$ of premises whose Voronoi polygons are contiguous with V_i . Finally, we define distance-decay weights

$$w_{ik} = \begin{cases} [1 + (u_{ik}/\delta)^2]^{-1} & \text{if } k \in N(i), \delta > 0 \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

where u_{ik} is the distance between premises i and k , and δ is a scaling parameter with units of distance. We then specify the conditional distribution of each $S_{i,t}$ given all other $S_{k,t}$ to be Normal with mean ρm_{it} where

$$m_{it} = \frac{\sum_{k \in N(i)} w_{ik} S_{k,t}}{\sum_{k \in N(i)} w_{ik}}, \text{ for all } k \neq i \quad (4)$$

and variance $\sigma^2 / \sum_{k \in N(i)} w_{ik}$. Together, these modelling assumptions imply that the so-called full conditional distributions of the $S_{i,t}$ that together determine the joint distribution of S are of the form

$$S_{i,t} | S_{k,t}, S_{k,t-1} \sim N \left(\rho m_{it} + \varphi \rho m_{it-1}, \frac{\sigma^2}{\sum_{k \in N(i)} w_{ik}} \right), \text{ for all } k \neq i \quad (5)$$

Using these full conditional distributions, we can simulate from the Bayesian predictive distribution of the random effects $S_{i,t}$ using an MCMC algorithm based on auxiliary variable techniques as described in Section 4.3 of Rue & Held²³. Our system is intended to be run in near-real-time, but the MCMC computations eventually become prohibitive as the time-span of the data, τ , grows. To counteract this, we run the MCMC algorithm on a moving nine-day window, which is long enough to capture the temporal correlation in our data; the magnitude of the within-premise autocorrelation of $S_{i,t}$ for a time lag of eight days is typically around 0.09. Over a time-window of this size, the effects of any systematic time-trend or seasonal effect on the fraction of GI consultations are negligible, which removes the need to include these as explicit terms in the model; see also section below on selection of explanatory variables.

We adopt the following set of mutually independent priors for the model parameters:

$$\theta \sim \text{MVN}(0, 10^3 I); \log \sigma^2 \sim N(-5, 9); \rho \sim \text{Uniform}(0, 1); \varphi \sim \text{Uniform}(0, 1); \delta \sim \text{Uniform}\{1, 2, \dots, 100\}$$

These were chosen to be vague, in the sense that they have little influence on the predictive inferences for the random effects $S_{i,t}$ that constitute the primary goal of the analysis. However, if inferences about the model parameters are required, samples from their Bayesian joint posterior distribution are produced automatically as a by-product of the MCMC algorithm.

Outbreak detection. Let $e_{i,t}$ denote the exceedence probability for premise i on day t , i.e. the probability that $S_{i,t} > l$ conditional on all available data up to and including day t , where l is the user-specified threshold value. To calculate the $e_{i,t}$, we generate M posterior samples $S_{i,t}^{(1)}, \dots, S_{i,t}^{(M)}$ from the joint predictive distribution of the random effects $S_{i,t}$ using an MCMC algorithm, and calculate

$$e_{i,t} = \frac{1}{M} \sum_{m=1}^M \mathbb{I} \left(S_{i,t}^{(m)} > l \right) \quad (6)$$

where $\mathbb{I}(S_{i,t}^{(m)} > l)$ takes the value 1 if $S_{i,t}^{(m)} > l$ and 0 otherwise. For this calculation to be accurate, we need the MCMC algorithm first to run for a sufficiently long time, called the burn-in period, to have reached convergence and then for a further M iterations to feed Eq. (6), where M is sufficiently large that the sampling error on the right-hand-side of (6) is negligible. We used a burn-in period of 5,000 iterations, followed by $M = 50,000$ iterations.

The spatio-temporal model was fitted using the R package ‘caramellar’²⁴.

Explanatory variable selection. Generalised Linear Models (GLMs) are unsuitable for outbreak detection modelling because the parameter estimates and standard errors assume that the observations are independent;

hence, they do not take account of spatial and/or temporal correlation. Nevertheless, we can use a standard probit regression model to establish whether there is a prima-facie case for including each explanatory variable in our outbreak prediction model, Eq. (1), using the following rule. We retained an explanatory variable if its effect was nominally significant at the conventional 5% level. This inclusion rule is conservative in the sense that in the presence of spatial or temporal correlation the standard probit regression analysis is likely to over-state the significance of individual regression effects. For both species, this led us to discard the explanatory variables pet insurance, micro-chipping and neutering status and to retain the following:

- the three-level factor ‘COUNTRY’ for the pet owner’s home address (i.e. England, Scotland or Wales);
- the two-level factor ‘WEEKDAY’ with values 0 and 1 indicating if the consultation date is a weekend day (Saturday, Sunday or public holiday) or a working weekday (Monday to Friday), respectively -we considered using day of the week as a factor on 7 levels, but this did not improve the fit significantly using a likelihood ratio (deviance difference) test;
- the two-level factor ‘GENDER’ with values 0 and 1 corresponding to ‘female’ and ‘male’, respectively;
- the two-level factor ‘PUREBRED’ with values 0 and 1 corresponding to crossbred or purebred, respectively;
- the continuous variable ‘AGE’ denoting the animal’s age, in years and $AGE^2 = AGE \times AGE$, both included because the quadratic term improves the model fit;
- the continuous variable ‘IMD’, is the rescaled deprivation measure relating to the pet owner’s home address (as described above in our section on data sources).

As noted earlier, fitting the model to moving nine-day windows of data removes any long-term trend or seasonal effects. The resulting provisional GLM is

$$\Phi^{-1}(p) = \alpha_{COUNTRY} + \beta_{COUNTRY} \times IMD + \theta_1 \times WEEKDAY + \theta_2 \times GENDER + \theta_3 \times PUREBRED + \theta_4 \times AGE + \theta_5 \times AGE^2 \quad (7)$$

where p denotes the probability that a presentation of a dog or cat (depending on the species evaluated) to a SAVSNET veterinary premise is recorded as a GI disease consultation. The first two terms on the right-hand side of Eq. (7) capture the interaction between country and IMD, so as to account for the fact that the three countries use different IMD measures, whilst $\theta_1, \theta_2, \dots, \theta_5$ are regression parameters for the remaining explanatory variables in the model. The GLM outputs for dogs and cats can be found as Supplementary Tables S1 and S2, respectively.

All computation was carried out using R version 3.4.0²⁵.

Outbreak Simulations

Our model’s ability to identify an outbreak, i.e. its sensitivity, is influenced by factors including the outbreak’s duration, spatial extent and the number of infected animals presenting at premises in the locality. In each of our simulations, we construct an outbreak by adding varying numbers of aberrant GI disease to the actual (baseline) SAVSNET-recorded cases in a specified set of premises over a specified number of consecutive days.

Simulation model. We use the actual SAVSNET total consultations for dogs during February 2016, together with their associated explanatory variables, to simulate a step increase in the proportion of GI disease cases affecting one or more premises from a given day t_0 , corresponding to 15 February 2016, by augmenting Eq. (1) with an extra term as follows

$$\Phi^{-1}(p_{j,i,t}) = d_{j,i,t}^T \theta + S_{i,t} + \gamma I_i(t \geq t_0), \quad (8)$$

where the indicator function I_i for premise i has value 1 for premise i and all days $t \geq t_0$ if premise i is affected by the outbreak, and has value 0 otherwise. By varying the value of γ we can control the probability of a GI case at an affected premise.

