



HHS Public Access

Author manuscript

Ann Work Expo Health. Author manuscript; available in PMC 2018 August 01.

Published in final edited form as:

Ann Work Expo Health. 2017 August 01; 61(7): 773–783. doi:10.1093/annweh/wxx046.

A Bayesian Approach for Summarizing and Modeling Time-Series Exposure Data with Left Censoring

E. Andres Houseman^{1,*} and M. Abbas Virji²

¹Oregon State University, College of Public Health and Human Sciences, 101 Milam Hall, 2520 SW Campus Way, Corvallis, OR 97331, USA

²National Institute for Occupational Safety and Health, Respiratory Health Division, 1095 Willowdale Rd, Morgantown, WV 26505, USA

Abstract

Objective—Direct reading instruments are valuable tools for measuring exposure as they provide real-time measurements for rapid decision making. However, their use is limited to general survey applications in part due to issues related to their performance. Moreover, statistical analysis of real-time data is complicated by autocorrelation among successive measurements, non-stationary time series, and the presence of left-censoring due to limit-of-detection (LOD). A Bayesian framework is proposed that accounts for non-stationary autocorrelation and LOD issues in exposure time-series data in order to model workplace factors that affect exposure and estimate summary statistics for tasks or other covariates of interest.

Method—A spline-based approach is used to model non-stationary autocorrelation with relatively few assumptions about autocorrelation structure. Left-censoring is addressed by integrating over the left tail of the distribution. The model is fit using Markov-Chain Monte Carlo within a Bayesian paradigm. The method can flexibly account for hierarchical relationships, random effects and fixed effects of covariates. The method is implemented using the *rjags* package in R, and is illustrated by applying it to real-time exposure data. Estimates for task means and covariates from the Bayesian model are compared to those from conventional frequentist models including linear regression, mixed-effects, and time-series models with different autocorrelation structures. Simulations studies are also conducted to evaluate method performance.

Results—Simulation studies with percent of measurements below the LOD ranging from 0 to 50% showed lowest root mean squared errors for task means and the least biased standard deviations from the Bayesian model compared to the frequentist models across all levels of LOD. In the application, task means from the Bayesian model were similar to means from the frequentist

* Author to whom correspondence should be addressed. Tel: +1 541 737 2643; Fax: +1 541 737 6914; andres.houseman@oregonstate.edu.

Supplementary data

Supplementary data are available at *Annals of Work Exposures and Health* online.

Disclaimer

Mention of a specific product or company does not constitute endorsement by the National Institute for Occupational Safety and Health. The findings and conclusions in this report are those of the authors and do not necessarily represent the views of the National Institute for Occupational Safety and Health.

The authors declare no conflict of interest relating to the material presented in this article.

models, while the standard deviations were different. Parameter estimates for covariates were significant in some frequentist models, but in the Bayesian model their credible intervals contained zero; such discrepancies were observed in multiple datasets. Variance components from the Bayesian model reflected substantial autocorrelation, consistent with the frequentist models, except for the auto-regressive moving average model. Plots of means from the Bayesian model showed good fit to the observed data.

Conclusion—The proposed Bayesian model provides an approach for modeling non-stationary autocorrelation in a hierarchical modeling framework to estimate task means, standard deviations, quantiles, and parameter estimates for covariates that are less biased and have better performance characteristics than some of the contemporary methods.

Keywords

autocorrelation; Bayesian; limit of detection; non-stationary time series; spline model

Introduction

Direct reading real-time instruments are valuable tools for measuring exposure, and have traditionally been utilized as early warning devices to detect leaks, evaluate accidental exposures, or for emergency responses (Pearce and Coffey, 2011). They are appealing as they provide real-time or near-real time measurements eliminating the time-lag between sample collection and laboratory analysis, thus enabling rapid decision making. Real-time data also provide important information on the short-term variability of exposure within a work-shift which may be important for identification of exposure excursions, development of control strategies, and metrics of peak exposures for epidemiologic studies. Generally, the use of direct reading instruments for aerosols, gases, and vapors have been limited to general survey or screening applications in part due to their lack of specificity or due to issues related to their performance such as validity, precision, calibration etc. (Coffey and Pearce, 2010). Presently, direct reading aerosol instruments are increasingly used to assess exposures to engineered nanomaterials, primarily as screening tools (Kuhlbusch *et al.*, 2011; Ostraat *et al.*, 2013). With the recent advances in sensor technology, especially nanotechnology-enabled sensors, there is renewed and growing interest in the development of direct reading instruments that have improved sensitivity, detection limit, specificity, multiplexing capability, and other performance characteristics (Sadik *et al.*, 2009; van Zee *et al.*, 2009). Direct reading instruments hold tremendous promise of novel exposure metrics for use in epidemiologic studies of acute effects or sensitization where peak exposures and patterns of exposure are important (Mihlan *et al.*, 2000; Preller *et al.*, 2004), as well as of identifying important time-varying factors such as tasks that affect exposure levels, thus offering opportunities for targeted control measures. To emphasize the importance of developing and using direct reading real-time instrumentation in exposure assessment, the National Institute for Occupational Safety and Health (NIOSH) launched the direct reading exposure assessment methods (DREAM) initiative in 2008 with the goal of enhancing the use of these technologies to improve worker health and safety (NIOSH, 2014). As part of this effort, NIOSH established the Center for Direct Reading and Sensor Technologies (NCDRST) in 2014 to coordinate research and develop guidance on direct reading and

sensor technologies. Despite their historical use and the renewed interest in direct reading instruments for exposure assessment, there is a paucity of published literature on the use of these instruments for quantitative occupational exposure assessment and epidemiology.

