

ORIGINAL ARTICLE

Bayesian spatiotemporal modelling for the assessment of short-term exposure to particle pollution in urban areas

Monica Pirani¹, John Gulliver², Gary W. Fuller¹ and Marta Blangiardo²

This paper describes a Bayesian hierarchical approach to predict short-term concentrations of particle pollution in an urban environment, with application to inhalable particulate matter (PM₁₀) in Greater London. We developed and compared several spatiotemporal models that differently accounted for factors affecting the spatiotemporal properties of particle concentrations. We considered two main source contributions to ambient measurements: (i) the long-range transport of the secondary fraction of particles, which temporal variability was described by a latent variable derived from rural concentrations; and (ii) the local primary component of particles (traffic- and non-traffic-related) captured by the output of the dispersion model ADMS-Urban, which site-specific effect was described by a Bayesian kriging. We also assessed the effect of spatiotemporal covariates, including type of site, daily temperature to describe the seasonal changes in chemical processes affecting local PM₁₀ concentrations that are not considered in local-scale dispersion models and day of the week to account for time-varying emission rates not available in emissions inventories. The evaluation of the predictive ability of the models, obtained via a cross-validation approach, revealed that concentration estimates in urban areas benefit from combining the city-scale particle component and the long-range transport component with covariates that account for the residual spatiotemporal variation in the pollution process.

Journal of Exposure Science and Environmental Epidemiology (2014) **24**, 319–327; doi:10.1038/jes.2013.85; published online 27 November 2013

Keywords: Bayesian analysis; exposure modelling; geostatistics; time series; urban particle pollution

INTRODUCTION

Concern over short- and long-term effects of outdoor air pollution on health has led to great efforts to study appropriate exposure assessment methods aimed at quantifying health risks.

Geographic information system (GIS)-based methods like land-use regression models^{1,2} have been successfully applied to estimate long-term (e.g. annual) ambient concentrations,³ but these techniques are not appropriate for short-term modelling as they do not include the influence of both changing source emissions or meteorology.

To provide accurate estimates of short-term air pollution concentrations to use in health effect studies, researchers are increasingly turning to statistical or deterministic dispersion models. The former approach typically considers series of data collected at monitoring sites and characterises these with spatial or spatiotemporal structures; in this context, the Bayesian paradigm has experienced a substantial increase in usage in the past decade.^{4–10} The latter approach simulates the dispersion of air pollution concentrations through deterministic models based on mathematical description of physical–chemical processes taking place in the atmosphere. Because the measurements at ambient monitoring stations can be sparse and irregularly spaced, as well as affected by missing data, the use of deterministic dispersion models has become increasingly popular owing to their more comprehensive coverage over space and time. However, deterministic models are affected by different sources of uncertainty when compared with measurements, because the output depends not only on accurately characterising source emissions,

meteorology and geographical features of the dispersion environment but also on the model configuration options selected by the user (for instance, several numerical models present options to apply diurnal, weekly and monthly profiles to the emission sources).

With respect to the issue of numerical uncertainty in deterministic models, a critical role is played by the ambient measurements as they are frequently used to develop, evaluate and calibrate the air quality models. The process of calibration is somewhat contentious but it is widely accepted that the use of measurements can lead to improved model performance where some inputs are not fully parameterised.¹¹

Recently, approaches to tackle this problem have been focussed on the combination of deterministic model output with observed monitoring data.^{12–18} We follow these current approaches, but differently from the main literature on this topic that is best suited for combining data collected at different spatial resolutions (observed concentrations collected at point level and output from deterministic models at grid level), we suggest a methodology for exposure assessment operating exclusively at point level scale. In particular, we provide an approach, under the Bayesian framework, to obtain enhanced predictions of particle pollution (PM) for use in short-term health effect studies.

In the context of exposure assessment, the successful modelling of PM at a local or city scale is a frequent technical challenge that requires information about regional background pollutant concentrations. This reflects the complexity of ambient PM that comprises of primary particulates arising from local traffic and

¹MRC-PHE Centre for Environment and Health, King's College London, Franklin Wilkins Building, 150 Stamford Street, SE1 9NH London, UK and ²Department of Epidemiology and Biostatistics, Centre for Environment and Health, Imperial College London, School of Public Health, Norfolk Place, W2 1PG London, UK. Correspondence to: Monica Pirani, MRC-PHE Centre for Environment and Health, King's College London, Franklin Wilkins Building, 150 Stamford Street, SE1 9NH London, UK.

Tel.: +44 (0)20 78484044. Fax: +44 (0)20 78484045.

E-mail: monica.pirani@kcl.ac.uk

Received 18 January 2013; accepted 24 September 2013; published online 27 November 2013

non-traffic emissions, and secondary particles formed by atmospheric physical and chemical processes, such as condensation of vapourised material or by-product of the oxidation of gases, mainly during the course of long-range transport of pollutants.

In this paper, we worked with point-referenced time series (daily) and considered several hierarchical models that differently accounted for the relative contribution of regional and local sources affecting the spatiotemporal properties of PM. Specifically, we considered:

- (1) A time-varying latent regional process for capturing the long-range transport of PM. In our study, regional PM concentrations were estimated through direct measurement of rural background concentrations, using an additive approach as suggested by Lenschow *et al.*¹⁹
- (2) A spatial local process for capturing the additional urban and local primary PM component. In our application, a local-scale air pollution dispersion model was used to describe this component.

Moreover, we accounted for selected space- or time-varying factors, which could have a direct influence on the pollution process or could be used as proxies for other unmeasured variables.

We applied our proposed methodology to model inhalable PM in Greater London (UK), namely particles with a diameter $\leq 10 \mu\text{m}$ (PM_{10}), one of the air pollutants of most concern for public health that has been linked to a range of serious cardiovascular and respiratory health effects.^{20–22}

We assessed the predictive performance of the models using a cross-validation approach.