For each simulation, we proceed as follows:

- (1) use the actual SAVSNET consultations during February 2016 to fit the no-outbreak model using Eq. (1) and to generate simulated realisations of $S_{i,t}$;
- (2) for $t \geq t_0$, use the actual explanatory variables and the simulated $S_{i,t}$ to compute $p_{j,i,t}$ using Eq. (8) with $\gamma > 0$;
- (3) use the computed values of $p_{j,i,t}$ to simulate case and control flags (1 or 0 respectively) and use these to reassign each actual SAVSNET data consultation as either a case or control.

See Supplementary Material for detailed R-code.

Simulation scenarios. We applied our simulation model to three sets of premises, which we selected based on their numbers of neighbours, defined to be other premises within an 8km radius, with the additional constraint that none of the sets of premises were within each other’s 8km radius. The selected sets of premises, which we designated as *dense*, *medium* and *sparse*, had 6, 3 and 0 neighbours, respectively. The SAVSNET data gave no indication that these selected premises are atypical or that they experienced a genuine outbreak during February 2016. See Figs 1 or 2, in each of which the top row, labelled ‘baseline’, is the actual SAVSNET data prior to simulating an outbreak. The premises at the centres of the three sets reported similar total numbers of consultations

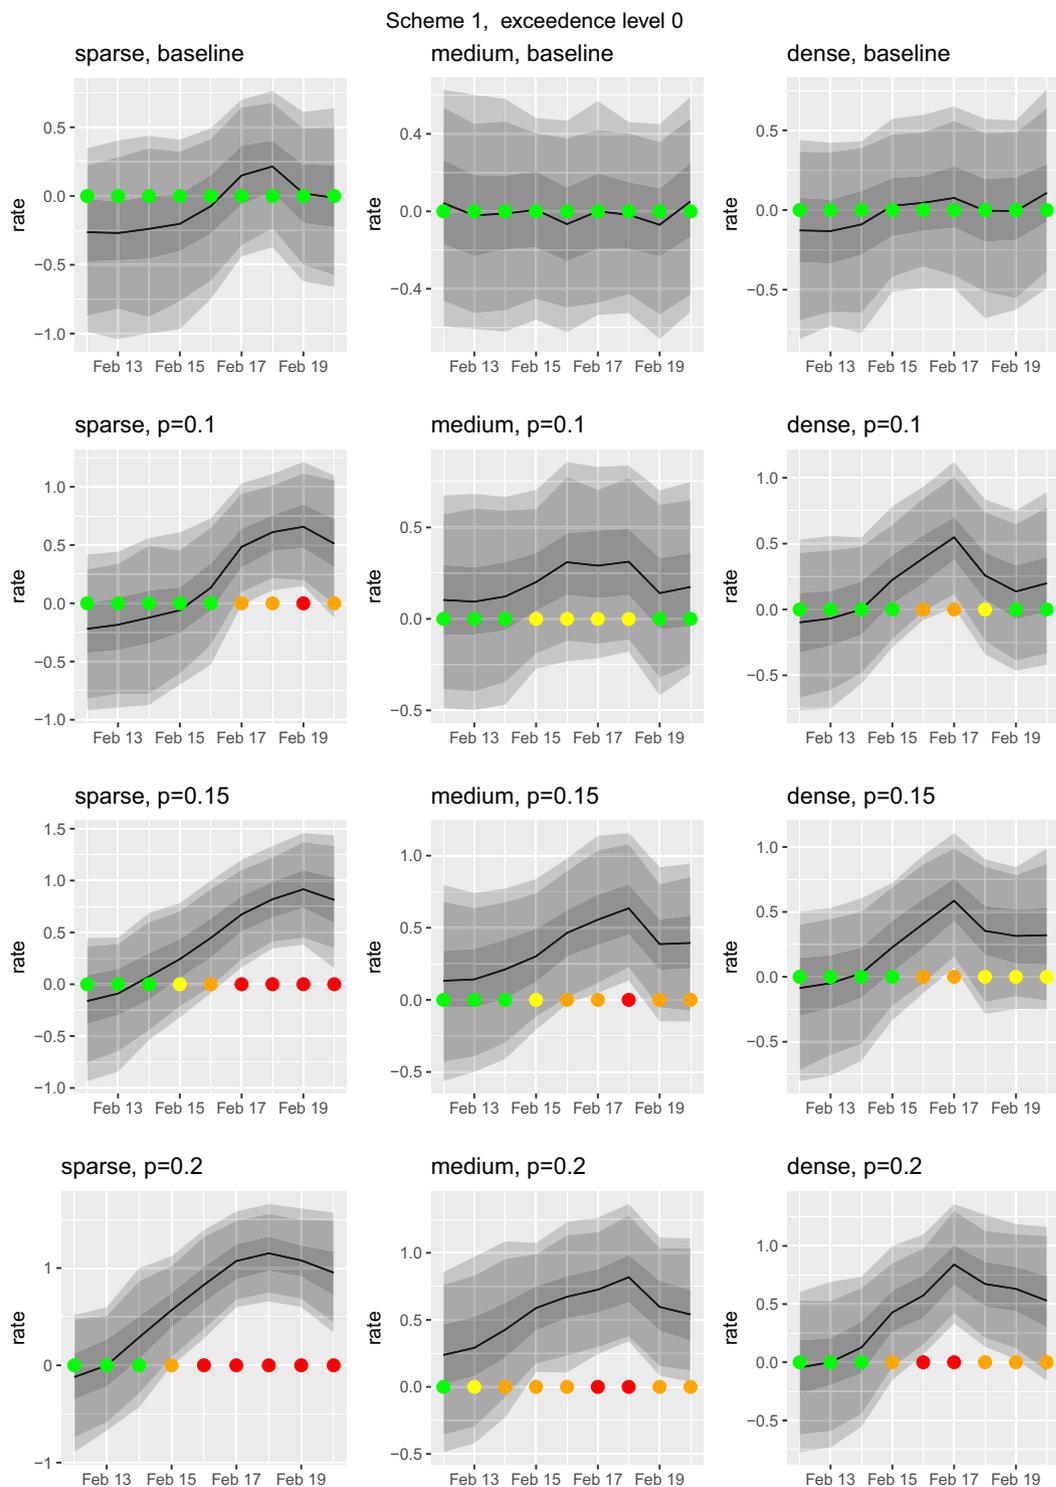


Figure 1. The results from our outbreak simulation study when using Scheme 1. In this scheme a single premise i at the centre of each region experiences an outbreak. Here we choose an exceedence level of $l=0$ (see supplementary material for other levels). This figure shows the results of 9 simulations plus the baseline level. The top row of timeseries plots is the 'baseline', that is the actual SAVSNET data without any simulated outbreak i.e. $\gamma=0$. The subsequent rows from top to bottom depict increasing severities of simulated outbreak labelled according the probability of a case at premise i e.g. $p=0$ and so on. The columns, from left to right, relate to the density of the region; 'sparse', 'medium' and 'dense' respectively. For each simulation we plot the timeseries of the predicted distribution of $S_{i,t}$ for premise i . In each time timeseries the solid black line is the predicted value of $S_{i,t}$, shaded areas are pointwise 50%, 90% and 95% predictive intervals. As an aid to rapid interpretation, we use a traffic-light system: if the predictive probability, q_t , is above 0.99 (defined as 'very high') the light shows red, if above 0.9 ('high') orange, if above 0.8 (medium) yellow, otherwise ('low') green (no outbreak). The outbreak commences on 15th February. The more intense the outbreak is the more the traffic light system tends towards red.

