A methodological framework for environmental public health surveillance with a practical example in wildfire smoke

Kathryn Teresa Morrison
Department of Epidemiology, Biostatistics and Occupational Health
McGill University, Montreal
December 2016

A thesis submitted to McGill University in partial fulfilment of the requirements for the degree of Doctor of Philosophy (Ph.D.)

Abstract

Wildfire smoke is considered a globally important cause of mortality by the World Health Organization, with an estimated 339,000 deaths from landscape fire smoke each year. The health effects from wildfire smoke are expected to intensify as changes in land use and climate are increasing the frequency and severity of wildfires, exposing more individuals to the harmful effects of smoke each year. The objective of an environmental public health surveillance (EPHS) system is to detect potential changes in the health status of the population, and to provide timely evidence for public health intervention during periods of hazardous exposure. However, the methodological and conceptual frameworks for surveillance are generally designed for infectious disease. EPHS poses some unique challenges, as exposures like air pollution are difficult to accurately measure and there are many indicators of health impact. Public health officials in environmental health have called for better decision-making surveillance tools.

To address these challenges, in Chapter 3 I proposed a methodological framework for forecasting two health indicators against a common imperfectly measured exposure. My objectives were to identify a statistical approach that would be appropriate EPHS, be flexible so that changes to data characteristics could be accommodated (e.g., additional indicators or changes in exposure metrics), and address the computational constraints of requiring an "online" daily surveillance system. I demonstrated how the proposed model could be implemented using integrated nested Laplace approximations (INLA), a cutting-edge approach to approximate Bayesian inference.

In Chapter 4, I explored the challenges of using the proposed EPHS system across an entire state or province with data aggregated by administrative boundaries, where some communities may have very small absolute populations. As an alternative to aggregating administrative units, I added spatial smoothing to the previously proposed model. The spatial smoothing stabilized the prediction variance in smaller regions, in exchange for a small loss of accuracy. In regions with larger populations, the smoothing was generally not found to be beneficial or necessary. The decision of whether to include spatial smoothing can be made within context of characteristics of the regions.

Finally, I demonstrated how the model proposed in Chapter 3 could be used for surveillance. I used the severe wildfire smoke event in July 2015 in southern British Columbia as a case study, where smoky conditions resulted in $PM_{2.5}$ concentrations up to ten times higher than baseline levels. Asthma-related physician visits and dispensations of respiratory relief medications increased by 34% and 27%, respectively, during the five smokey day period. I created a surveillance scenario by applying the comprehensive surveillance framework for wildfire smoke using nationally available smoke forecasts to predict the health impacts and estimate the reductions in morbidity that may have been achieved if such a system had been used to support intervention decisions. I found that the approach did not require precise smoke forecasts to produce reasonable health indicator forecasts, but it was negatively affected when smoke exposure models did not foresee highly smoky conditions. Based on the theoretical intervention assessment, I found that simple, early interventions with lower effectiveness, such as public health messaging about steps to reduce exposure, were preferable to delayed interventions with higher effectiveness, such as evacuation. Further research is needed on spatial interpolation and future prediction for air pollution and smoke, as well as further research on interventions to facilitate a more integrated decision making framework.

This work provides a flexible methodological approach that can be extended and improved as related research advances, such as improvements in exposure estimation and availability of additional real-time surveillance health data. The methods developed and evaluated in the first two manuscripts and implemented in the third are currently being integrated into the existing public health surveillance system at the BC Centre for Disease Control, and will be used during the wildfire season of 2017.

Résumé

La fumée des feux de forêt est considérée comme une cause de mortalité importante à travers le monde, avec 339 000 décès causés par la fumée de feux de paysage chaque année. On s'attend à ce que les effets sur la santé de la fumée de feux de forêt s'intensifient, à mesure que les changements dans l'utilisation des terres et le climat augmentent la fréquence et la gravité des feux de forêt, exposant plus d'individus aux effets nocifs de la fumée chaque année. L'objectif d'un système de surveillance de la santé publique environnementale (EPHS) est de détecter les changements potentiels dans l'état de santé de la population et de fournir des preuves en temps opportuns pour le choix d'intervention de santé publique pendant les périodes d'exposition dangereuse. Cependant, les cadres méthodologiques et conceptuels de la surveillance sont généralement conçus pour les maladies infectieuses. L'EPHS pose des défis uniques, car les expositions comme la pollution de l'air sont Mesurer avec précision et il ya de nombreux indicateurs de l'impact sur la santé. Les responsables de la santé publique en matière de santé environnementale ont réclamé de meilleurs outils de surveillance pour la prise de décisions.