In the practice of industrial hygiene, measurement data are used for multiple purposes including estimating exposures for jobs or tasks for evaluating compliance with regulatory standards, for epidemiologic studies, or for assessing engineering controls or respiratory protection requirements. Real-time measurement with observations or self-reported activity diaries can be particularly useful for identifying high-exposure tasks to target interventions, however, reliably estimated summary measures such as the mean, geometric mean (GM), geometric standard deviation (GSD), and various quantiles such as the 95th percentile (P95) are needed to fully utilize these data. Modeling the determinants of real-time exposure measurements can provide task-specific exposure estimates taking into account fixed-effects covariates that may explain exposure variation within tasks, e.g. enclosures, engineering controls or materials used, and random effects of clustering around measurements taken on an individual, or within a specific workplace. In addition to the task-specific mean or GM, estimates of task-specific variation (e.g. GSD) and quantiles (e.g. P95) are essential to inform the appropriate selection of control measures (e.g. ventilation equipment or respirator selection) or in assigning short-duration task exposures to participants in epidemiologic studies, e.g. for exposure to cleaning chemicals among healthcare workers. However, statistical analysis of real-time data is made difficult by several issues in addition to the analytical issues associated with the performance of direct reading instruments such as lack of specificity or the presence of interferences. One is the high potential for non-stationary autocorrelation among successive measurements, especially when accounting for a variable microenvironment (i.e. task). Another is the presence of left-censoring due to limit-of-detection (LOD). Thus there is a need for a statistical tool for modeling the determinants of real-time exposure data that can account for non-stationary autocorrelation, LOD measurements, fixed-effects covariates, hierarchical random effects, and provide a range of summary measures such as the mean, standard deviation, and various quantiles of interest.

A variety of approaches have been used to analyze autocorrelated exposure data in the occupational and environmental fields ranging from relatively simple approaches such as graphical presentation, time-weighted averages (Dodson *et al.*, 2007), regression approaches, and some nonparametric tests (Brook *et al.*, 2007), and time-series models (Oka *et al.*, 2010; Entink *et al.*, 2011) to more complex approaches recently proposed using Bayesian methods. For example, Oka *et al.* (2010) proposed an algorithm for estimating volatile organic compound (VOC) concentrations from time-series data using simple moving averages. Entink *et al.* (2011) proposed an ARIMA time-series approach using autoregressive correlation structure to analyze time-series data. More recently, Entink *et al.* (2015) proposed a Bayesian approach for analyzing nanoparticle data, fitting a multilevel model with first-order autoregressive error structure to address temporal autocorrelation. Clerc *et al.* (2013) proposed a Bayesian approach for analyzing nanoparticles, where the probability distribution of differences between source and far-field sampling points were estimated via Bayesian methods, but the approach does not provide data-generating statistical models.

In general, time series obtained from personal real-time monitors can be thought of as either functional data (Ramsay and Silverman, 2002) or repeated measures data (Diggle *et al.*, 2002). The former approach, which assumes a uniform time domain over which the characteristics of a stochastic process are to be estimated, is complicated by the irregular lengths of the series over which the personal real-time data have been collected. The latter approach, which attempts to model autocorrelation within each series but does not require time domains of uniform length, is a more natural choice. However, most available software provide a limited number of choices in modeling the correlation within a series: compound symmetry (every measurement is equally correlated with every other measurement within a series); first-order autoregressive models (AR-1), which assume a correlation that exponentially decays with a larger time interval and is characterized by a single correlation parameter, e.g. AR-1 models, which assume equally spaced intervals, or conditional autoregressive (CAR) models which account for irregularity in interval spacing; or more general autoregressive-moving average (ARMA) models (Entink *et al.*, 2011). These options are typically available in software for estimating linear mixed effects (LME) models, e.g. the `lme` function in R (`nlme` package) or the MIXED procedure in SAS; however, they all assume stationary time-series (i.e. time-series where the autocovariation parameters are constant throughout the series). Alternatives based on estimating marginal means (i.e. generalized estimating equations, or GEEs, models that do not require detailed specification of within-cluster correlation structure) are also available; these approaches employ simple yet unbiased models for estimating the fixed effects of interest but provide ‘robust’ standard error estimates that account for autocorrelation within a series. These methods tend to underperform when there are a small number of long series (Houseman *et al.*, 2002); limitations can be mitigated by time-series bootstrap methods (Heagerty and Lumley, 2000), which partition a series into ‘windows’ that are resampled for standard error estimation, although it can be difficult to choose a window size and there are no standard software packages for implementation.