Finally, we compared our approach with the one typically used in the literature on spatiotemporal modelling of air pollution, including random intercepts to account for spatial and temporal dependencies.

MATERIALS AND METHODS

Data Description

The PM_{10} data ($\mu\text{g}/\text{m}^3$) were daily average concentrations (midnight to midnight) collected in the years 2002–2003 (728 days). This period was selected to include several pollution episodes (i.e. periods of elevated PM_{10}) and the 2003 European heat wave.²³ The data were log-transformed to achieve a Gaussian distribution. They came from three sources:

(1) *Mass concentration measurements from the London Air Quality Network (LAQN; www.londonair.org.uk)*: This monitoring network had 76 PM_{10} sites in 2002–2003, with some of these also affiliated with the National Automatic Urban and Rural Network (AURN; <http://uk-air.defra.gov.uk>). Out of these sites, we selected 45 for which the proportion of missing data, in each year, did not exceed 20%. The missing observations were assumed to be missing at random. The average proportion of missing data for the 45 sites in the study period was 5.1% (range: 0–17.4). The mean distance between the selected sites was 16,967 m (range: 657–45,298). The majority of measurements were made using the Tapered Element Oscillating Microbalance method using TEOM 1400a and 1400ab analysers (Thermo Scientific). These instruments are known to underestimate the PM_{10} concentrations owing to losses of semivolatile constituents (such as ammonium nitrate and organic aerosols);^{24,25} therefore, a default adjustment factor of 1.3 was applied.²⁶ Eight sites were equipped with Beta Attenuation Monitors (Met-One BAM), where a correction factor of 0.82 was applied according to the results of UK trails.²⁷

(2) *Output from the high spatial resolution Air Dispersion Modelling System (ADMS-Urban; CERC, Cambridge, UK)*:^{28,29} ADMS-Urban was used to represent the local primary component of PM_{10} . It simulates the dispersion into the atmosphere of pollutants released from road traffic, industrial and domestic sources across urban areas and integrates emissions inventories with meteorological data. We obtained emissions factors from the London Atmospheric Emissions Inventory (LAEI), which contains data on road network geometry comprising about 60,000 individual road links attributed with traffic flows and composition. Roads are represented as line sources in ADMS-Urban with a spatial accuracy of <1 m. Point and area source emissions were aggregated in the LAEI to

1 km resolution grids. This is a relatively quick method for modelling poorly defined or diffuse sources in the dispersion model. Dispersion from road sources used a Gaussian plume model with a non-Gaussian plume profile in convective conditions to account for the skewed structure of the vertical component of turbulence. Grid sources were modelled using a simple trajectory model. Output from both models was combined to predict pollutant concentrations at point locations, namely air pollution monitoring sites. Meteorological data included in ADMS-Urban were obtained from the British Atmospheric Data Centre for the nearest site.

(3) *Mass measurements from rural monitoring sites*: Background concentrations, as proxy of the long range transport of PM_{10} , were sourced from two rural monitoring stations, approximately equidistant from London: Harwell (near Didcot, Oxfordshire; 81 km north-west of London, towards the West Midland conurbation) and Detling (Kent; 50 km south-east of London towards continental pollution sources areas). The Pearson's correlation coefficient between the two time series was 0.6 (P -value <0.001), indicating that these measurement sites provided different information about the long-range transported air pollution affecting London.

In addition, we considered the following set of covariates:

(1) *Type of site*, which accounted for different environmental conditions. The LAQN monitoring sites were classified into different types, depending on their location. Of the 45 sites selected for the study, 8 were suburban sites (located in residential areas on the outskirts of London), 13 were urban background sites (located away from major sources and broadly representative of city-wide background concentrations), 20 were roadside sites (located from 1 to 5 m from a major carriageway) and 4 were kerbside sites (located within 1 m of a major road carriageway).

(2) *Day of the week*, which accounted for unknown changes in emissions between weekdays and weekend days, because emission inventories are not time-varying but only contain annual totals. The indicator variable for day of the week was categorised as Monday–Friday, Saturday and Sunday or Public Holiday.

(3) *Average daily temperature* to describe seasonal changes in chemistry between primary and regional secondary PM_{10} . Other meteorological variables were not considered as these are used in the ADMS-Urban model; however, this does not include secondary PM_{10} formation, and thus daily mean temperature was used as a surrogate for such processes. Over the 2002–2003 years, the average temperature, recorded at London Heathrow, was 11.9°C , with daily mean ranging between -1.3 and 28.2°C .

Exploratory Data Analysis

Figure 1 presents the geographical location of the 45 monitoring sites across Greater London by site type. As we found little difference between the PM_{10} concentrations at suburban and urban background sites, we aggregated these two categories.

Figure 2 shows the correlation of daily data for pairs of monitoring sites as a function of their separation distance. The correlations were generally high, also over long distances ($\geq 30,000$ m), indicating that factors other than distance may have a role in explaining the spatial variability of PM_{10} levels.

Figure 3 presents the daily levels of PM_{10} across the 45 monitoring sites sorted from the top to the bottom by decreasing longitude, during the 2 year in study.³⁰ The daily values are displayed according to the tertiles computed on the global data set to ensure the comparability of the time series and assigned to low (brown), medium (pale green) and high (green) categories of PM_{10} concentrations. The bottom of the plot shows the daily median values across all the time series. The PM_{10} pollution episodes that London experienced during February, March, April, August, September and November 2003 are clearly visible. These episodes were mainly caused by secondary PM_{10} from distant sources, with summer episodes also being linked to photochemistry.³¹ The November 2003 episode was associated with Guy Fawkes Night fireworks and bonfires.³²

The analysis via cross-correlogram of the time series of PM_{10} concentrations observed in Greater London and the local component of PM_{10} captured from ADMS-Urban output, presented in Figure 4, shows that the correlation was relatively high and positive at lag 0 (same day pollution levels), suggesting that the numerical model captured the time variation of PM_{10} observed at monitoring sites.