Pour répondre à ces défis, nous avons proposé dans au chapitre 3 un cadre méthodologique pour la prévision de deux indicateurs de santé contre une exposition commune et imparfaitement mesurée. Nos objectifs étaient d'identifier une approche statistique qui ferait de l'EPHS appropriée, qui serait suffisamment souple pour permettre de tenir compte des changements apportés aux caractéristiques des données (par exemple, des indicateurs supplémentaires ou des changements dans les paramètres d'exposition), et qui permettrait de résoudre les contraintes informatiques liées aux besoins d'un système de surveillance quotidienne ńen ligneż. Nous avons démontré comment le modèle proposé pourrait être mis en ¡uvre en utilisant des approximations intégrées de Laplace (INLA), une approche de pointe pour l'approximation de l'inférence bayésienne. Au chapitre 4, nous avons exploré les défis liés à l'utilisation du système EPHS proposé dans l'ensemble d'un état ou d'une province, avec des données agrégées par des territoires administratifs, où certaines communautés peuvent avoir de très petites populations absolues. Comme alternative à l'agrégation des territoires administratifs, nous avons ajouté le lissage spatial au modèle proposé précédemment. Le lissage spatial a stabilisé la variance de prédiction dans les

régions plus petites, en échange d'une petite perte de précision. Dans les régions où la population est plus importante, le lissage n'a généralement pas été jugé bénéfique ou nécessaire. La décision d'inclure le lissage spatial peut être prise dans le contexte des caractéristiques des données d'un ensemble de régions administratives.

Enfin, nous avons décrit comment le modèle proposé au chapitre 3 peut être utilisé dans un système de surveillance en temps réel. Nous avons utilisé l'événement de la fumée de feu de forêt de juillet 2015, dans le sud de la Colombie-Britannique, où les concentrations de PM2,5 étaient jusqu'à dix fois plus élevées que les niveaux de référence. Les consultations médicales liées à l'asthme et les dispensations de médicaments de soulagement ont augmenté respectivement de 34% et de 27% pendant cette période. Nous avons créé un scénario de surveillance en appliquant notre cadre compréhensif de surveillance de la fumée de feu de forêt, en utilisant des prévisions de fumée disponibles à l'échelle nationale, afin de prévoir les impacts sur la santé et d'estimer les réductions de morbidité qui auraient pu être obtenues si un tel système avait été utilisé pour appuyer les décisions d'intervention. Nous avons constaté que notre approche ne nécessitait pas de prévisions de fumée précises pour produire des prévisions d'indicateurs de santé raisonnables, mais elle a été négativement affectée lorsque les modèles d'exposition à la fumée n'a pas prévu de conditions très enfumées. D'après notre évaluation théorique de l'intervention, nous avons constaté que des interventions simples et précoces à plus faible efficacité, comme la messagerie de santé publique sur les mesures à prendre pour réduire l'exposition, étaient préférables aux interventions plus tardives mais avec une efficacité plus élevée, comme l'évacuation. Des recherches supplémentaires sont nécessaires sur l'interpolation spatiale et la prévision future de la pollution atmosphérique et de la fumée, ainsi que sur les interventions visant à faciliter un cadre décisionnel plus intégré.

Ce travail offre une approche méthodologique souple qui peut être étendue et améliorée à mesure que progressent les recherches connexes, telles que l'amélioration d'estimation de l'exposition et la disponibilité de données supplémentaires sur la surveillance de la santé en temps réel. Les méthodes développées et évaluées dans les deux premiers manuscrits et mises en ¡uvre dans le troisième sont actuellement intégrées dans le système existant de surveillance de la santé publique au Centre for Disease Control de la Colombie-Britannique et seront utilisées pendant la saison des feux de forêt de 2017.