None of the methods described above simultaneously account for left-censoring in a rigorous manner, in every case requiring the ad-hoc and arbitrary substitution of some fraction of the LOD for left-censored values. Such substitution can lead to bias in statistical estimates (Antweiler and Taylor, 2008). It is possible to use multiple imputation to impute the left-censored values (Hopke *et al.*, 2001), or else via quadrature methods (Thiébaud and Jacqmin-Gadda, 2004). These latter methods are relatively computationally intensive as they depend on resampling of one form or another.

In this article, we present a systematic Bayesian framework for analyzing exposure time-series arising from real-time monitors. The proposed model, illustrated in Figure 1, has two components: (1) fixed-effects terms representing variables of interest (e.g. task indicators, microenvironment characteristics, or instrument indicators); and (2) a non-parametric term representing autocorrelated error processes. Because the model is constructed within a Bayesian framework, it can easily integrate over the censored part of the modeled distributions and can easily be extended to include hierarchical random effects such as worker nested within a workplace, nested within an industry. Our method is similar to that proposed by Entink *et al.* (2015), but we employ a spline-based approach to model autocorrelation, which allows for potential non-stationary autocorrelation and entails fewer

assumptions about the autocorrelation structure; we also address left-censoring due to LODs, and provide flexibility in accounting for potentially complex hierarchical relationships, fixed effects of exposure determinants and other covariates of interest. Although the computation is of complexity similar to those based on multiple imputation or quadrature, we have implemented it using the *rjags* package in R; this package is an R interface for the software *Just Another Gibbs Sampler (JAGS)*, which provides comprehensive support for sampling from the posterior distribution of a Bayesian model, allowing for flexible modification of the model to fit specific exposure assessment needs, as well as flexible data manipulation via the functionality available in R.

We address the non-stationary autocorrelation using a spline approach that is common semi-parametric modeling (Ruppert *et al.*, 2003). Splines have been used extensively in the environmental statistics literature to account for autocorrelation (Paciorek and Schervish, 2006; Gryparis *et al.*, 2007; Torabi and Rosychuk, 2011). We apply the basic idea here to account for autocorrelation within the measurement time-series. One important issue is the selection of tuning parameters to properly constrain the spline coefficients, and to that end we employ a method proposed by Brumback *et al.* (1999), where the spline coefficients are embedded in a random effects model and the corresponding tuning parameter is represented by a variance component parameter.

In the following sections, we present a brief description of the approach, followed by an illustration of the method by applying it to real-time data, and ending with simulation studies to evaluate method performance.

Methods

Details of our proposed model, depicted schematically in Figure 1, appear in the Supplementary Methods section available online (Supplementary Online Material Part 1); we briefly describe its principal features here.

Bayesian spline model

We assume that n time-series Y_i , $i \in \{1, \dots, n\}$, have been collected, each series of potentially varying length, and with each measurement corresponding to one of m designated occupational tasks. We also allow for the possibility that each measurement corresponds to a vector of covariates. Each measured value Y_{ir} (where r indexes individual sequential measurements) is assumed to be normally distributed, with mean μ_{ir} depending upon a number of fixed effects (task-specific means and coefficients corresponding to covariates), task-specific measurement errors, and random series-specific effects. The latter random series-specific effects are represented as a stochastic error process, implemented using a B-splines basis, and may also include a random series-specific intercept to account for heterogeneity in mean response across individual profiles. Additionally, we support task-specific standard deviations. Finally, a variant of our model accounts for left censoring by LOD, which is a common feature in many exposure assessment settings.

While our proposed model is complicated to fit within a frequentist framework, it is relatively straight-forward to fit this model in a Bayesian paradigm using standard software

such as JAGS. JAGS compiles a model described using a standardized language that supports flexible Bayesian model specification, together with data informing the model and initial guesses at values representative of the posterior distribution, and returns a Markov chain representing values sampled from the posterior distribution. The fundamental principle used by JAGS is Gibbs Sampling, a well-known Markov-Chain-Monte-Carlo (MCMC) technique. In particular, Gibbs Sampling seamlessly integrates over unobserved portions of the model, e.g. hierarchical random effects and the left tail of below-detection observations.

Data source

We demonstrate our proposed analytical approach within the context of two exposure assessment scenarios. In our primary example, we used a dataset of real-time exposure to total volatile organic compounds (TVOC) among 141 workers over 207 days covering 14 healthcare occupations at several hospitals described in detail by LeBouf *et al.* (2014). Full-shift monitoring using seven real-time TVOC monitors (ppbRAE Plus monitor, RAE Systems, San Jose, CA, USA) as mobile-area samples was conducted to measure ppb-level TVOC concentrations. During exposure monitoring, information on exposure determinants such as tasks and activities, products used, tools used, control technologies etc., was recorded at 5-min intervals on standardized form for each study participant for the entire sampling duration as described by Saito *et al.* (2015). A wide range of tasks are performed by healthcare occupations from cleaning and laboratory tasks to patient care and administrative tasks. The averaging time of the sampling instruments was set to 10 s, which was summarized to 5-min averages to match the observation data. The instrument LOD was 1 ppb; 8.3% of the 10-s readings were below the LOD, resulting in 7.9% of the 5-min averages falling below detection. In this example, our approach illustrates the summarization of real-time exposure data accounting for autocorrelation within series defined by employee–date combinations, as well as non-stationary data, and LOD measurements. In particular, we present summary exposures for tasks, accounting for fixed effects of covariates such as instruments.