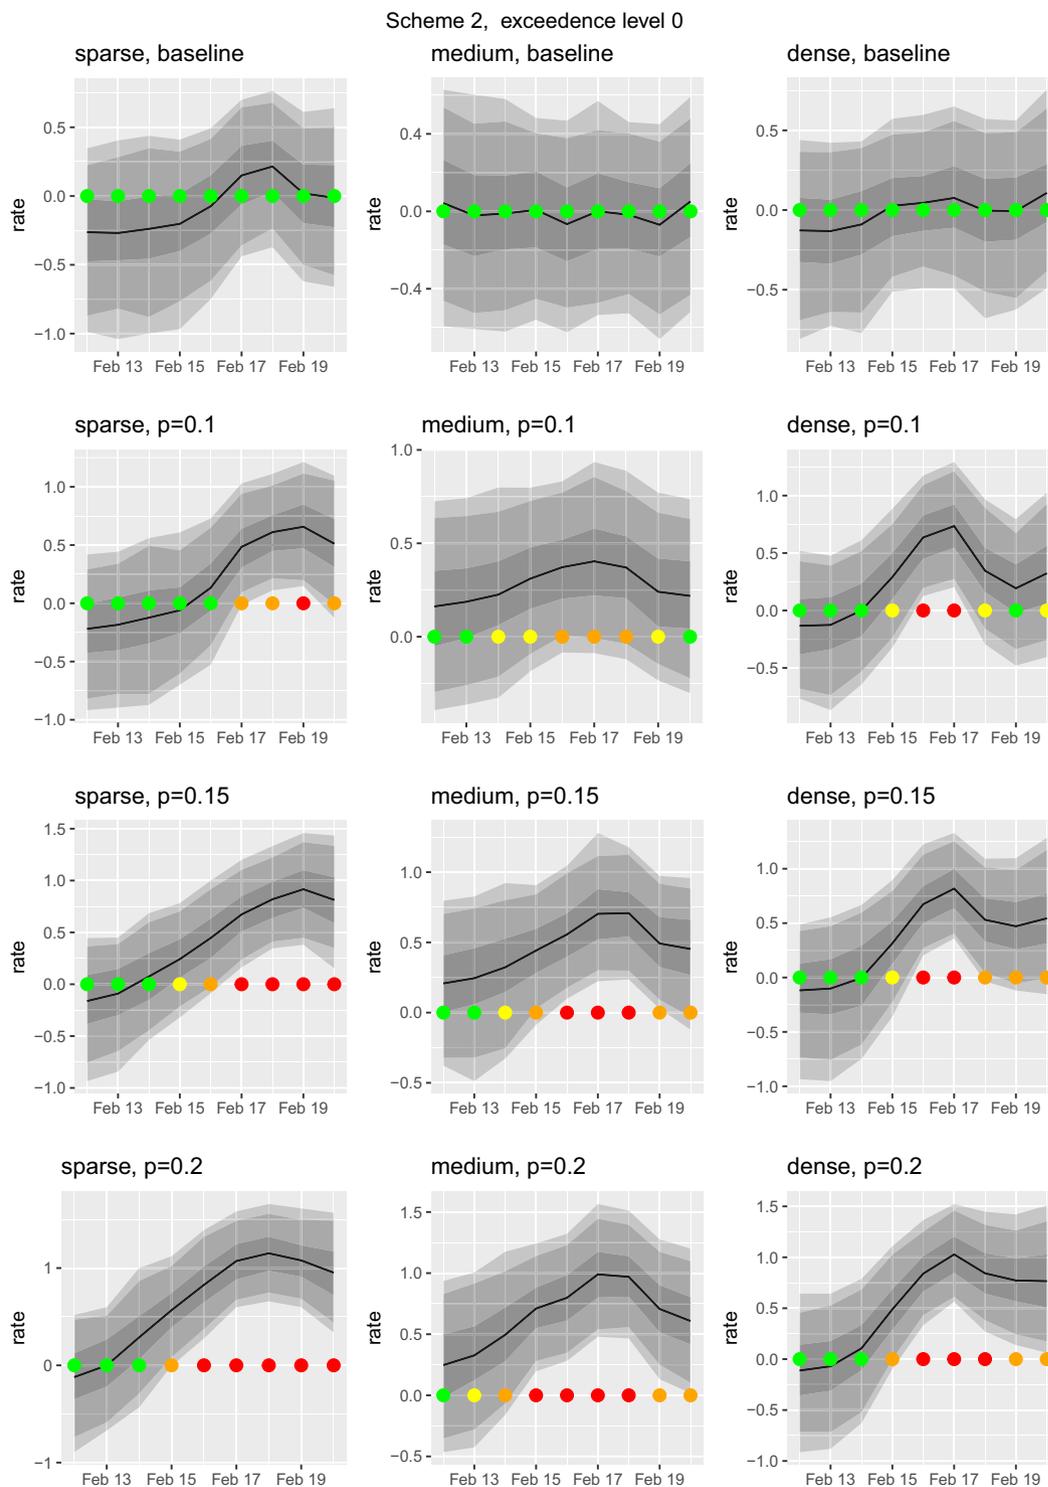


Figure 2. The results from our outbreak simulation study under Scheme 2. The overall layout and format of timeseries plots is the same as Fig. 1, for details see its caption. The simulated outbreak begins on 15th February and the timeseries plots are for premise i at the centre of each region. Here we depict results using Scheme 2, that is premise i and its neighbours, within an 8 km radius, experience an outbreak. Again we choose an exceedence level $l = 0$ (see Supplementary Material for other levels).

during February 2016 (349, 268 and 350 for dense, medium and sparse, respectively) and similar proportions of GI consultations (0.036, 0.055 and 0.042 for dense, medium and sparse, respectively). Using these three sets of premises, we simulated under 15 different scenarios as follows.

quantile	weekday (weekend)	weekday (workday)	gender (male)	purebred	age	age ²	IMD
0.025	-1.8	-1.9	-0.010	-0.210	-0.071	8.1e-05	0.018
0.5	-1.6	-1.8	0.060	-0.120	-0.042	2.2e-03	0.160
0.975	-1.5	-1.7	0.140	-0.038	-0.012	4.2e-03	0.310

Table 2. Regression parameters estimated by outbreak detection model given the baseline data during February 2016; our outbreak simulation results are based on this data. Note, the spatial overall domain of the outbreak simulations is the north west of England hence there is no country effect; see Eq. 7.

- Scheme 1: The outbreak only affects the central premise of each set. For each, we simulate outbreaks of different severities, in which the probability of a case is 0.1, 0.15 or 0.2. This gives a total of 9 scenarios.
- Scheme 2: The outbreak affects the central premise and all of its neighbouring premises. This leads to another 6 separate scenarios, as Schemes 1 and 2 are identical for the sparse set.