Acknowledgements

Well here I am...finishing 12+ years of university. McGill has been my most challenging and interesting adventure.

I was very fortunate to have two wonderful supervisors, David Buckeridge at McGill and Sarah Henderson at UBC/BCCDC. You have each provided invaluable mentorship and guidance for my thesis research, for PhD survival, and for figuring out the rest of my life. What I have learned from each of you over the past five years is not sufficiently summarized within these pages. I finally feel like a competent scientist, and this would not have been possible without you both.

Special thanks are in order to Gavin Shaddick, my unofficial 3rd supervisor. I can't thank you enough for going above and beyond for me these past years. I have so enjoyed learning from you and our chats in many cities. Most of my thesis would have been impossible without your support, wisdom, enthusiasm, and time.

I have learned so much from my classmates, many who are now close friends. Special thank yous and hugs to Miriam, Corinne, Ryan, Jenn, Samuel, Maggie, Aman, and Kate. I couldn't have made it through this program without you all.

I'd also like to thank Patricia Plouffe and Katherine Hayden for all of their help over the years.

To my friends and family in Victoria and Philadelphia, thank you for helping me maintain (and, at times, locate) my sanity. To my wonderful partner Matt, who deserves an honorary degree from living through this experience with me, thank you for everything. Your patience, pep talks, programming help, and especially your genuine interest in my work so often kept me going. Time for our next adventure!

Contributions of authors

The objectives and structure of the three manuscripts in this thesis developed out of collaboration between me and my co-authors. Under the guidance of my supervisory committee (Dr. Buckeridge, Dr. Henderson, and Dr. Shaddick), I developed the research objectives, reviewed methodological and substantive literature, performed all data management, data simulations and statistical analyses, and wrote the first draft of each manuscript.

All co-authors reviewed and approved the final version of each manuscript.

Manuscript 1: Morrison KT, Shaddick G, Henderson SB, Buckeridge DL.
A latent process model for forecasting multiple time series in environmental public health surveillance. Published in *Statistics in Medicine*, (2016) 35(18): 3085-3100.

Manuscript 2: Morrison KT, Shaddick G, Thomas ML, Henderson SB, Buckeridge DL.
Improving forecast precision in environmental public health surveillance: a spatio-temporal approach. To be submitted to *Spatial and Spatial-Temporal Epidemiology*

Manuscript 3: Morrison KT, Buckeridge DL, Shaddick G, Henderson SB.
One step ahead of the epidemiologic curve: forecasting the public health impact of a wildfire smoke event. To be submitted to *Environmental Health Perspectives*.

David L. Buckeridge is an Associate Professor in the Department of Epidemiology and Biostatistics at McGill University. Dr Buckeridge has researched extensively in public health informatics and was closely involved in all three manuscripts as my primary supervisor, providing support on the methodological and substantive aspects of my doctoral research.

Sarah B. Henderson is an Assistant Professor in the School of Population and Public Health at the University of British Columbia, and a Senior Scientist at the British Columbia Centre for Disease Control, Environmental Health Services division. Dr. Henderson has researched extensively in environmental epidemiology and wildfire smoke surveillance. She gave substantive guidance and feedback for the first and second manuscripts, and provided primary supervision for the third manuscript.

Gavin Shaddick is a Reader in the Department of Mathematical Sciences at the University of Bath. Dr. Shaddick is a biostatistician who has published extensively in environmental statistics. He provided primary methodological supervision and guidance for the first and second manuscripts.

Matthew L. Thomas is a doctoral candidate in the Department of Mathematical Sciences at the University of Bath. He provided methodological and computational support for the second manuscript during my time visiting the University of Bath.