Modelling Approach

We denoted $Z(t,s)$ as the log-transformed daily PM_{10} concentrations, with $t = 1, \dots, T = 728$ (days) and $s = 1, \dots, n = 45$ (sites of the pollutant monitoring




Figure 1. Location and siting characteristics of the air quality monitoring sites in Greater London selected for the study.




Figure 2. Correlation between pairs of monitoring sites as a function of their separation distance.

network). We assumed that the observed monitoring data were characterised by measurement error defined by a zero-mean Gaussian white noise process. We specified a Gaussian likelihood for $Z(t,s)$:

$$Z(t,s) \sim N(\mu(t,s), \sigma^2(s)) \quad (1)$$

where $\mu(t,s)$ represents the mean process driven by covariates varying over space and time and $\sigma^2(s)$ is the site-specific measurement error variance.⁴ We considered a class of different nested statistical formulations for the mean space-time process, $\mu(t,s)$, that differently accounted for factors affecting the spatiotemporal properties of particle concentrations.

Model I represented a simple statistical structure where the daily measurements at each monitoring site were assumed to be a function of a residual mean concentration across the urban area and a latent pollutant process described by the long-range transported proportion of particulate. The time-varying latent regional process was included assuming that concentrations at the city scale derive largely from information borrowed from rural measurements. It assumed the form:

Model I:

$$\mu(t,s) = \alpha + \mu_{lt}(t) \quad (2)$$

where α is the residual intercept and $\mu_{lt}(t)$ represents the mean of the latent process.

In particular, let j denote several available rural background monitoring sites around the metropolitan area, with $j = 1, \dots, J$. We assumed that the time series of pollution data from the rural monitoring sites are a reflection of an underlying long-range transport of particles into the urban area, measured with error:

$$l_{rt}(t,j) \sim N(\mu_{lt}(t), \sigma_{lt}^2(j)) \quad (3)$$

In our application, this latent process was driven by the two time series of PM₁₀ measured at the Harwell and Detling rural background sites ($j = 1, 2$).

This simple model accounted for the temporal variability of the pollution process, but did not incorporate a spatial structure. The model describes the main hypothesis in the definition of air pollution exposure in ecological time series studies, where the pollution estimates for a given study region, are generally free from a spatial dimension, although these studies typically use averaged ambient pollutant levels from one or more background monitoring stations to represent the exposure experienced by a study population.

Model II added to the constant, α , the local city primary PM₁₀ component described by ADMS-Urban, $\ell(t,s)$:

Model II:

$$\mu(t,s) = \alpha + \beta_1(s)\ell(t,s) \quad (4)$$

A spatially varying coefficient model³³ was used for this component to capture the effect of site. We assumed $(\beta_1 = \beta_{1,1}, \dots, \beta_{1,n})^T$ to be distributed according to a zero-centred multivariate Gaussian distribution $\beta_1 \sim MVN(0, \omega^2 H(\phi))$, where $\omega^2 > 0$ is the spatial effect variance parameter and H represents the spatial correlation matrix. In this paper, $H_{ss'}$ is




Figure 3. Daily particle concentrations for the 45 monitoring sites sorted from the top to the bottom by decreasing longitude (from west to east). Each time series is assigned to low (brown), medium (pale green) and high (green) category of pollution levels (i.e. tertiles based on data from all the 45 time series); missing data are denoted by the colour white. The bottom panel shows the daily median concentrations across the time-series.




Figure 4. Cross-correlogram between the time series of particle concentrations in Greater London and the ADMS-Urban output (on log-scale). The graph shows that at lag 0 (same day pollution levels), there is a positive contemporaneous correlation between observed data at monitoring sites and ADMS-Urban output.

described by an exponential function³⁴ $f(d, \varphi) = \exp(-\varphi d)$, where $d = \|s - s'\|$ and $\| \cdot \|$ indicates the Euclidean distance between two generic sites s and s' , and φ is the (non-negative) decay parameter that represents the rate of decline of spatial correlation among sites over distance. This spatial structure for the ADMS-Urban output provided a realistic representation of the spatial variability observed in the explorative analysis. However, we would expect a poor performance of this model as it did not account for the temporal variability in the pollution process.

Model III included both the regional and the local primary PM_{10} components:

Model III:

$$\mu(t, s) = \alpha + \mu_{lt}(t) + \beta_1(s)\ell(t, s) \quad (5)$$

Model IV was performed to explore the effect of the set of covariates (without regional and local PM_{10} components):

Model IV:

$$\mu(t, s) = \alpha + \beta_{2, type(s)} + \beta_{3, dow(t)} + \beta_4(t)temp(t) \quad (6)$$

where *type* is the type of site, *dow* is the day of the week and *temp* is the temperature. In particular, we used site type to represent possible difference in concentration levels, as road and kerb sites are likely to have higher concentrations as they are closer to traffic source of pollution;

daily mean temperature to describe chemical processes affecting local PM₁₀ concentrations that are not considered in local-scale dispersion models and day of week to account for time-varying emission rates that are not described in emissions inventories. In Eq. (6), the fixed-effects coefficients β_2 and β_3 are unknown parameters for the variables site type and day of the week. The vector $\beta_4(t) = (\beta_4(1), \dots, \beta_4(T))^T$ is the dynamic parameter associated with the temperature built according to a Gaussian second-order random walk (RW2), which was found provide the best smoothness prior for this variable. It assumed the form: $\beta_4(t+2) \sim N(2\beta_4(t+1) - \beta_4(t), \sigma_v^2)$ for $t = 1, \dots, T-2$, where σ_v^2 is the variance. A RW2 acts as a smoothness prior based on the second difference and penalises deviations from a linear trend.³⁵ This prior, for regular time-point, provides enough flexibility because of its invariance under addition and it is computationally convenient because of its Markov properties.³⁶ The choice of this prior followed an initial explorative analysis where we found that the relationship between temperature and PM₁₀ concentrations (not shown) was potentially well described by a cubic smoothing spline. The RW2 is a discrete-time analogue of a cubic smoothing spline.³⁷