We also used a second, smaller dataset of real-time exposure to nanoparticles collected during a walkthrough visit at an ultrafine titanium dioxide (TiO₂) and lithium titanate (Li₂TiO₃ or LTO) manufacturing facility. The purpose of this analysis was to inform the parameters of a simulation study, described in the latter portion of this article. Summary statistics are presented in Table 1. Details of the analysis are described in the Supplementary Methods section (Supplementary Online Material Part 1).

Results

Examples

We now demonstrate the proposed methods with an analysis of real-time TVOC monitoring data from healthcare workers. Summary characteristics of this data set appear in Table 1. Many of the 5-min average TVOC concentrations fell below the detection limit of $\delta = 1$ ppb and the number of series present in the data set was large enough to warrant the inclusion of random series intercepts. Thus, this analysis demonstrates the full potential of our proposed methodology. The model was fit using JAGS implemented in the R package *rjags* (version

3.4.0) run in R (version 3.2.2). The JAGS model used is described in the Supplementary Methods under *Code for Example 1* (in Supplementary Online Material part 1).

Posterior statistics are shown in Table 2 for a selection of tasks related to cleaning and disinfecting activities. Figure 2 shows the observed data for a time single series Y_j (an LPN sampled on August 13, 2010) overlaid with the mean profile μ_{jt} . Essentially, the solid curve represents the realization of a smooth stochastic process modeled by the spline term, while differences between the dashed and solid curves represent the independent error process. Supplementary Figure S1 (in Supplementary Online Material Part 2) shows plots of observed measurements Y_{jt} by expected value μ_{jt} along with a 95% credible intervals for μ_{jt} ; in general, there was substantial agreement between observed and expected, although some of the below-detection values were predicted to be very close to zero.

Corresponding estimates from frequentist methods are shown in Supplementary Table S1 (in Supplementary Online Material Part 2) for all the tasks evaluated. We note that estimates from frequentist methods were similar to, but not identical with, the posterior means and medians from the proposed method. Supplementary Table S2 (in Supplementary Online Material Part 2) provides Pearson and Spearman correlation coefficients between frequentist estimates of task-specific means and the corresponding posterior means from our proposed model. No method returns results that are perfectly correlated with another. In general, posterior means from our method, AR-1 estimates, and CAR estimates were correlated with each other to a moderately high degree (Spearman 0.66–0.75, Pearson 0.73–0.91) and uncorrelated with OLS estimates. The OLS model assumes independent measurements and additionally, LOD measurements were replaced by LOD/2 in this model. Under these circumstances, we would not expect OLS to perform well. The lack of correlation reinforces the observation that OLS is inferior to our proposed method as well as some of the other commonly used frequentist methods we investigated.

We also fit a model that adjusts for instrument-specific effects, as described in the Supplementary Methods (in Supplementary Online Material part 1). Posterior statistics are shown in Supplementary Table S3 (in Supplementary Online Material part 2), and reveal task-specific effects very similar to those shown in Table 1. Note that the random effect SD approximately doubles after adding an instrument effect, reflecting a greater differentiation of individual profiles after deconvolving the instrument effect (which is otherwise mixed into the profile effects). As a final note, we assessed convergence of the MCMC chains by applying the Gelman and Rubin convergence diagnostic (available in the R package *coda*) to two independent chains for each data analysis performed. The convergence statistic, \hat{R} , should be close to one, preferably less than 1.2 or (a more stringent criterion) less than 1.1. As it was not possible to use all variables in the multivariate chain to calculate \hat{R} , we assessed convergence only for the α variables, i.e., those governing the task-specific means. We calculated $\hat{R} = 1.01$ for our primary example.

Detailed results of our second analysis appear in Supplementary Tables S4–S6 and Supplementary Figure S3–S6 (in Supplementary Online Material part 2). In summary, we observed many of the same features described for the primary data analysis, including the fact that frequentist estimates were similar to, but not identical with, estimates from our

proposed model. In addition, parameter estimates for covariates, e.g. source enclosure were significant in some frequentist models, but in the Bayesian model their credible intervals contained zero; such discrepancies with the covariates were observed in multiple datasets. Note that in the Bayesian context, the credible intervals describe the bounds around the parameter estimate, and the comparison to null is not relevant; we do that here solely to facilitate comparison to the classical models.