Statement of originality

The work contained in this thesis represents an original contribution to the field of environmental public health surveillance. To my knowledge, we have presented a novel methodological approach to forecasting multiple health outcomes with a latent exposure in a computationally efficient inferential framework. We extended this approach to include spatio-temporal smoothing to allow the proposed methods to be more useful in regions with smaller populations, where instability in the prediction variance may reduce the utility of the forecasts. I provide a computationally efficient way to obtain statistically valid predication intervals, to increase the utility of the forecasting beyond a single value. Finally, I implement the proposed framework into a surveillance system for wildfire smoke surveillance in British Columbia, demonstrating how such a system could have provided timely information for public health. I estimate the preventable increases in acute respiratory morbidity indicators had such a system been in place during the severe smokey period in July 2016.

As a whole, this work provides a flexible methodological approach that can be extended and improved as related research advances, such as improvements in exposure estimation and availability of additional real-time surveillance health data. The methods developed and evaluated in the first two manuscripts and implemented in the third are currently being integrated into the existing public health surveillance system at the BC Centre for Disease Control, and will be used during the wildfire season of 2017.

While I have received guidance from my supervisory committee on methodological and substantive areas, I declare that the conception, execution, and drafting of the work in this thesis were my own.

Statement of financial support

I am very grateful for the financial support that I received over the course of my research and studies at McGill University, making this degree possible. The first year of my Ph.D. was jointly funded by my primary supervisor, Dr. David Buckeridge, and an entrance award provided by the department of Epidemiology, Biostatistics and Occupational Health. The following three years were supported by a Frederick Banting and Charles Best Doctoral Award from the Canadian Institutes of Health Research (CIHR). Dr. Henderson kindly provided support for me to spend time with her at the BCCDC in Vancouver. Dr. Shaddick generously supported my time in Bath, England in 2016 for work on my second manuscript and to attend a environmental statistics workshop. The final eighteen months of my degree have been generously supported by a stipend from Dr. Buckeridge.

I received a Graduate Research Enhancement and Travel Awards travel award from McGill University and the Department of Epidemiology and Biostatistics in 2015-2016 academic year, which allowed me to present my first manuscript at the *XXVIII*th *International Biometric Conference* in Victoria, BC.

Table of contents

Tables

Figures

Abbreviations

BC	British Columbia
BCAMS	BC Asthma Monitoring System
BCCDC	BC Centre for Disease Control
BUGS	Bayesian inference using Gibbs sampling
CAR	Conditional autoregressive
CI	Credible interval
EPHS	Environmental public health surveillance
FRP	Fire radiative power
GLM	Generalized linear model
GMRF	Gaussian Markov random field
ICD-9	International Disease Classification, 9th edition
iid	Independently identically distributed
INLA	Integrated nested Laplace approximations
JAGS	Just another Gibbs Sampler
LGM	Latent Gaussian model
LHA	Local Health Authority
MAE	Mean absolute error
MAPE	Mean absolute percentage error
MCMC	Markov chain Monte Carlo
MODIS	Moderate Resolution Imaging Spectroradiometer
MSE	Mean square error
MSP	Medical services plan
OSSEM	Optimized Statistical Smoke Exposure Model
PI	Prediction interval
$PM_{2.5}$	Particulate matter less than 2.5 μm in diameter
PSA	Public safety announcement
PST	Pacific Standard Time
WHO	World Health Organization

1 Introduction

Acute environmental exposures, such as wildfire smoke and extreme hot weather, are increasing in frequency and severity along with their population health impacts. Wildfire smoke is considered a globally important cause of mortality by the World Health Organization (WHO), with an estimated 339,000 deaths from landscape fire smoke each year [1]. Epidemiological research has also shown that short-term increases in ambient particulate matter (PM) lead to increased rates of healthcare utilization for respiratory conditions, including prescription drug dispensations, outpatient physician visits, ambulance dispatches, and hospital admissions [2, 3, 4]. These health effects from wildfire smoke are expected to increase because changes in land use and climate are increasing the frequency and severity of wildfires, exposing more individuals to the harmful effects of smoke each year [5].