Model V finally represented the full model that accounted for the regional and local PM₁₀ components and for the covariates:

Model V:

$$\mu(t, s) = \alpha + \mu_{\text{int}}(t) + \beta_1(s)\ell(t, s) + \beta_{2, \text{type}(s)} + \beta_{3, \text{dow}(t)} + \beta_4(t)\text{temp}(t) \quad (7)$$

Parameter Priors and Implementation

A Gaussian prior distribution with mean zero and variance 10^2 was assigned to the intercept α , and to the fixed-effects coefficients β_2 and β_3 . To ensure identifiability, we fixed the first category of these two parameters as zero ($\beta_{2,1} = 0$ and $\beta_{3,1} = 0$). The same Gaussian prior was chosen for the mean of the latent background process. Weakly informative inverse-Gamma hierarchical priors were specified for the error variance variables $\sigma^2(s) \sim \text{IG}(a, b)$, $s = 1, \dots, n$, and $\sigma_{\text{int}}^2(j) \sim \text{IG}(c, d)$, $j = 1, \dots, J$, setting the hyperpriors (a, b, c, d) as $\text{IG}(1, 0.1)$. Similarly, inverse-Gamma priors were specified for the between-site variance component ω^2 and the variance of the RW2 σ_v^2 with hyperparameters $\text{IG}(1, 0.1)$.

We assumed a discrete uniform prior distribution for the decay parameter ϕ as suggested by Diggle and Ribeiro,³⁸ with range chosen based on prior beliefs about the minimum and maximum correlation at the smallest and largest distances. Typically, locations close in space are assumed to be characterised by a stronger degree of correlation, but we did not want to assume a strong prior on it and we allowed for a range of correlation between 0.10 and 0.99. For large separation distances, we specified a range between 0.01 and 0.65.

The models were implemented in WinBUGS,³⁹ a freely available software to perform Bayesian inference via Markov chain Monte Carlo (MCMC) simulative method.⁴⁰ Two parallel MCMC chains with different starting values were run for each model. We ran 60,000 iterations with 50,000 burn-in and thinned the Markov chains by a factor of 10, resulting in samples of size 2000 to estimate the posterior distributions for the parameters of interest. Posterior correlation was reduced by a grand mean centring of the covariates.⁴¹ Convergence was assessed by checking the trace plots of the samples, the estimated kernel density plots, the autocorrelation functions, and a Monte Carlo errors <5% of the posterior standard deviation.

Comparison with Models Implemented with Varying Intercepts

The model formulation proposed in our paper deviates from the standard spatiotemporal statistical models that include varying intercepts (baseline concentrations) that are spatially or temporally correlated.^{4,6,8,14,16} The most common setting assumes that the spatial and temporal dependences are introduced into the modelling in the form of random effects. Thus, pollution concentrations characterised by a Gaussian likelihood are typically related to a trend surface model together with additive-independent random spatiotemporal effects that in a simple implementation can assume the form:

$$\mu(t, s) = \beta X(t, s) + \theta(t) + \eta(s) + \varepsilon(t, s) \quad (8)$$

Here, β is a vector of regression coefficients associated with the covariates $X(t, s)$. The residual is partitioned into a temporal, $\theta(t)$, a spatial, $\eta(s)$, and an independent process $\varepsilon(t, s)$, which is Gaussian with zero mean and $\sigma^2(s)$ variance. As a comparison with our approach, we have considered a model implementation within this classical framework using the same set of data.

We developed five nested hierarchical structures that incorporated separable random space and time effects.

The first model was given by:

Model Ia:

$$\mu(t, s) = \theta(t) + \eta(s) \quad (9)$$

The parameters $\theta(t) = (\theta_1, \dots, \theta_T)^T$ should capture the residual temporal dynamics characterising the pollutant process. This temporal process was described using an autoregressive first-order non-stationary model as daily dependence in air particulate concentrations can be expected,⁸ and was built as $\theta(t+1) \sim N(\theta(t), \sigma_\theta^2)$ for $t = 1, \dots, T-1$. The term $\eta(s) = (\eta_1, \dots, \eta_n)^T$ represents a spatially varying intercept that we assumed described by a zero-centred Gaussian process with variance σ_η^2 and an exponential correlation function that depend upon the intersite distance and the parameter ϕ quantifying the correlation decay.

Model IIa also included the latent regional process defined as in Eq. (3):

Model IIa:

$$\mu(t, s) = \theta(t) + \eta(s) + \mu_{\text{int}}(t) \quad (10)$$

Model IIIa added to the random effects of the urban local component of PM₁₀ described by ADMS-Urban:

Model IIIa:

$$\mu(t, s) = \theta(t) + \eta(s) + \beta_1(s)\ell(t, s) \quad (11)$$

The space-varying slope $\beta_1 = (\beta_{1,1}, \dots, \beta_{1,n})^T$ was built according to a Bayesian isotropic kriging,³⁴ as specified in our main analysis.