Simulations

Our proposed method is cast within a Bayesian framework, and thus, from a strict philosophical perspective, does not admit frequentist interpretations. Nevertheless, we anticipate readers will be interested in comparisons of our method with commonly used frequentist methods. Consequently, we conducted several simulation experiments to investigate the properties of the proposed method when applied with a frequentist interpretation, using posterior mean as an analogue of the frequentist estimate, posterior standard deviation as an analogue of the frequentist standard error, and credible interval as an analogue of the frequentist confidence interval. For each experiment, we used posterior statistics obtained from one of the nanoparticle data sets analyzed in our second example. Details are described in the Supplementary Methods section (in Supplementary Online Material part 1). We compared our proposed method to three commonly used frequentist methods: OLS, CAR, and ARMA, each with the value of below-detection data set to half the LOD (on the original scale before applying the logarithmic transform). In general, our proposed method appeared to be more efficient than the competing methods. For example, our method generated the lowest values of root mean square error (RMSE), as shown in Figure 3, which, for a range of assumed detection limits, displays the RMSE values for the intercept and for one of the five tasks assumed in the simulation. Results for all tasks are shown in Supplementary Figure S8 (in supplementary online material Part 2). Additionally, our method often demonstrated more accurate coverage, and frequentist methods showed biased estimates of sampling standard deviations. We also examined the robustness of our method when the true data-generating mechanism followed a more conventional error model. We fit four AR(1) models, with autocorrelation parameters 0.10, 0.25, 0.50, and 0.75 respectively, and fit three AR(2) models, with autocorrelation vectors (0.25, -0.50), (0.50, -0.25), and (0.50, -0.50). Figure 4 displays the results for the AR(1) simulations; AR(2) results appear in Supplementary Figure S12, in Supplementary Online Material part 2. RMSE values were about the same for all methods (except OLS in some cases); in terms of interval coverage our proposed method performed about as well as CAR and ARIMA in the AR(2) models and in the AR(1) models with lower autocorrelation, but tended to break down at the highest level of autocorrelation (0.75). We surmise that the reason for this is that the dense placement of knots was inconsistent with the high levels of autocorrelation, leading to an insufficiently high value for the maximum possible autocorrelation that could be modeled; the remedy, in such a situation, would be to use a sparser set of knots.

Discussion and conclusions

There is a growing interest in developing and using direct reading instruments to assess exposures as they minimize analytical cost and the eliminate the time-lag between sampling

and receiving results after analysis. For example, the 2014 Mine Safety and Health Administration's Respirable Coal Dust Rule (30 CFR Parts 70, 71, 74, and 90) (USDOL 2014) requires the use of continuous personal dust monitors (CPDM) to assess respirable dust exposures of underground coal miners, and recommends their use (optional) for other circumstances such as surface miners, workers in non-production areas of the underground coal mining operations, or underground anthracite mines using certain methods, to enable rapid identification of out-of-control situation and enable corrective actions to be taken in a timely manner. Thus it is foreseeable that large quantities of real-time data will be generated in the near future. Without readily available analytical methods, it is likely that all the real-time data will be summarized into a single time-weighted average (TWA) summary as has been traditionally done with real-time exposure data, e.g. noise data. Once the summary data have been stored in databases, the real-time data may no longer be available to examine and extract information on short-term exposure variability and peak exposures. Systematic evaluation of such data will likely provide extremely valuable information to plan for and develop effective strategies to prevent overexposures to augment the on-the-spot corrective actions. Thus, there is a need to develop readily available methods and tools to analyze real-time exposure data.

Another rapidly growing area that will require these new statistical analytical tools is related to the field of sensor development (Sadik *et al.*, 2009; van Zee *et al.*, 2009). With new multiplexing platforms, large datasets will be generated which will likely have left-censoring issues for some of the analytes because of the different proportions of chemicals present in a mixture. These types of sensors are already being used in a variety of exposure and health assessment settings, e.g. the use of electronic noses where an array of sensors are employed to detect patterns of mixtures to distinguish signature patterns of different products, health outcomes, or signatures of different exposure scenarios. There is already interest in extending the capabilities of these sensor arrays beyond pattern recognition to quantitative determination of the different exposure components in real time. This calls for an extension of the analytical methods to develop multivariate methods that address the issue of correlation among multiple outcome variables in addition to all the issues identified with analysis of real-time data. Entink and colleagues have started to address this aspect of multiple correlated outcome variables datasets containing size distribution of particles in real-time (Entink *et al.*, 2015), but did not incorporate issues related to left censoring and non-stationary autocorrelation. Using additional random effects, our method can be extended in a relatively straightforward manner to incorporate correlation across different measurements collected simultaneously.

With regards to the simulation studies, we note that all methods tended to break down with high autocorrelation. In cases where high autocorrelation is anticipated (or judged on the basis of simpler methods such as ARMA), potential remedies are longer averaging intervals (e.g. averaging over 60 min rather than 5 min) or else using sparser knot placement in our proposed model. Additional work is needed in estimating optimal knot selection in this context.

We also note that our method entails distribution assumptions that may not be exactly correct, e.g. lognormality of the response. Note, however, that non-normality of posterior

distributions may be driven by the skewed nature of the variance component parameters. Note also that the available methods commonly employed for real-time data entail the same distribution assumptions, with little flexibility for altering them without extensive reworking of algorithms. In contrast, our Bayesian formulation can relatively flexibly be altered to account for different distributions (e.g. gamma responses), although with some considerable computational cost. In general, our method provides a flexible methodology for addressing non-stationary stochastic error in real-time sampling. Our model accounts for potential task-specific variation in error, at the cost of slightly complicating the interpretation of variance components; however, as we have demonstrated in the results presented in Supplementary Material, it is still possible to use graphical means to communicate the time-dependent autocovariance implied by the posterior distributions of the parameters (see Supplementary Figures S2 and S3 in the Supplementary Online Material part 2). We note that additional heterogeneity in the variance of the stochastic error can be addressed using additional layers of variation, as can more complex hierarchical relationships among the employees. Given our Bayesian framework and use of the R library *rjags*, it is relatively straightforward to implement such extensions.