Similar to infectious disease surveillance, the objective of an environmental public health surveillance (EPHS) system is to detect potential changes in the health status of the population, and to provide timely evidence for public health intervention during periods of hazardous exposure [6]. However, conventional surveillance methods and conceptual frameworks were generally designed for infectious disease. The implementation of EPHS poses some unique challenges, like accurately measuring exposures such as air pollution across space and time, and modelling how exposures influence multiple potential measures of health.

Public health decision makers who are responsible for recommending interventions to mitigate the dangers of wildfire smoke exposure have called for timely population health surveillance tools to inform their decisions [7]. However, EPHS activities thus far have focussed on monitoring the exposure and the various health impacts independently, without linking them to one another. Integrated modelling of these data sources within a surveillance system would facilitate the transformation of data into actionable information.

1.1 Research objectives

The overall goal of this dissertation was to propose and evaluate a methodological framework for EPHS, using wildfire smoke as an application. To achieve this goal, three primary objectives were addressed:

i Propose a statistical framework for forecasting multiple health outcomes arising from a common latent exposure, addressing some methodological and computational challenges in EPHS.

ii Evaluate the proposed framework using data from British Columbia (BC) by (a) extending the proposed model to include spatio-temporal smoothing to stabilize the prediction variance for application to regions with small populations, as is common in EPHS, and (b) evaluating the validity and computational efficiency of prediction intervals around forecasts.

iii Propose and evaluate the a surveillance framework for wildfire smoke, using nationally available smoke forecasts to predict the anticipated health impacts, and estimate the potential reductions in morbidity that may have been achieved if such a system had been operational during the severe urban wildfire event of July 2015 in southwestern BC.

1.2 Organization of the thesis

Chapter 2 presents background information on wildfire smoke and EPHS. Chapter 3 presents a methodological and inferential approach to forecasting multiple health outcomes with a common latent exposure, and evaluates the approach with a simulation study and a small case study from one region of BC. Chapter 4 extends the proposed latent process approach to spatio-temporal data and compares forecast performance for the temporal and spatio-temporal approaches across BC administrative units. Chapter 5 explores the implementation of the proposed methodological approach within a surveillance framework and estimates the potential impact of timely forecast-based interventions for the severe wildfire season of 2015. Chapter 6 provides a summary of conclusions from the three research manuscripts, with a focus on implications for public health and priorities for future research.

2 Background

In the first half of this chapter, I present biologic and epidemiologic evidence about the harmful effects of wildfire smoke on human health, and summarize the public health burden of wildfire smoke. In the second half of this chapter, I outline the major methodological challenges in wildfire smoke surveillance, the current state of research on exposure estimation, and the contribution of my work to moving EPHS research forward. Throughout this chapter I focus primarily on the Canadian context, where the most cutting edge research is being done.

2.1 Respiratory health effects of acute wildfire smoke

Wildfire smoke is an environmental exposure that has significant acute impacts on respiratory health [8, 3, 4]. In this thesis, the word *exposure* refers to the estimated individual-level inhalation of particulate matter based on a proxy ecological measure, as it is used throughout the relevant epidemiologic literature. The acute respiratory health effects of wildfire smoke exposure are an important public health concern because entire populations across large geographic areas can be exposed, including those hundreds or thousand of of kilometres away from the fire [9, 10]. The effects are especially strong in individuals with pre-existing respiratory conditions, such as asthma or chronic obstructive lung disease (COPD).

2.1.1 Smoke toxicology

Wildfire smoke is a complex mixture of solids, liquids, and gasses, many of which are known to be harmful to humans when inhaled. Of these substances, PM has been most extensively studied, often measured as PM_{10} (particles less than 10 microns in size) or $PM_{2.5}$ (less than 2.5 microns in size). In general, $PM_{2.5}$ is considered to be more harmful because it penetrates deeper into human lungs and can potentially enters the bloodstream [11, 12]. From a biological standpoint, there is a considerable literature on the toxicology of controlled exposure to woodsmoke PM [9, 8]. However, it is unclear which compounds in smoke actually cause all of the observed respiratory symptoms in humans [9]. Acute exposure to carbon monoxide (CO), nitrogen oxides (NO), and ozone (O_3) in wildfire smoke may also impact population health, based on human and animal toxicology studies [13]. Furthermore, some components of smoke are known carcinogens and many chronic effects have been associated with long term exposure to $PM_{2.5}$ [14], suggesting that populations regularly exposure to wildfire smoke may be at higher acute and chronic risk.