Model IVa incorporated both the long-range and the local components of PM₁₀:

Model IVa:

$$\mu(t, s) = \theta(t) + \eta(s) + \mu_{\text{int}}(t) + \beta_1(s)\ell(t, s) \quad (12)$$

Model Va included exclusively the spatiotemporal random intercepts and the covariates type site, day of the week and daily mean temperature:

Model Va:

$$\mu(t, s) = \theta(t) + \eta(s) + \beta_{2, \text{type}(s)} + \beta_{3, \text{dow}(t)} + \beta_4(t)\text{temp}(t) \quad (13)$$

Similar to the main analysis, we also tried to implement a full model as:

Model VIa:

$$\mu(t, s) = \theta(t) + \eta(s) + \mu_{\text{int}}(t) + \beta_1(s)\ell(t, s) + \beta_{2, \text{type}(s)} + \beta_{3, \text{dow}(t)} + \beta_4(t)\text{temp}(t) \quad (14)$$

However, it resulted overparameterised and yielded implausible predictions, and thus we decided not to present the results from this model.

Models Ia–Va were specified assuming for the variance parameters σ_θ^2 and σ_η^2 inverse-Gamma priors $\text{IG}(1, 0.1)$. The other priors were specified as in the main analysis.

Performance Assessment

To compare our models, we partitioned the monitoring network into three sets of sites following these steps: (i) we stratified the 45 sites by type (urban/suburban, roadside and kerbside sites), (ii) we chose a random sample of nine sites, representative of the entire network (with respect to the number of sites of each type) as validation data for testing the models, and (iii) we retained the other 36 sites as training data to fit the models. We repeated steps (i)–(iii) three times (so each site entered into the validation data once).

To evaluate the predictive performance of the models, we compared the predicted PM₁₀ concentrations against the observed measurements on the validation set via the following indices: the empirical coverage of 90% credible intervals (90% CI) coupled with their average length, the squared correlation coefficient (R^2) and the root mean square error (RMSE).¹⁴ Lower values of RMSE indicate more similarity among observed measurements and predicted values. To obtain these indices, for each model we used the full posteriors from each Markov chain and we combined the predicted values from the three sets. This same procedure was used to summarise the results for the parameters evaluation.

Sensitivity Analysis

Sensitivity analysis was performed in order to:

(1) Assess the performance of our modelling approach in urban environments that have a monitoring network less dense than in London. The EU Air quality directive (2008/50/EC) stipulates the minimum population-dependent measurement requirements for EU cities. With 36 European cities with populations above 1 million and 9 above 2 million,⁴²

we considered that testing the methodology on a sample of 10 measurement sites (matching the minimum number of monitoring sites for a city of 2.75 million population) would provide an assessment of applicability in a typical city. A city of 2.75 million would be smaller than the total area of Greater London. To this end, we considered the north-west boroughs in Greater London only and selected 10 sites as training set and 3 sites as validation set, representative of 3 site types, following the methodology described for the main analysis.

(2) Investigate whether results remained essentially unchanged in the presence of different prior distributions. We considered commonly used inverse-Gamma priors for the variance parameters (measurement errors) $\sigma^2(s)$ and $\sigma^2_{lr}(j)$: $IG(0.5, 0.0005)$ and $IG(0.1, 0.1)$. For the spatial effect variance parameter, ω^2 , and the random walk variance parameter, σ^2_v , we tested the prior $IG(0.001, 0.001)$.

RESULTS

Predictive Performance

Table 1 shows the cross-validation summary statistics. The results are reported on the original scale correcting for bias after logarithmic transformation.⁴³ Moving from model I to model V, we noted a progressive improvement in the prediction capability, with exception of model II. However, we found that the validation indices improved heavily when the site-specific local component, described by ADMS-Urban output, was included in addition to the regional background component (as an example, the RMSE decreased from 11.11 for model II to 5.11 for model III). The incorporation of the selected covariates in models IV and V produced an additional increase in the cross-validation performance.

Figure 5 shows the Taylor diagrams^{44,45} for the models, over (a) the whole study period and (b) a 2003 heat-wave event (days from 4 to 13 August 2003). This diagram represents a useful method for evaluating predictive performance, as it visualises simultaneously the centred RMSE (it is centred because the mean values of the observed and predicted data are subtracted first), the correlation coefficient (R) and the standard deviation of the observed and

predicted values. In detail, the observed variability (i.e. the standard deviation) is plotted on the x-axis (specifically, the magnitude of the variability is measured as the radial distance from the origin of the plot), R is shown on the grey arc, whereas the RMSE is indicated by the concentric brown dashed lines emanating from the observed point. The Taylor diagram performed on the entire study period (plot a) showed a quite similar performance of the models from 3 to 5, with model V being the best as presenting the highest correlation, the least RMSE and a reasonable similar variability compared with the observations, and the poor performance of model II was also confirmed. However, the Taylor diagram obtained on a 10-day heat-wave event (plot b) to assess how the models performed in capturing these events, pointed out differences, with models II and model V performing worst in comparison to models I and III. This result could be explained by the fact that the heat-wave events of 2003 were dominated by the long-range transport component.

Predictive Performance of Models Implemented with Varying Intercepts

Table 2 presents the predictive ability of the models implemented using the classical approach given by space- and time-varying intercepts. Generally, the validation indices showed slightly worst values when compared with the cross-validation results from our modelling approach. However, for model IIIa including the spatiotemporal random effects and the urban local component of PM, we found lower prediction errors in comparison to model II of our main analysis. This result confirmed that without temporal dependencies, the predictive capability of ADMS-Urban yield poor performance.

Parameters Evaluation

In our modelling approach, we found that the time-varying latent regional process described by $\mu_{lr}(t)$ had similar behaviour in

Table 1. Predictive performance by model (on original scale).

Models	Average width 90% CI	Coverage 90% CI	RMSE	R^2
Model I	23.67	0.91	5.26	0.58
Model II	45.55	0.88	11.11	0.04
Model III	21.51	0.91	5.11	0.61
Model IV	22.20	0.89	5.04	0.61
Model V	20.40	0.89	4.75	0.63

Abbreviations: CI, credible intervals; RMSE, root mean square error.