Performance evaluation. We use each scenario to generate a simulated set of consultations for February 2016, to which we fit our model using Eq. (1). To assess the capability of our model to detect outbreaks we then use the predictive distribution $S_{i,t}$ from which we compute summary statistics, including exceedence probabilities and times to detection. We set the positive predictive value of the system at $q_0 = 0.9$. We set values of the reporting threshold at $l = 0, 0.3$ and 0.6 . Note that $l = 0$ corresponds to an observed pattern exactly equal to expectation and is analogous to, although formally different from, using statistical rather than clinical significance in hypothesis testing. We do not recommend using $l = 0$ in practice, but use it here only as a benchmark to compare the system's performance under different scenarios. In a genuine application, the threshold value l would be chosen to represent a clinically significant increase in reporting rate, and the positive predictive value q_0 to balance sensitivity against specificity. Note, in this context, that because $S_{i,t}$ is measured on the probit scale, the increase in the fraction of GI cases corresponding to a fixed increase in $S_{i,t}$ necessarily depends on the baseline fraction. For example, if the expected fraction is 0.5, which corresponds to setting $d_{j,i,t}^T \theta = 0$ and $S_{i,t} = 0$ in Eq. (1), then a $\log(2)$ threshold for $S_{i,t}$ represents a fraction $\log(2) = 0.756$ i.e. an increase of 0.256. In contrast, for a baseline fraction 0.1, a $\log(2)$ threshold now represents a fraction 0.278, i.e. an increase of 0.178.

Simulation results. For each of the three regions (sparse, medium, dense) we ran our model a hundred times on the baseline data, where each run had a different random seed; we did not detect any false-positives with $l = 0$. Given the February 2016 baseline data, in Table 2 we report the credible intervals of the regression parameters estimated from the outbreak detection model's MCMC samples.

Our model detected a simulated outbreak in 14 out of the 15 outbreak scenarios when the reporting threshold was set at $l = 0$ (Table 3). The model detected an outbreak on the first day of its actual onset in six scenarios, one day after onset in a further seven scenarios and two days after onset in a further one scenario (Table 3). Alerting timeliness was inversely related to outbreak severity (Table 3).

Figures 1 and 2 give a more detailed illustration of the performance of our outbreak detection methodology in response to a step change in the proportion of cases, for Schemes 1 and 2 respectively and with the threshold value $l = 0$. These figures also illustrate the use of a traffic-light system whereby, rather than fixing a single value for the positive predictive probability, q , we report a categorised value of the exceedence probabilities at each premise on each day to indicate the strength of the evidence for an outbreak.

We focus on the sparse and dense sets of premises since the central premises of these two sets had almost identical numbers of consultations. Recall that under Scheme 1 the outbreak affects only the central premise of each set. Also, the prediction algorithm exploits the estimated spatial correlation amongst the fractions of GI cases at different premises. As a consequence, the system is better able to detect an outbreak at a single premise when this premise does not have close 'outbreak-free' neighbours whose fractions of GI cases are as expected. In effect, the model smooths its predictions over a range corresponding to its estimated correlation range; Fig. 3 shows an example of this phenomenon. This explains why, under Scheme 1 (Fig. 1), the system delivers a stronger detection signal for the sparse than for the dense set. Under Scheme 2 (Fig. 2), the results for the sparse and dense sets are more similar. Also, because the outbreak affects more premises in the medium, and dense sets, their results show generally stronger detection signals than in Scheme 1, as indicated by the increased number of traffic-lights tending towards red in Fig. 2 compared with Fig. 1.

Results of our model's performance using the reporting thresholds $l = 0.3$ and $l = 0.6$ are available in the Supplementary Files; see Table S3 and Figs S1 and S3, and Table S4 and Figs S2 and S4, respectively. For example, given Scheme 1 (density sparse and $p = 0.15$) then: with $l = 0$ we detect an outbreak over the period 16 to 20 February (see Fig. 1); with $l = 0.3$ we also detect an outbreak, albeit less strongly, over the period 17 to 20 February (see Fig. S1 in Supplementary Material); with $l = 0.6$ we do not detect the outbreak (see Fig. S2). An increase in the reporting threshold value l necessarily reduces the probability that an outbreak will be declared and increases its time to detection (Tables S3 and S4, Figs S1–S4). This emphasises that the choice of l must be made in context and is unrelated to the inherent quality of the outbreak detection algorithm.

Setting the probability of a case to 0.1 and with $l = 0$, the model's performance was compared with similar models in the sparse, medium and dense regions:

- (a) *Model without covariates* $\Phi^{-1}(p_{j,i,t}) = S_{i,t}$. All the variation is accounted for by the latent term $S_{i,t}$ so in a real-world application this model would be more prone to false-positives; in the context of Scheme 1 our

Spatial geometry	Extent	Severity (fraction of GI cases)	Timeliness (days to detection since start of outbreak)
Sparse	Confined to central premise	0.1	2
Sparse	Confined to central premise	0.15	1
Sparse	Confined to central premise	0.2	0
Medium	Confined to central premise	0.1	NA
Medium	Confined to central premise	0.15	1
Medium	Confined to central premise	0.2	0
Dense	Confined to central premise	0.1	1
Dense	Confined to central premise	0.15	1
Dense	Confined to central premise	0.2	0
Medium	Extending to neighbouring premises	0.1	1
Medium	Extending to neighbouring premises	0.15	0
Medium	Extending to neighbouring premises	0.2	0
Dense	Extending to neighbouring premises	0.1	1
Dense	Extending to neighbouring premises	0.15	1
Dense	Extending to neighbouring premises	0.2	0

Table 3. Timeliness of a spatio-temporal Bayesian mixed effects regression model at detecting a simulated outbreak in 15 different gastrointestinal disease outbreak scenarios, at a reporting threshold $l = 0$. In one scenario (NA: not applicable) timeliness could not be calculated because no outbreak was detected.

simulations showed this model to be more sensitive. Comparing this model with the full model (Eq. 1) we find they are identical in terms of timeliness but the model without covariates shows more strength of the evidence for the outbreak in that the exceedence probabilities are higher overall.

- (b) *Model without spatial correlation* – Scheme 1. In the presence of the outbreak only occurring at the central premise we found this model to be more sensitive at detecting outbreaks since the surrounding premises will not influence, and hence reduce, the inferred effects of the outbreak at the single central premise. Compared with the full model (with spatial correlation) we find this model to be identical in terms of timeliness for the sparse and dense regions, but the outbreak is now detected in the medium region with a one-day lag. Overall, the exceedence probabilities are higher in all regions.
- (c) *Model without spatial correlation* – Scheme 2. With the outbreak spread over the neighbouring premises, this model was less sensitive as the neighbours did not influence, and therefore support, the detection of the outbreak. In particular we did not detect the outbreak in the medium and dense regions.

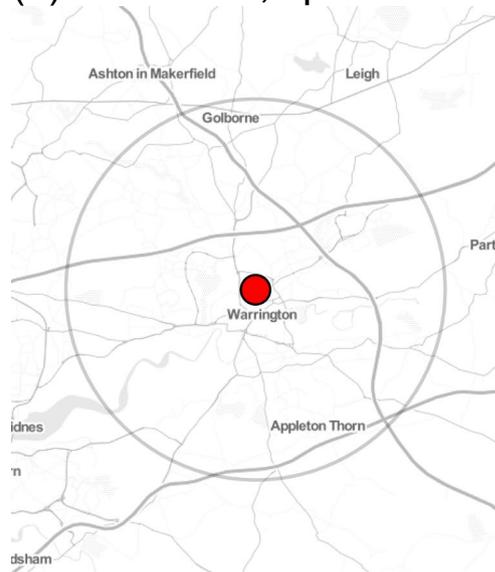
Discussion

Syndromic surveillance systems offer the opportunity to enhance the public and animal health community's ability to detect, and respond quickly to, disease outbreaks⁵. The last decade has seen a growth in the field of disease surveillance in companion animals, notably in the UK^{9,26} and in the USA^{27,29}. However, to the best of our knowledge, this is the first surveillance system that conducts integrated spatio-temporal analysis of data from a national network of veterinary practices so as to enable real-time detection of spatially and temporally localised changes in reporting rate patterns across the network.