In Canada, most cities have mean annual background concentrations of PM_{10} below 30 $\mu g/m^3$, and background $PM_{2.5}$ below 20 $\mu g/m^3$ [15]. Rural areas often experience lower background concentrations, however the air quality can vary depending on neighbouring industry and transportation networks [15]. During wildfires, PM_{10} and $PM_{2.5}$ can easily reach or exceed ten times the background concentrations, and large fires can affect air quality at local, regional, continental, and global scales [1, 16]. In Canada, wildfires cause the worst acute air quality that most people are likely to experience within the country, with daily maximum $PM_{2.5}$ sometimes exceeding 250 $\mu g/m^3$.

2.1.2 Epidemiology of acute wildfire smoke exposure on respiratory health

The harmful respiratory effects of wildfire smoke exposure have been observed across many epidemiological studies in many contexts, showing a consistently significant association between increased exposure and respiratory health outcomes (Reid 2016). These results have been replicated in numerous studies using different study designs, various forms of exposure measurements, different levels of exposure severity, and in study areas from Canada, the United States, Australia, Europe, and Asia.

The health outcomes measured in association with smoke exposure range from mild (medication dispensations) to moderate (physician and emergency department visits) to severe (hospital admissions, mortality) [1]. The less severe respiratory health outcomes generally have larger observed impacts on population health because they occur more frequently and sooner after the exposure. As such, these outcomes have been the focus of systems used for public health surveillance of the smoke exposures. Respiratory health effects have been observed consistently across many studies with an approximately 5% increased risk of healthcare utilization per 10 $\mu g/m^3$ of increased $PM_{2.5}$, with smaller effects for more severe outcomes such as mortality, which increases by approximately 2% [17, 18]. Given that the concentrations of $PM_{2.5}$ can increase from typical background concentrations to over 200 $\mu g/m^3$ and that the prevalence of asthma is quite high (over 5% confirmed diagnoses in British Columbia, but often estimated to be closer to 10% [19, 20]) these effects can translate into large numbers of impacted individuals over short periods.

2.1.3 The public health burden of wildfire smoke

The WHO recognizes that smoke from wildfires, even far from the original source, can adversely affect the health of large populations, including increased daily mortality [21]. Wildfire smoke is considered a globally important cause of mortality and morbidity, with an estimated 339,000 deaths annually attributed to *smoke* from all landscape fires [1]. Low and middle income countries are more severely affected, but large fires also occur in high income areas such as parts of Canada, the United States, and Europe [1, 22]. Wildfire seasons in North America have become longer and more severe, with more fires and larger fires over time, concurrent with increases in global temperature (Figure 1) [23]. For example, Figure 1 shows the frequency of large fires (greater than 400 hectares) in the Western US since 1970.

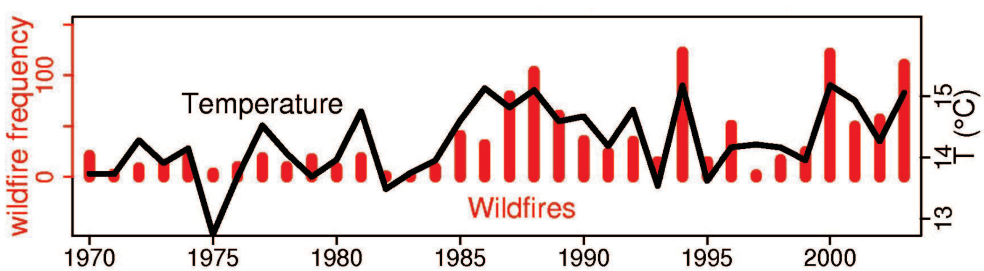

Figure 2–1: Western US forest fire frequency
Source: Westerling et al. 2006 [23]