Table 2. Predictive performance of the models implemented using spatiotemporal varying intercepts (on original scale).

Models	Average width 90% CI	Coverage 90% CI	RMSE	R^2
Model Ia	28.58	0.89	7.37	0.64
Model IIa	29.90	0.88	7.58	0.64
Model IIIa	28.43	0.92	6.84	0.65
Model IVa	28.59	0.91	6.89	0.64
Model Va	27.18	0.91	6.05	0.64

Abbreviations: CI, credible intervals; RMSE, root mean square error.




Figure 5. Taylor diagrams showing the predictive performance of the five hierarchical models related to: (a) the entire period of study and (b) a 2003 heat-wave event (from 4 to 13 August 2003).

Table 3. Posterior mean and 90% CI for the fixed effects and for the variance parameters by model (on log-scale).

Parameters	Model I		Model II		Model III		Model IV		Model V	
	Mean	90% CI	Mean	90% CI	Mean	90% CI	Mean	90% CI	Mean	90% CI
α (intercept)	3.243	3.242, 3.244	3.315	3.309, 3.322	3.251	3.252, 3.253	3.325	3.302, 3.347	3.253	3.251, 3.254
$\beta_{2,2}$ (road site) ^a	—	—	—	—	—	—	0.185	0.179, 0.192	0.143	0.142, 0.144
$\beta_{2,3}$ (kerb site) ^a	—	—	—	—	—	—	0.282	0.269, 0.294	0.283	0.281, 0.284
$\beta_{3,2}$ (Saturday) ^b	—	—	—	—	—	—	−0.215	−0.276, −0.157	−0.032	−0.033, −0.030
$\beta_{3,3}$ (Sunday or Public Holiday) ^b	—	—	—	—	—	—	−0.201	−0.335, −0.074	−0.080	−0.082, −0.079
$\sigma^2(s)$ (range among sites of the posterior mean of variance)	0.061–0.202		0.163–0.168		0.038–0.074		0.048–0.052		0.033–0.050	
ω^2 (spatial effect variance for the local PM component)	—	—	0.066	0.024, 0.152	0.041	0.040, 0.042	—	—	0.042	0.041, 0.043
σ_v^2 (second order random walk variance for the temperature)	—	—	—	—	—	—	1.111	0.950, 1.300	0.006	0.005, 0.007

Abbreviations: CI, credible intervals.

^aReference category: $\beta_{2,1}$ (suburban/urban site).^bReference category: $\beta_{3,1}$ (weekday).**Table 4.** Predictive performance by model obtained in the sensitivity analysis (on original scale).

Models	Coverage 90% CI	Average width 90% CI	RMSE	R^2
Model I	0.93	31.52	6.91	0.52
Model II	0.87	47.01	11.36	0.02
Model III	0.92	29.62	6.65	0.57
Model IV	0.89	28.54	6.65	0.53
Model V	0.88	23.29	5.38	0.61

Abbreviations: CI, credible intervals; RMSE, root mean square error.

models I, III and V. However, a visual inspection of the plot of the posterior mean of $\mu_{irt}(t)$ (not shown) pointed out a more evident daily variability in model V.

The range (on log-scale) of the posterior mean of the spatial coefficients, $\beta_1(s)$, associated with ADMS-Urban output, varied in model II from 0.005 to 0.333, in model III from 0.005 to 0.371, whereas in model V this ranged from −0.001 to 0.238. This suggested a weaker effect of the local PM₁₀ component when the covariates were included in model V.

Through the analysis of the decay parameter, ϕ , we found coherent results in models II, III and V, across all sets, for the spatial correlation among sites. Specifically, we observed high correlation at minimum distance between sites, ~ 0.97 , that decayed progressively, being ~ 0.50 at mean distance, and ~ 0.24 at maximum distance.

Table 3 presents the posterior mean estimates and their 90% CI for the fixed effects and for the variance parameters. The residual mean concentration, α , remained constant among the models. Instead, we found a strong effect of the site type described by β_2 , indicating that PM₁₀ concentrations were greater for road and kerb sites than for suburban/urban sites, as expected. A negative relationship was estimated between PM₁₀ and day of the week (described by β_3), as the concentrations were lower on the weekends than on weekdays.

The effect of the temperature on PM₁₀ showed a considerable variability, especially for coefficients related to the spring days. The RW2 variance parameter σ_v^2 was definitely bigger in model IV in comparison to model V.

Finally, with the exception of model II, we noted a progressive decrease in the measurement error variance across the models. This reduction underlined the contribution given by the adjustment for covariates to explain part of the variability in the estimated PM₁₀ concentrations.

Sensitivity Analysis Results

Table 4 describes the results related to the predictive ability of our modelling approach on a restricted number of monitoring sites in north-west London. We found that the indices were consistent with those reported in the main analyses (Table 1).

We performed also an assessment of the sensitivity of findings to prior details and these analyses showed that the results were quite robust to these choices.

DISCUSSION

We have presented a Bayesian spatiotemporal approach for modelling particulate pollution concentrations in urban area for health risk studies. We combined air monitoring data with the output from a local-scale air pollution model and explicitly solved the problem of incorporating regional pollution concentrations within the city-scale assessment. Moreover, we assessed the effects of covariates to account for the residual spatiotemporal variation of particle concentrations. We evaluated the predictive performance of these statistical structures through a robust procedure of cross-validation that allowed us to compare the daily predictions with the observed PM₁₀ concentrations within three validation sets of sites, which represented different urban environment (i.e. site types).

In particular, we applied our modelling approach to predict PM₁₀ concentrations in Greater London, using a latent regional pollution process derived by rural sites to describe the long-range transport PM₁₀ component and the output from ADMS-Urban to capture the local primary PM₁₀ component.