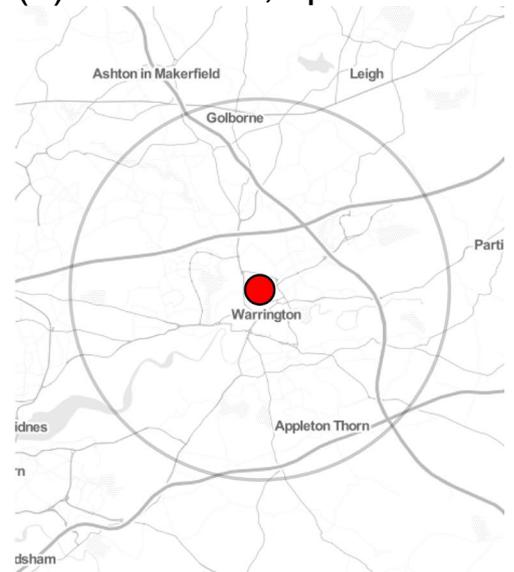
We have illustrated the applicability of our proposed surveillance system using gastrointestinal disease syndrome in dogs and cats as an example. The system is fed with electronic health records (EHRs) collected in real-time through SAVSNET from volunteer veterinary premises across the UK. We applied our system to 15 simulated GI disease outbreaks of varying spatial extent and severity, amongst which the system was able to detect 14 of the 15. Had these been real outbreaks, the proposed surveillance system would have triggered timely investigations, which ultimately would have aided control strategies. The system requires the user to specify a reporting threshold corresponding to an increase in case incidence (reporting rate) that would be considered large enough to be of practical importance. Given this reporting threshold, the system delivers the predictive probability, q , at each location (here, veterinary premise), that the threshold is currently exceeded. Declaring an outbreak when this probability is greater than a specified value q_0 is equivalent to fixing the positive predictive value of the system (per location, per day) at q_0 . Alternatively, reporting the actual value of q gives an indication of the strength of evidence for an outbreak. Increasing the value of the reporting threshold, l , necessarily reduces the value of q and consequently increases the average time to detection of an outbreak at a fixed value of q_0 .

A critical component of a syndromic surveillance system is the application of optimal disease aberration detection methods. Most of the methods used in veterinary and public health surveillance systems are concerned with detecting disease-outbreaks and health-related threats in time rather than in space^{30–38}. However, disease incidences vary naturally in both space and time. Thus, for example, these techniques may be late at detecting outbreaks that start locally when the surveillance region is large³⁹. In contrast, our proposed method has the advantage of being able to directly incorporate data for each individual animal's consultation, including the date of the visit and the location of the pet's owner. In temporal aberration detection algorithms, explanatory variables such as seasonality and day-of-the-week effects would generally be incorporated, but most of these methods cannot easily include individual-level explanatory variables.

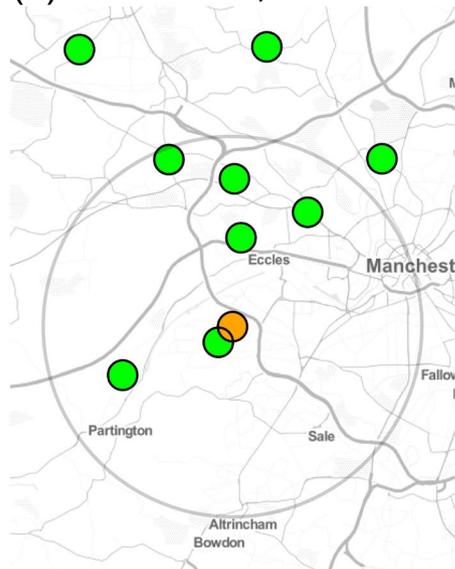
(a) Scheme 1, sparse



(b) Scheme 2, sparse



(c) Scheme 1, dense



(d) Scheme 2, dense

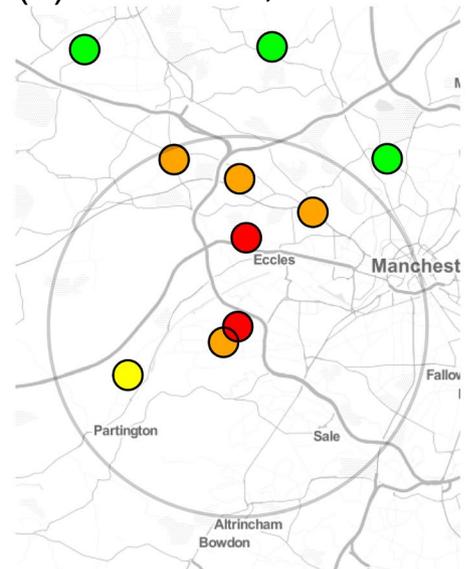


Figure 3. Maps of regions in which we simulated outbreaks where a premise is located at a coloured dot. These premises were selected for illustrative purposes, the actual SAVSNET data shows no indication that they are atypical or that they experienced a genuine outbreak during February 2016. As the base layer we use map tiles by Stamen Design, under CC BY 3.0: data by OpenStreetMap, under ODbL. The premise at the centre of each outbreak region is in the middle of the large light grey circle (8km radius). This figure shows the results of 4 simulations for 17th February 2016 when we use an exceedance level of $l=0$; n.b. the corresponding temporal results are given in Fig. 1. and 2. The top and bottom rows relate to the density of the region, ‘sparse’ and ‘dense’, respectively, and the left and right columns relate to simulation Schemes 1 and 2 respectively. The simulated probability of a case at the premise in the centre of each region is $p = 0.15$. To aid interpretation, we use the traffic-light system described in Fig. 1 caption, as such each coloured circle on the map is derived from the predicted distribution of $S_{i,t}$ at each corresponding premise. Panels (a,c) show when the central premise has neighbours who are not experiencing an outbreak it is less able to detect the outbreak, panel (c), when compared to a premise without neighbours, panel (a). If the neighbours also experience an outbreak the system is then better able to detect this outbreak at central premise, panel (d), compared with when the neighbours did not experience an outbreak, panel (c).

Earlier spatio-temporal aberration detection methods have been introduced by Rogerson^{40,41}. However, these approaches lack measures of uncertainty associated with the identified clusters and are unable to account for covariate information. Also, they are based on an assessment of global pattern change throughout the

geographical area under study, as opposed to our method, which is used to detect the specific geographical location of an outbreak. Prospective space-time scan statistics have also been used in syndromic surveillance systems for the early detection of disease outbreaks^{39,42}. The space-time permutation scan statistic uses only case numbers, with no need for population-at-risk data³⁹ and, in contrast to Rogerson's methods, does operate locally in both space and time. This method may therefore be suitable for setting up surveillance systems in the small animal sector where only case numbers are available. However, it does not acknowledge the uncertainty associated with any identified clusters, cannot easily incorporate continuous covariates, and can only detect outbreaks characterised by excess cases within a specified, regular shaped affected area, for example a circle or ellipse. Also, in our context the number of veterinary premises participating in SAVSNET can change over time due to the ongoing process of recruiting new premises and/or as a result of premises that could potentially stop being part of the project. This can lead to biased results if a space-time permutation model is used, as the method cannot distinguish an increase in cases due to a local population increase versus an increase in disease risk.