In conclusion, the proposed Bayesian model provides an approach for modeling non-stationary autocorrelation in a hierarchical modeling framework to estimate task means, standard deviations, and parameter estimates for covariates that are less biased and have better performance characteristics than some of the contemporary methods.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

The authors thank Drs. Scott Hendricks and James T. Wassell for their technical review of the manuscript, Dr. Ethan Fechter-Leggett for assistance with running the R software, Ms. Xiaoming Liang for preparation of the datasets and checking the R-code, (available at <https://www.cdc.gov/niosh/topics/drst/>) and Dr. Aleksandr Stefaniak for making available the nanoparticle exposure data used in example 2.

References

1. Antweiler RC, Taylor HE. Evaluation of statistical treatments of left-censored environmental data using coincident uncensored data sets: I. Summary statistics. *Environ Sci Technol.* 2008; 42:3732–8. [PubMed: 18546715]
2. Brook JR, Burnett RT, Dann TF, et al. Further interpretation of the acute effect of nitrogen dioxide observed in Canadian time-series studies. *J Expo Sci Environ Epidemiol.* 2007; 17(Suppl. 2):S36–44. [PubMed: 18079763]
3. Brumback BA, Ruppert D, Wand MP. Comment. *J Am Stat Assoc.* 1999; 94:794–97.
4. Clerc F, Njiki-Menga G-H, Witschger O. Exploratory study on a statistical method to analyse time resolved data obtained during nanomaterial exposure measurements. *J Phys Conference Series IOP Publishing.* 2013; 429:012003.
5. Coffey CC, Pearce TA. Direct-reading methods for workplace air monitoring. *J Chem Health Safety.* 2010; 17:10–21.
6. Diggle, P., Heagerty, P., Liang, K-Y., et al. *Analysis of longitudinal data.* Oxford, UK: Oxford University Press; 2002.

7. Dodson RE, Houseman EA, Levy JI, et al. Measured and modeled personal exposures to and risks from volatile organic compounds. *Environ Sci Technol*. 2007; 41:8498–505. [PubMed: 18200885]
8. Entink RHK, Bekker C, Fransman WF, et al. Analysis of time series of particle size distributions in nano exposure assessment. *J Aerosol Sci*. 2015; 81:62–69.
9. Entink RHK, Fransman W, Brouwer DH. How to statistically analyze nano exposure measurement results: using an ARIMA time series approach. *J Nanoparticle Res*. 2011; 13:6991–7004.
10. Gryparis A, Coull BA, Schwartz J, et al. Semiparametric latent variable regression models for spatiotemporal modelling of mobile source particles in the greater Boston area. *J Royal Stat Soc Series C (Applied Statistics)*. 2007; 56:183–209.
11. Heagerty PJ, Lumley T. Window subsampling of estimating functions with application to regression models. *J Am Stat Assoc*. 2000; 95:197–211.
12. Hopke PK, Liu C, Rubin DB. Multiple imputation for multivariate data with missing and below-threshold measurements: time-series concentrations of pollutants in the Arctic. *Biometrics*. 2001; 57:22–33. [PubMed: 11252602]
13. Houseman EA, Ryan L, Levy JI, et al. Autocorrelation in real-time continuous monitoring of microenvironments. *J Appl Stat*. 2002; 29:855–72.
14. Kuhlbusch TA, Asbach C, Fissan H, et al. Nanoparticle exposure at nanotechnology workplaces: a review. *Part Fibre Toxicol*. 2011; 8:22. [PubMed: 21794132]
15. LeBouf RF, Virji MA, Saito R, et al. Exposure to volatile organic compounds in healthcare settings. *Occup Environ Med*. 2014; 71:642–50. [PubMed: 25011549]
16. Mihlan GJ, Todd LA, Truong KN. Assessment of occupational exposure patterns by frequency-domain analysis of time series data. *Appl Occup Environ Hyg*. 2000; 15:120–30. [PubMed: 10712067]
17. NIOSH. [Accessed 12 February 2015] Direct reading and sensor technologies center for disease control and prevention. 2014. Available at <http://www.cdc.gov/niosh/topics/drst/>
18. Oka K, Iizuka A, Inoue Y, et al. Development of a combined real time monitoring and integration analysis system for volatile organic compounds (VOCs). *Int J Environ Res Public Health*. 2010; 7:4100–10. [PubMed: 21317996]
19. Ostraat ML, Thornburg JW, Malloy QG. Measurement strategies of airborne nanomaterials. *Environ Eng Sci*. 2013; 30:126–32.
20. Paciorek CJ, Schervish MJ. Spatial modelling using a new class of nonstationary covariance functions. *Environmetrics*. 2006; 17:483–506. [PubMed: 18163157]
21. Pearce T, Coffey C. Integrating direct-reading exposure assessment methods into industrial hygiene practice. *J Occup Environ Hyg*. 2011; 8:D31–6. [PubMed: 21476167]
22. Preller L, Burstyn I, De Pater N, et al. Characteristics of peaks of inhalation exposure to organic solvents. *Ann Occup Hyg*. 2004; 48:643–52. [PubMed: 15280165]
23. Ramsay, JO., Silverman, BW. *Applied functional data analysis: methods and case studies*. New York: Springer; 2002.
24. Ruppert, D., Wand, MP., Carroll, RJ. *Semiparametric regression*. Cambridge, UK: Cambridge University Press; 2003.
25. USDOL. Lowering miners' exposure to respirable coal mine dust, including continuous personal dust monitors: final rule. 30 CFR Parts 70, 71, 74 and 90. Mine Safety and Health Administration. 2014
26. Sadik OA, Zhou AL, Kikandi S, et al. Sensors as tools for quantitation, nanotoxicity and nanomonitoring assessment of engineered nanomaterials. *J Environ Monit*. 2009; 11:1782–800. [PubMed: 19809701]
27. Saito R, Virji MA, Henneberger PK, et al. Characterization of cleaning and disinfecting tasks and product use among hospital occupations. *Am J Ind Med*. 2015; 58:101–11. [PubMed: 25351791]
28. Thiébaud R, Jacqmin-Gadda H. Mixed models for longitudinal left-censored repeated measures. *Comput Methods Programs Biomed*. 2004; 74:255–60. [PubMed: 15135576]
29. Torabi M, Rosychuk RJ. Spatio-temporal modelling using B-spline for disease mapping: analysis of childhood cancer trends. *J Appl Stat*. 2011; 38:1769–81.