ADMS-Urban is widely used for estimating urban-scale air pollution for regulatory purposes and in epidemiological air pollution studies.^{46,47} We found that the exclusive use of ADMS-Urban to predict the PM₁₀ concentrations produces poor results. So far, although the inclusion of ADMS-Urban, in addition to a regional latent process, increases the predictive performance of the models, we suggest that the use of this deterministic output to measure the population exposure to PM in short-term epidemiological studies should be enhanced with the combination of other information sources characterising the study area, such as site-type or time-varying emission factors linked to day of the week, as evidenced by the strength of the covariates in our models.

In this implementation, we adopted an indicator variable for site types that is actually quite crude. The use of a more localised index of sites better reflecting land use and building geometry (canyon orientation for example) by utilising GIS techniques may further improve our model performance.

The final aim of our study was to assess air PM exposure models to use in short-term health effects studies in London. We therefore worked with the dense monitoring network available owing to the city size and the legal structures for local air quality management. To assess the applicability of our approach in urban environment with smaller number of monitoring sites, we performed a sensitivity analysis restricting the study area to a part of London matching the minimum requirements in EU directives. The results suggested that our approach will also perform well in smaller urban environments with more sparse monitoring networks, which are typical of many European cities.

Methodologically, the models presented here deviate from the standard space–time statistical modelling approach, which typically presents varying intercepts.^{4,6,8,14,16} As we were including in our models a set of covariates characterised by spatial and temporal variation, we assumed only time- and space-varying regression coefficients. To assess the plausibility of our approach in comparison with a classical modelling scenario, we developed five models, with independent spatiotemporal random effects. We assessed the predictive capability of these structures, and we found that our methodology, applied in an urban environment, performed better than the classical approach. This evidence suggests that, in context where local and urban primary emissions together with regional background data are not available, the inclusion in the models of independent error distributions is able to capture spatial and temporal dependencies. However, in context of analysis, where the researchers can perform extra modelling efforts, our proposed models perform better than a classical approach.

Finally, the hierarchical methodology that we proposed in this study provided a flexible way to model daily PM. This approach could also be applied to other environmental space–time processes (e.g. to model time series of different ambient primary or secondary pollutants) and used to predict non-daily data (e.g. hourly).

CONFLICT OF INTEREST

The authors declare no conflict of interest.

ACKNOWLEDGEMENTS

This work was partially supported by (i) the Traffic Pollution and Health in London Grant (NE/I008039/1) from the National Environment Research Council, Medical Research Council, Economic and Social Research Council, Department of Environment, Food and Rural Affairs and Department of Health and (ii) the ALSPAC Grant (P13646_DFHM) from the Medical Research Council.