Our spatio-temporal model, in conjunction with a Bayesian inferential framework, takes account of all sources of uncertainty in both parameter estimation and prediction, and is able to accommodate spatial, temporal and individual-level covariate information. Other examples of Bayesian approaches include Markov models⁴³, Bayesian information fusion networks⁴⁴ and Bayesian hierarchical models^{45–47}.

An earlier near-real-time syndromic surveillance system in small animals has been developed in the USA utilising EHRs from a similar network of primary care veterinary hospitals²⁹. Briefly, in this approach the daily proportion of patients with a given clinical or laboratory finding was contrasted with an equivalent average proportion from a historical comparison period allowing construction of the proportionate diagnostic outcome ratio (PDOR)²⁹. Our surveillance system builds upon a similar epidemiological metric by modelling the spatio-temporal reporting rate of GI disease in dogs and cats as a proportion of all presentations. The two approaches use different inferential methods: the US study uses confidence intervals for recognising aberrant health events, whilst our approach uses predictive probabilities of exceeding policy-relevant thresholds. A more important difference is that we use a bespoke model that incorporates spatio-temporal covariance structure, with the aim of detecting outbreaks that are spatially and temporally localised without imposing any artificial assumptions on the geometrical shape of an outbreak or the extent of spatial correlation in disease incidence.

Our inferential paradigm of predictive inference within a generalized linear mixed model could equally be applied in purely temporal surveillance settings where the aim is the timely detection of area-wide increases in reporting rate, but in that context we cannot claim the same level of novelty.

Another USA study explored the feasibility of using veterinary laboratory test orders as one of the data sources for syndromic surveillance in companion animals²⁸. The inherent biases associated with the use of laboratory data in veterinary medicine have been described elsewhere^{28,48–50}. However, the results derived from Shaffer *et al.*²⁸ demonstrated the stability and timely availability of test order data for companion animals and the potential of using these data as a basis for outbreak detection. In addition to EHRs from veterinary practices, SAVSNET also receives routine downloads of diagnostic test results from commercial diagnostic laboratories throughout the UK⁹. Although laboratory test results are less timely than test orders, future research is warranted to explore whether the former data could be used to enhance the real-time syndromic surveillance system described here, which is based on real-time data from consultations in small animal premises.

Raising the reporting threshold, l , and/or the required positive predictive probability, q_0 , increases the specificity of the system at the cost of reducing its sensitivity, and conversely. In our analysis of the simulated outbreaks, we chose different reporting thresholds to illustrate the performance of our system. However, in any substantive application, the specified reporting threshold can and should be adjusted so as best to reflect end-users' (i.e. veterinary surgeons in practice) preferred balance between sensitivity and specificity. A pragmatic choice would be to set the threshold to some proportion above the historic average at each premise.

End-users (hereafter "analysts") of a real-time surveillance system will be responsible for receiving system outputs, interpreting them, and if necessary following up on alarms. Therefore, in addition to flexibility, another important attribute of a surveillance system should be that it reports outcomes in an easily interpretable manner. Our system generates outputs in the form of practice-specific time-series and maps that display the spatio-temporal evolution of GI disease risk over an area of interest in a user-friendly manner; see Fig. 3. Additionally, we have illustrated the use of a traffic-light device as a visual aid for analysts to quickly identify potential GI disease outbreaks on a given day at their own premises. The traffic-light device is based on predictive probabilities for exceedence of reporting thresholds that can be tailored to the analysts' needs.

We intend to integrate our daily model-based predictions into the SAVSNET system so as to make them available to each participating premise through their SAVSNET web interface. This implementation will include the other two syndromes with outbreak potential that are currently recorded by SAVSNET (respiratory disease and pruritus). This syndromic surveillance system should be a step towards facilitating the prompt detection and control of health threats in companion animals throughout the UK. In addition, the identified temporal and geographical trends in specific syndromes can be a valuable contribution to the evidence-base when veterinarians are deciding how to treat individual animals in their practice.

One of the challenges of conducting epidemiological studies in the small animal sector is that information about the population-at-risk (in our study defined as the overall population of small animals across the UK or target population) is generally lacking. This makes it impossible to measure parameters typically used in human health surveillance systems, such as the average incidence in a day or period of days. Other methods must therefore be employed to approximate, for instance, an incidence rate ratio. Evidence suggests that in countries with developed pet industries, a high proportion of owned pet animals (pets who may approximate the target population) attend a veterinary surgeon^{51,52}. Therefore, although no single data source can detect all outbreaks that may occur in companion animal populations, EHRs of the kind that are extensively collected from veterinary practices

in many developed countries may be the best available source to include in surveillance activities for increasing our capabilities to detect those outbreaks that result from both endemic and potential emerging pathogens.

One limitation of this study is that the veterinary practices contributing data to our system were selected by convenience, based on their use of a compatible version of PMS, and recruited on the basis of their willingness to take part in the SAVSNET project. Hence, the data used in our system might not be representative of the source population (in our study defined as the overall veterinary-visiting population across the UK). For this reason, we aimed to develop a syndromic surveillance system to detect changes in the relative, rather than absolute, incidence of GI disease presentations in the small animal veterinary premises participating in SAVSNET. Nevertheless, the practices included in the current study were widely distributed around the UK and represented 8.5% of those practices that constituted the source population in 2009⁵¹. Thus, the number and geographical extent of SAVSNET-participating practices is such that changes in the relative risk of GI disease in this large network of premises can act as a proxy for changes in the level of GI disease in the wider source population.

A further limitation relates to missing data. Over the spatial domain and time-period of the simulation we found that 9% of consultations do not record location and 13% do not record breed. As a result, in total about 20% of the data are discarded due to incomplete data, our methodology assumes that these data are missing completely at random so that there is no inherent bias in the spatial distribution of the available data.

Another limitation is that each animal was classified only by its breed-status (purebred or crossbred). As such, we were unable to adjust for breed-specific phenotypes that could have an impact on the incidence of GI disease presentations. However, overall the breed distribution in our study population is consistent with previous studies. Labrador Retriever was the most common dog breed in our population as it is in earlier studies^{8,51,53}. Also, nineteen out of the top twenty-six dog breeds in our study population were also in the top twenty breeds listed by The Kennel Club⁵³. In future work we aim to identify additional means by which breeds can be effectively summarised according to both shared genotype and phenotype.

We are aware that the detection of a high relative risk for GI disease could trigger a false alarm if it is due to a localised decrease in the incidence of diagnosing other syndrome/s and routine veterinary interventions, leading to a higher than expected fraction of GI disease consultations. Conversely, a localised increase in the incidence of diagnosing other syndromes could conceal a genuine GI disease outbreak. If the goal is to detect anomalous patterns of absolute incidence rather than relative risk, then provided that data are available to calculate any changes in the population base of each premise our approach can be modified accordingly, for example by using a Poisson log-linear version of our spatio-temporal mixed model rather than the current binomial probit-linear version.

In order to understand and mitigate shared GI disease aetiologies between humans and animals it would be necessary to develop a 'One Health' surveillance system that integrates human and veterinary healthcare databases. In future work, we intend to adapt the approach described in this paper to human GI disease surveillance by re-calibrating the model against data relating to human GI disease presentations at general practitioner surgeries. A further extension of the approach would then be to a bivariate model for the joint surveillance of veterinary and human GI disease risk. A suitable starting point for this would be to replace the single Eq. (1) by a pair of equations,

$$\Phi^{-1}(p_{j,i,t}) = d_{j,i,t}^T \theta + S_{i,t} \quad (9)$$

and

$$\Phi^{-1}(p'_{j,k,t}) = e_{j,k,t}^T \theta' + S'_{k,t}, \quad (10)$$

where Eqs. (9) and (10) describe the relative risk of GI at veterinary premise i and GP surgery k , respectively. A bivariate model would allow non-zero correlations between the $S_{i,t}$ and $S'_{k,t}$ corresponding to closely located pairs of veterinary premises and GP surgeries.