30. van Zee, RD., Pomrenke, GS., Evans, HM. Nanotechnology-Enabled Sensing: Report of the National Nanotechnology Initiative Workshop. Evans, Heather M., editor. Arlington, VA: National Science and Technology Council; 2009. Available at <http://www.dtic.mil/cgi-bin/GetTRDoc?Location=U2&doc=GetTRDocpdf&AD=ADA523650> [Accessed 3 March 2017]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

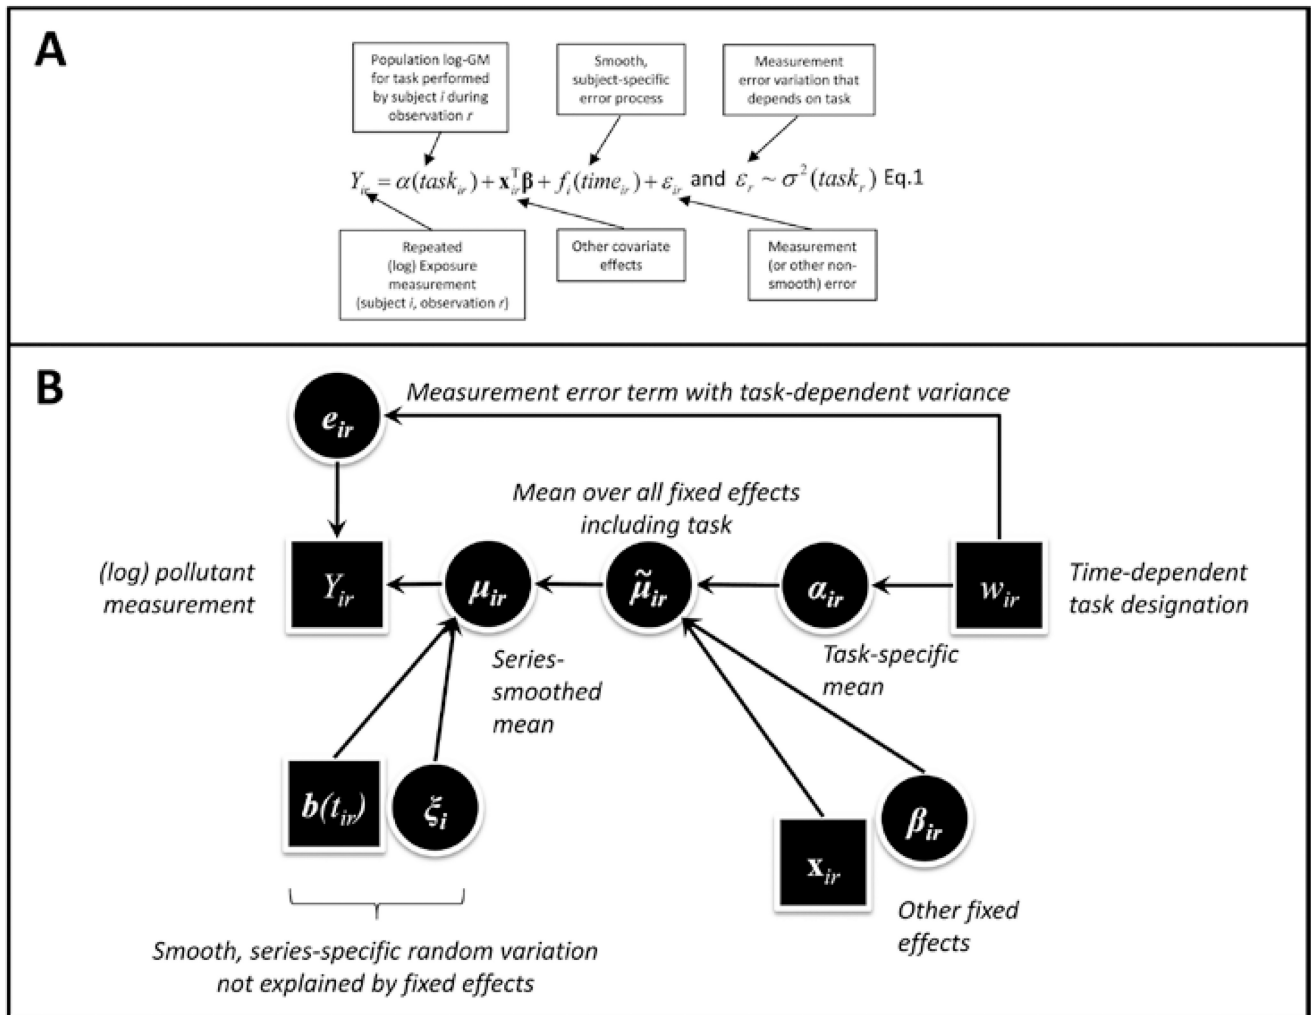


Figure 1.
Overview of the Bayesian spline model.




Figure 2. Example time series gray region shows the 2.5th and 97.5th percentile of 1000 MCMC samples from the posterior distribution of μ_{it} .




Figure 3.

Root mean square error for simulation experiment Simulation, described in detail in the Supplement, assumed five separate tasks. Shading represents detection limit used in simulation (higher detection limit corresponds to larger proportion of missing data). Coefficient from task #1 (based on Receiving Powder task from nanoparticle data set described in Supplement) displayed the largest bias resulting from missing data. Results from other tasks are shown in Supplementary Figure S8.




Figure 4. Interval coverage for simulation experiment with AR(1) errors. Y-axis shows coverage probabilities for our method ('proposed') compared with three commonly used frequentist methods. X-axis shows assumed autoregressive parameter. Individual line styles show five distinct tasks, described in detail in the Supplement.

Table 1

Characteristics of data used in examples

	Meas. (n)	ND (n)	Series (n)	GM	Median	GSD	IQR		Range		Time		
							25%ile	75%ile	Min	Max	Min	Max	