REFERENCES

- Briggs DJ, de Hoogh C, Gulliver J, Wills J, Elliott P, Kingham S et al. A regression-based method for mapping traffic related air pollution: application and testing in four contrasting urban environments. *Sci Total Environ* 2000; **253**: 151–167.
- Hoek G, Beelen R, de Hoogh K, Vienneau D, Gulliver J, Fischer P et al. A review of land-use regression models to assess spatial variation of outdoor air pollution. *Atmos Environ* 2008; **42**: 7561–7578.
- Gulliver J, Morris C, Lee K, Vienneau D, Briggs D, Hansell A. Land use regression modeling to estimate historic (1962–1991) concentrations of black smoke and sulfur dioxide for Great Britain. *Environ Sci Technol* 2011; **45**: 3526–3532.
- Shaddick G, Wakefield J. Modelling daily multivariate pollutant data at multiple sites. *J R Stat Soc Ser C* 2002; **51**: 351–372.
- Huerta G, Sanso B, Stroud JR. A spatiotemporal model for Mexico City ozone levels. *J R Stat Soc Ser C* 2004; **53**: 231–248.
- Sahu SK, Gelfand AE, Holland DM. Spatio-temporal modeling of fine particulate matter. *J Agric Biol Environ Stat* 2006; **11**: 61–86.
- Sahu SK, Gelfand AE, Holland DM. High resolution space–time ozone modeling for assessing trends. *J Am Stat Assoc* 2007; **102**: 1221–1234.
- Cocchi D, Greco F, Trivisano C. Hierarchical space–time modelling of PM10 pollution. *Atmos Environ* 2007; **41**: 532–542.
- Chang HH, Peng RD, Dominici F. Estimating the acute health effects of coarse particulate matter accounting for exposure measurement error. *Biostatistics* 2011; **12**: 637–652.
- Cameletti M, Ignaccolo R, Bande S. Comparing spatio-temporal models for particulate matter in Piemonte. *Environmetrics* 2011; **22**: 985–996.
- National Research Council. *Models in Environmental Regulatory Decision Making*. The National Academies Press: Washington, DC, USA, 2007. Available at www.nap.edu/catalog/11972.html.
- Fuentes M, Raftery AE. Model evaluation and spatial interpolation by Bayesian combination of observations with outputs from numerical models. *Biometrics* 2005; **61**: 36–45.
- McMillan N, Holland DM, Morara M, Feng J. Combining numerical model output and particulate data using Bayesian space–time modeling. *Environmetrics* 2010; **21**: 48–65.
- Sahu SK, Yip S, Holland DM. Improved space–time forecasting of next day ozone concentrations in the Eastern US. *Atmos Environ* 2009; **43**: 494–501.
- Sahu SK, Gelfand AE, Holland DM. Fusing point and areal level space–time data with application to wet deposition. *J R Stat Soc Ser C* 2010; **59**: 77–103.
- Berrocal VJ, Gelfand AE, Holland DM. A spatio-temporal downscaler for output from numerical models. *J Agric Biol Environ Stat* 2010; **15**: 176–197.
- Berrocal VJ, Gelfand AE, Holland DM. A bivariate space–time downscaler under space and time misalignment. *Ann Appl Stat* 2010; **4**: 1942–1975.
- Gelfand AE, Sahu SK. Combining monitoring data and computer model output in assessing environmental exposure. In: O'Hagan A, West M (eds). *Handbook of Applied Bayesian Analysis*. Oxford University Press: Oxford, UK, 2010, pp 482–510.
- Lenschow P, Abraham HJ, Kutzner K, Lutz M, Preuß JD, Reichenbacher W. Some ideas about the sources of PM10. *Atmos Environ* 2001; **35**: S23–S33.
- Samet JM, Dominici F, Currier FC, Coursac I, Zeger SL. Fine particulate air pollution and mortality in 20 US cities, 1987–1994. *N Engl J Med* 2000; **343**: 1742–1749.
- Atkinson R, Anderson H, Sunyer J, Ayres J, Baccini M, Vonk J et al. Acute effects of particulate air pollution on respiratory admissions: results from APHEA 2 project. *Am J Respir Crit Care Med* 2001; **164**: 1860–1866.
- Analitis A, Katsouyanni K, Dimakopoulou K, Samoli E, Nikoloulopoulos AK, Peta-sakis Y et al. Short-term effects of ambient particles on cardiovascular and respiratory mortality. *Epidemiology* 2006; **17**: 230–233.
- Solberg S, Coddeville P, Forster C, Hov Ø, Orsolini Y, Uhse K. European surface ozone in the extreme summer 2003. *Atmos Chem Phys* 2005; **5**: 9003–9038.
- Allen G, Reiss R. Evaluation of the TEOM method for measurement of ambient particulate mass in urban areas. *J Air Waste Manage* 1997; **47**: 682–689.
- Green D, Fuller G, Barratt B. Evaluation of TEOMTM 'correction factors' for assessing the EU stage 1 limit values for PM10. *Atmos Environ* 2001; **35**: 2589–2593.
- DETR. *Assistance with the Review and Assessment of PM10 Concentration in Relation to the Proposed EU Stage 1 limit values*. HMSO: London, UK, 1999.
- Harrison DUK. *Equivalence Programme for Monitoring of Particulate Matter*. Ref: BV/AQ/AD202209/DH/2396DEFRA: London, UK, 2006.
- McHugh CA, Carruthers DJ, Edmunds HA. ADMS and ADMS-Urban. *Int J Environ Pollut* 1997; **8**: 438–440.
- Carruthers DJ, Edmunds HA, Lester AE, McHugh CA, Singles RJ. Use and validation of ADMS-Urban in contrasting urban and industrial locations. *Int J Environ Pollut* 2000; **14**: 364–374.
- Peng RD. A method for visualizing multivariate time series. *J Stat Softw* 2008; **25**: 1–17.
- Fuller G, Green D. Evidence for increasing concentrations of primary PM10 in London. *Atmos Environ* 2006; **40**: 6134–6145.
- Fuller G. *Air Quality in London 2003. Final Report*. King's College London, Environmental Research Group: London, UK, 2005. Available at: www.londonair.org.uk/london/reports/AirQualityInLondon2003.pdf (last accessed date 25 June 2013).
- Gelfand AE, Kim H-J, Sirmans CF, Banerjee S. Spatial modeling with spatially varying coefficient processes. *J Am Stat Assoc* 2003; **98**: 387–396.
- Banerjee S, Gelfand AE, Carlin BP. *Hierarchical Modeling and Analysis for Spatial Data*. Chapman & Hall/CRC: Boca Raton, FL, 2004.
- Lee D, Shaddick G. Modelling the effects of air pollution on health using Bayesian dynamic generalised linear models. *Environmetrics* 2008; **19**: 785–804.
- Lindgren F, Rue H. On the second-order random walk model for irregular locations. *Scand J Stat* 2008; **35**: 691–700.
- Chiogna M, Gaetan C. Dynamic generalized linear models with applications to environmental epidemiology. *J R Stat Soc Ser C* 2002; **51**: 453–468.
- Diggle P, Ribeiro P. *Model-Based Geostatistics*. Springer: New York, NY, USA, 2007.
- Lunn DJ, Thomas A, Best N, Spiegelhalter D. WinBUGS — a Bayesian modelling framework: concepts, structure, and extensibility. *Stat Comput* 2000; **10**: 325–337.
- Gilks WR, Richardson S, Spiegelhalter DJ. *Markov Chain Monte Carlo in Practice*. Chapman & Hall: Boca Raton, FL, USA, 1998.

- 41 Hox J. *Multilevel Analysis. Techniques and Applications*. Lawrence Erlbaum Associates: Englewood Cliffs, NJ, 2002.
- 42 City Mayors Statistics. Europe's largest cities — Cities ranked 1 to 100. City Mayors 2012, Available at: http://www.citymayors.com/features/euro_cities1.html (last accessed date 2 August 2013).
- 43 Beauchamp JJ, Olson JS. Corrections for bias in regression estimates after logarithmic transformation. *Ecology* 1973; **54**: 1403–1407.
- 44 Taylor K. Summarizing multiple aspects of model performance in a single diagram. *J Geophys Res* 2001; **106**: 7183–7192.
- 45 Carslaw DC, Ropkins K. Openair — an R package for air quality data analysis. *Environ Model Softw* 2012; **27–28**: 52–61.
- 46 Laurent O, Pedrono G, Segala C, Filleul L, Havard S, Deguen S *et al*. Air pollution, asthma attacks, and socioeconomic deprivation: a small-area case-crossover study. *Am J Epidemiol* 2008; **168**: 58–65.
- 47 Blangiardo M, Richardson S, Gulliver J, Hansell A. A Bayesian analysis of the impact of air pollution episodes on cardio-respiratory hospital admissions in the Greater London area. *Stat Methods Med Res* 2011; **20**: 69–80.



This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs 3.0 Unported License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/3.0/>