Conclusions

We have demonstrated the feasibility of a real-time spatio-temporal syndromic surveillance system using as an example small animal veterinary premises in the UK. Our detection algorithm uses Bayesian predictive inference within a spatio-temporal model. The method demonstrated promising performance in detecting simulated outbreaks signals of varying spatial extent and severity at different reporting thresholds. The system is flexible: the reporting threshold of elevated risk and the positive predictive probability per premise and day may be set to whatever levels best meet the needs of a particular application; the system estimates the parameters of the model from historical data rather than imposing specific values for these, and can therefore be re-calibrated to detect outbreaks of any syndrome of interest. A traffic-light system based on exceedence probabilities offers a visual aid to rapid identification of potential outbreaks on a given day at each premise. We intend to implement the system on SAVSNET servers for the early detection of outbreaks in GI and in other syndromes that have outbreak potential and are routinely recorded in SAVSNET.

Data availability

The datasets generated and/or analysed during the current study are not publicly available due to issues of companion animal owner confidentiality, but are available on request from the SAVSNET Data Access and Publication Panel (savsnets@liverpool.ac.uk) for researchers who meet the criteria for access to confidential data. The R scripts used for pre-processing and analysing the data supporting this article can be found as Supplementary Material online. The R package 'precara' developed for pre-processing the data supporting this article is publicly available from the Zenodo repository (<https://doi.org/10.5281/zenodo.812822>)⁵⁴. The R package 'caramellar' developed to run the spatio-temporal model is publicly available from the GitHub repository (<https://github.com/barryrowlingson/caramellar/tree/master>)²⁴.

Received: 10 July 2017; Accepted: 29 October 2017;

Published online: 28 November 2019

References

- Jones, K. E. *et al.* Global trends in emerging infectious diseases. *Nature* **451**, 990–993 (2008).
- Bravata, D. M. *et al.* Systematic review: surveillance systems for early detection of bioterrorism-related diseases. *Ann. Intern. Med.* **140**, 910–922 (2004).
- May, L., Chretien, J. P. & Pavlin, J. A. Beyond traditional surveillance: applying syndromic surveillance to developing settings—opportunities and challenges. *BMC Public Health* **9**, 242 (2009).
- Fricke, R. D. Syndromic Surveillance in *Encyclopedia of Quantitative Risk Analysis and Assessment* (eds Melnick, E.L. & Everitt, B.S.) 1743–1752 (John Wiley and Sons Ltd, 2008).
- Dórea, F. C., Sanchez, J. & Revie, C. W. Veterinary syndromic surveillance: current initiatives and potential for development. *Prev. Vet. Med.* **101**, 1–17 (2011).
- Dupuy, C. *et al.* Inventory of veterinary syndromic surveillance initiatives in Europe (Triple-S project): current situation and perspectives. *Prev. Vet. Med.* **111**, 220–229 (2013).
- Taylor, L. H., Latham, S. M. & Woolhouse, M. E. Risk factors for human disease emergence. *Philos. Trans. R Soc. Lond. B Biol. Sci.* **356**, 983–989 (2001).
- O'Neill, D. G., Church, D. B., McGreevy, P. D., Thomson, P. C. & Brodbelt, D. C. Approaches to canine health surveillance. *Canine Genet. Epidemiol.* **1**, 2 (2014).
- Sánchez-Vizcaíno, F. *et al.* Small animal disease surveillance. *Vet. Rec.* **177**, 591–594 (2015).
- Stavisky, J. *et al.* A case-control study of pathogen and lifestyle risk factors for diarrhoea in dogs. *Prev. Vet. Med.* **99**, 185–192 (2011).
- Smith, S. *et al.* Value of syndromic surveillance in monitoring a focal waterborne outbreak due to an unusual *Cryptosporidium* genotype in Northamptonshire, United Kingdom, June–July 2008. *Euro. Surveill.* **15**, 19643 (2010).
- Hubbard, K., Skelly, B. J., McKelvie, J. & Wood, J. L. Risk of vomiting and diarrhoea in dogs. *Vet. Rec.* **161**, 755–757 (2007).
- Jones, P. H. *et al.* Surveillance of diarrhoea in small animal practice through the Small Animal Veterinary Surveillance Network (SAVSNET). *Vet. J.* **201**, 412–418 (2014).
- Neill, O. *et al.* Prevalence of disorders recorded in dogs attending primary-care veterinary practices in England. *PLoS One* **9**, e90501 (2014).
- Stavisky, J. *et al.* Prevalence of canine enteric coronavirus in a cross-sectional survey of dogs presenting at veterinary practices. *Vet. Microbiol.* **140**, 18–24 (2010).
- Sánchez-Vizcaíno, F. *et al.* Demographics of dogs, cats, and rabbits attending veterinary practices in Great Britain as recorded in their electronic health records. *BMC Vet. Res.* **13**, 218 (2017).
- Department for Communities and Local Government. English indices of deprivation 2010: indices and domains, <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010> (2011).
- The Scottish Government. Part 2 – SIMD 2012 Data – Overall ranks and domain ranks, <http://simd.scotland.gov.uk/publication-2012/download-simd-2012-data/> (2012).
- Welsh Government. WIMD 2011 individual domain scores and overall index scores for each Lower Layer Super Output Area (LSOA), http://gov.wales/statistics-and-research/welsh-index-multiple-deprivation/?lang=en#?tab=previous&lang=en&_suid=1434998682468013747584768796872 (2011).
- Department for Communities and Local Government. English indices of deprivation 2010, <https://www.gov.uk/government/statistics/english-indices-of-deprivation-2010> (2011).
- The Scottish Government. Overview of the SIMD, <http://simd.scotland.gov.uk/publication-2012/introduction-to-simd-2012/overview-of-the-simd/what-is-the-simd/> (2012).
- Welsh Government. Welsh Index of Multiple Deprivation, 2011: Summary report, http://gov.wales/statistics-and-research/welsh-index-multiple-deprivation/?lang=en#?tab=previous&lang=en&_suid=1434998682468013747584768796872 (2011).
- Rue, H. & Held, L. In *Gaussian Markov Random Fields: Theory and Applications* 280 pp (Chapman and Hall/CRC, 2005).
- Rowlingson, B., Giorgi, E. & Hale, A.C. Conditional Auto-regressive Space-Time Model (caramellar). GitHub, <https://github.com/barryrowlingson/caramellar/tree/master> (2017).
- R Core Team. R: language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria, <http://www.R-project.org/> (2017).
- VetCompass. VetCompass, <http://www.rvc.ac.uk/vetcompass> (2016).
- Glickman, L. T. *et al.* Purdue University-Banfield National Companion Animal Surveillance Program for emerging and zoonotic diseases. *Vector Borne Zoonotic Dis.* **6**, 14–23 (2006).
- Shaffer, L.E. *et al.* Early outbreak detection using an automated data feed of test orders from a veterinary diagnostic laboratory in *Intelligence and Security Informatics: Biosurveillance* (eds Zeng, D., Gotham, I., Komatsu, K., Lynch, C., Thurmond, M., Madigan, D., Lober, B., Kvach, J. & Chen, H.) 1–10 (Springer, 2007).
- Kass, P. H. *et al.* Syndromic surveillance in companion animals utilizing electronic medical records data: development and proof of concept. *PeerJ* **4**, e1940 (2016).
- Stroup, D. F., Williamson, G. D., Herndon, J. L. & Karon, J. M. Detection of aberrations in the occurrence of notifiable diseases surveillance data. *Stat. Med.* **8**, 323–329 (1989). discussion 331–322.
- Nobre, F. F. & Stroup, D. F. A monitoring system to detect changes in public health surveillance data. *Int. J. Epidemiol.* **23**, 408–418 (1994).
- Farrington, C. P., Andrews, N. J., Beale, A. D. & Catchpole, M. A. A statistical algorithm for the early detection of outbreaks of infectious disease. *J. R. Stat. Soc. Series A.* **159**, 547–563 (1996).
- Hutwagner, L. C., Maloney, E. K., Bean, N. H., Slutsker, L. & Martin, S. M. Using laboratory-based surveillance data for prevention: An algorithm for detecting salmonella outbreaks. *Emerg. Infect. Dis.* **3**, 395–400 (1997).
- Simonsen, L. *et al.* A method for timely assessment of influenza-associated mortality in the United States. *Epidemiology* **8**, 390–395 (1997).
- Stern, L. & Lightfoot, D. Automated outbreak detection: a quantitative retrospective analysis. *Epidemiol. Infect.* **122**, 103–110 (1999).
- Reis, B. & Mandl, K. Time series modeling for syndromic surveillance. *BMC Med. Inform. Decis. Mak.* **3**, 2 (2003).
- Sonesson, C. & Bock, D. A review and discussion of prospective statistical surveillance in public health. *J. R. Stat. Soc. Series A.* **166**, 5–21 (2003).
- Hutwagner, L. C., Thompson, W. W., Seaman, G. M. & Treadwell, T. A simulation model for assessing aberration detection methods used in public health surveillance for systems with limited baselines. *Stat. Med.* **24**, 543–550 (2005).
- Kulldorff, M., Heffernan, R., Hartman, J., Assunção, R. & Mostashari, F. A space-time permutation scan statistic for disease outbreak detection. *PLoS Med.* **2**, e59 (2005).
- Rogerson, P. A. Surveillance systems for monitoring the development of spatial patterns. *Stat. Med.* **16**, 2081–2093 (1997).
- Rogerson, P. A. Monitoring point patterns for the development of space-time clusters. *J. R. Stat. Soc. Series A.* **164**, 87–96 (2001).
- Kulldorff, M. Prospective time periodic geographic disease surveillance using a scan statistic. *J. R. Stat. Soc. Series A.* **164**, 61–72 (2001).