Changes in the frequency, severity, and behaviour of wildfires have been attributed to climate change, landscape and land use changes, prescribed burning practices, and the potential interactions between these complex phenomena [23]. Wildfires are expected to continue to increase in frequency and severity over the coming decades. Given the increasing prevalence of asthma, concerns about other sources of chronic and acute air pollution (e.g., urban), the impact that climate changes have on temperature (e.g., heat waves, droughts), and the potential interaction between these factors, the health effects of wildfires are expected to increase dramatically. Furthermore, evidence suggests that smoke from wild and prescribed fires will play an increasingly important role in lifetime air pollutant exposures as emissions from other sources, such as vehicles and industry, come under tighter control [5].

Research on the economic impacts of wildfires has historically been focused on the loss of timber for the forestry industry, and the destruction of property. However, the health costs associated with wildfire smoke exposure are considerable, and research in the United States estimated that cost per exposed person per day could be as high as $85. Economics are an important facet of wildfire impacts, given that hundreds of thousands or even millions of individuals can be exposed per year in California alone, and that the fire season continues to lengthen [24]. Canadian research estimated the cost of a 2001 Alberta wildfire at approximately $12 million for the 670,000 individuals exposed to concentrations of $PM_{2.5}$ in the 35 to 55 $\mu g/m^3$ range for one day, which was a conservative estimate. These economic costs are generally based on the increased healthcare utilization due to increased risk of illness during high levels of PM exposure, although more sophisticated assessments also consider losses of productivity. Cost-benefit analyses could seek to balance these costs against the cost of an intervention (such as evacuation).

To briefly demonstrate the public health impacts of smoke exposure, consider a hypothetical cohort of one million individuals. Assume that the prevalence of asthma in this population is 10% [19], and that on any given day in the summer season the expected number of salbutamol dispensations is 500, the expected number of asthma-related physician visits is 200, and expected all-cause mortality is 20 individuals, with all estimates based on empirical data from BC. Research has shown (with replicated results) that for each 10 $\mu g/m^3$ increase in $PM_{2.5}$ there is an expected increase in healthcare utilization of approximately 5% and an expected increase in all-cause mortality of approximately 2% [17, 12]. Therefore, given a scenario where the PM increased from background $PM_{2.5}$ concentrations of 10 $\mu g/m^3$ to concentrations of 50 $\mu g/m^3$, I would expect 125 additional salbutamol dispensations, 50 additional asthma-related physician visits, and 2 additional deaths per day at this level of exposure. If the $PM_{2.5}$ increased to 150 $\mu g/m^3$, which has occurred during past fire seasons in BC, I would expect an additional 375 salbutamol dispensations, 150 additional asthma-related physician visits, and 6 additional deaths per day at this level of PM exposure. These simple calculations assume a linear dose-response relationship at the concentrations explored here, which is supported by research and empirical data for acute PM exposure [25].

2.2 Challenges in EPHS research

As changes in land use and climate drive increasingly frequent environmental events, such as extreme heat and acute episodes of air pollution, environmental surveillance becomes an increasingly important tool. Surveillance is a core public health function, defined by the Public Health Agency of Canada (PHAC) as (bold added for emphasis):

*"The **tracking and forecasting** of any health event or health determinant through the on-going collection of data, the **integration, analysis and interpretation of data** into surveillance products, and the **dissemination** of those surveillance products to those who need to know in order to undertake necessary actions or responses. **The overarching purpose of surveillance is to generate information and knowledge for public health action** in the short and long term*" [26].

Environmental surveillance of an acute exposure, such as air pollution, involves identification and collection of data, analysis of data, implementation of a surveillance system to give the results context, and use of the information therein to inform decision making. Expert methodological and substantive knowledge informs every stage.