TVOC	ppbRae	16683	1313	205	124.6	192.7	8.4	58.9	434.8	<0.5	28088	0:00	23:59
Nano	CPC	143	0	3	23518	24882	1.9	17428	32174	7296	169529	8:28	16:38
	NSAM	168	0	3	49	44	1.7	33	71	20	167	8:13	15:41
	OPC	187	0	3	275	365	2.0	173	451	38	925	8:13	16:38
	SMPS	185	0	3	15141	15852	2.0	11172	21440	4158	162754	8:13	16:38

CPC, OPC, and SMPS: # cm⁻³; NSAM: μm² cm⁻³; TVOC: ppb.

Table 2

Model posterior statistics for TVOC for selected cleaning and disinfecting tasks

#	Task means (α_i)					Task SDs (σ_i)				
	Mean	Std	Med	2.5%	97.5%	Med	2.5%	97.5%	97.5%	
Buff floor	135	4.78	0.15	4.78	4.50	5.06	0.559	0.475	0.656	
Clean bathroom	187	4.77	0.14	4.77	4.51	5.05	0.338	0.283	0.405	
Cleaning	1022	4.74	0.13	4.74	4.48	5.01	0.621	0.588	0.660	
Clean spill	69	4.52	0.21	4.53	4.12	4.93	1.000	0.806	1.273	
Clean window	48	4.72	0.16	4.73	4.40	5.02	0.318	0.241	0.434	
Disinfect machine	22	5.01	0.25	5.01	4.53	5.50	0.723	0.515	1.072	
Floor strip	32	4.66	0.25	4.66	4.16	5.14	0.944	0.727	1.337	
Hand wash	38	4.76	0.14	4.75	4.49	5.04	0.084	0.049	0.200	
Mix product	33	4.77	0.17	4.78	4.43	5.11	0.349	0.214	0.578	
Mop floor	599	4.72	0.13	4.72	4.47	4.98	0.575	0.536	0.620	
Pour/mix product	77	4.79	0.14	4.79	4.52	5.08	0.317	0.249	0.407	
Scope clean	51	4.68	0.28	4.67	4.12	5.24	0.961	0.752	1.293	
Sterilize disinfect	172	4.77	0.15	4.77	4.49	5.06	0.814	0.718	0.933	
Sweep/vacuum/dust	152	4.70	0.14	4.69	4.42	4.98	0.455	0.379	0.554	
Wash equipment	220	4.70	0.15	4.70	4.42	4.99	0.285	0.249	0.327	
Wipe with alcohol	19	4.29	0.26	4.28	3.77	4.82	0.877	0.650	1.246	
Spline SD (σ_i^2)							1.68	1.65	1.72	
Random effect SD (σ_i)							0.82	0.67	1.03	

All coefficients represent statistics for natural log-transformed concentrations