43. Le Strat, Y. & Carrat, F. Monitoring epidemiologic surveillance data using hidden Markov models. *Stat. Med.* **18**, 3463–3478 (1999).
44. Mnatsakanyan, Z. R., Burkom, H. S., Coberly, J. S. & Lombardo, J. S. Bayesian information fusion networks for biosurveillance applications. *J. Am. Med. Inform. Assoc.* **16**, 855–863 (2009).
45. Sebastiani, P., Mandl, K.D., Szolovits, P., Kohane, I.S. & Ramoni, M.F. A Bayesian dynamic model for influenza surveillance. *Stat. Med.* **25**, 1803–1816. discussion 1817–1825 (2006).
46. Zou, J., Karr, A. F., Datta, G., Lynch, J. & Grannis, S. A Bayesian spatio-temporal approach for real-time detection of disease outbreaks: a case study. *BMC Med. Inform. Decis. Mak.* **14**, 108 (2014).
47. Chan, T.-C. *et al.* Probabilistic daily ILI syndromic surveillance with a spatio-temporal Bayesian hierarchical model. *PLoS One.* **5**, e11626 (2010).
48. Power, C. Passive animal disease surveillance in Canada: A benchmark in *Proceedings of a CAHNet Workshop* (Ottawa, Canada, 1999).
49. Stone, M. The potential for exotic disease syndromic surveillance within veterinary laboratory submissions data in *Proceedings of the Food Safety, Animal Welfare & Biosecurity, Epidemiology & Animal Health Management, and Industry branches of the NZVA, Proceedings of the Epidemiology & Animal Health Management Branch of the NZVA* 91–102 (2007).
50. Vourc'h, G. *et al.* Detecting emerging diseases in farm animals through clinical observations. *Emerg. Infect. Dis.* **12**, 204–210 (2006).
51. Asher, L. *et al.* Estimation of the number and demographics of companion dogs in the UK. *BMC Vet. Res.* **7**, 74 (2011).
52. Lund, E. M., Armstrong, P. J., Kirk, C. A., Kolar, L. M. & Klausner, J. S. Health status and population characteristics of dogs and cats examined at private veterinary practices in the United States. *J. Am. Vet. Med. Assoc.* **214**, 1336–1341 (1999).
53. Kennel Club. Top 20 Breeds 2013–2014. In: Breed registration statistics. 2015, http://www.thekennelclub.org.uk/media/350279/2013_2014_top_20.pdf.
54. Hale, A. C. Processes spatially referenced data (precara). Zenodo Repository, <https://doi.org/10.5281/zenodo.812822> (2017).

Acknowledgements

We wish to thank data providers both in practice (VetSolutions, Teleos, CVS and non-corporate practitioners) and in diagnostic laboratories, without whose support and participation, this research would not be possible. SAVSNET is supported and major funded by BBSRC and BSAVA, with additional sponsorship from the Animal Welfare Foundation. F.S.-V. was fully supported and A.R. partly supported by the National Institute for Health Research Health Protection Research Unit (NIHR HPRU) in Emerging and Zoonotic Infections at University of Liverpool in partnership with Public Health England (PHE), in collaboration with Liverpool School of Tropical Medicine. A.C.H. was funded by the Medical Research Council (MRC), grant number MR/N015266/1. B.R., P.J.D. and A.D.R. were supported by the Wellcome Trust and Department of Health through the Health Innovation Challenge Fund: Theme 5 Infections Response Systems; grant number HICF-T5-354; additionally A.C.H. and S.J.O'B. were also partially supported by HICF-T5-354. S.J.O'B. was also supported by the NIHR HPRU in Gastrointestinal Infections, which is a partnership with PHE in collaboration with the Universities of East Anglia and Oxford and the Quadram Institute. E.G. was funded by the Medical Research Council, grant number MR/M015297/1. The views expressed are those of the authors and not necessarily those of the MRC, the NHS, the NIHR, the Department of Health or Public Health England.

Author contributions

The study was conceived and designed by A.D.R., S.J.O'B., P.J.D., F.S.-V., A.C.H. and B.R. The financial support for the projects leading to this publication was acquired by A.D.R., S.J.O'B. and P.J.D. The data were acquired by A.D.R. and F.S.-V. The data curation was carried out by F.S.-V. The exploratory analysis and data interpretation to inform model design was conducted by A.C.H., F.S.-V. and B.R. The designs of the spatio-temporal model, outbreak simulations and software implementation were performed by A.C.H., B.R. and E.G. The manuscript was drafted by the joint first authors F.S.-V. and A.C.H. The manuscript was revised critically for important intellectual content by P.J.D., S.J.O'B., A.D.R., B.R. and E.G. All authors gave final approval for publication.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information is available for this paper at <https://doi.org/10.1038/s41598-019-53352-6>.

Correspondence and requests for materials should be addressed to A.C.H. or F.S.-V.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2019