In the surveillance and health informatics literature, using non-traditional data sources to monitor population health is referred to as *syndromic surveillance* [27]. For example, conventional surveillance may monitor laboratory-confirmed cases of an infectious disease, whereas a syndromic surveillance system may passively monitor sales of over-the-counter pharmaceuticals [28, 29]. Other examples of data used in syndromic surveillance include prescription drug dispensations, online search trends, physician visits, emergency department visits, hospitalizations, and cause-specific or all-cause mortality [30, 31]. In infectious disease surveillance, syndromic surveillance data can potentially facilitate early outbreak detection at a lower cost than conventional methods [32]. However, because these data are not collected for the purposes of surveillance, they contain more artifacts and noise than laboratory-confirmed data collected for classic surveillance. For example, individuals are more likely to see a doctor on a weekday than a weekend. Syndromic surveillance data generally require sophisticated methods to disentangle the health effect of interest (the signal) from the other effects (the noise) inherent to the data. In an environmental surveillance setting, outbreak detection is not the primary goal because the precipitating event (e.g., a wildfire or extreme heat event) is generally known to occur or not occur. However, syndromic surveillance data can provide rapid information on 1) whether the event is having a measurable impact on the health of the exposed population, or 2) is likely to have an impact in the near future. Forecasting the likely health effects under different exposure scenarios can reduce uncertainty and facilitate decision making related to the implementation of interventions to reduce exposure and mitigate the public health impact.

Most of the methodological tools used in public health surveillance were developed or adapted for conventional systems based on infectious disease models [33]. Environmental public health surveillance shares many objectives with infectious disease surveillance, but because it seeks to detect and forecast the short term *health impact of an exposure*, it presents some different methodological challenges [6, 34]. I have identified six major challenges for timely environmental public health surveillance:

1. Exposure estimation
2. Health outcomes measurement
3. Statistical modelling
4. Implementation
5. Heterogeneous populations
6. Interventions and assessment

I elaborate on each of these challenges, summarizing the existing state of the research and identifying the gaps that I sought to address in this thesis.

2.2.1 Exposure estimation: measuring air quality across large geographic areas

Ambient PM is a challenging environmental exposure to measure from an epidemiological perspective. Once smoke from a wildfire is emitted, the smoke moves across space over time, and the PM concentrations depend on many factors, such as temperature, humidity, and wind. Individual-level exposure to PM is usually estimated from environmental measurements, because collecting true individual exposure measurements is simply not feasible for an entire population. Exposure assignment can be very simple, such as a binary (smoke or no smoke) classification based on presence of smoke or fire in a given region using remotely sensed data (satellite imagery), or it can be a continuous measurement of PM_{10} or $PM_{2.5}$ from air quality monitors or an estimate from predictive models.

All of the available exposure measurements have strengths and limitations, and there is no established gold standard. Air quality monitors provide measurements multiple times per day but are located sporadically across provinces or states, providing a low spatial resolution. Remotely sensed data from platforms such as the Moderate Resolution Imaging Spectroradiometer (MODIS) have a high spatial resolution, but do not directly measure PM and may only allow binary measurements (e.g., smoke presence or absence per pixel) in some cases. Research has moved away from single exposure measurements to the development of predictive models that integrate data from air quality monitors, remote sensing, and meteorological data to estimate PM [35, 4]. Predictive models can be validated using air quality monitors in regions where monitors are available, which approximates an ecological level gold standard in those areas. Model-based measurements that combine air quality monitors with additional predictors are considered to be more flexible, provide more stable estimates over time, allow finer spatial and temporal resolution, and potentially facilitate air quality forecasting.

In BC, researchers at the BCCDC have explored several methods of estimating PM exposure for the purpose of wildfire smoke surveillance in the province. The regulatory air quality network of monitors is the most readily available source of air quality assessment [36]. The location and number of monitoring stations varies between wildfire seasons; there were 43 in 2010 providing hourly $PM_{2.5}$ ground-level measurements, and roughly equal numbers measuring PM_{10} [37]. Environmental scientists at the BCCDC developed the Optimized Statistical Smoke Exposure Model (OSSEM), which predicts ground-level $PM_{2.5}$ per 5×5km grid in populated areas throughout BC [37]. The model was fit using data from $PM_{2.5}$ air pollution monitors, remotely sensed aerosol and fire location data, expert tracings of smoke plumes from satellite imagery, a fire danger rating, and the atmospheric venting index provided by the BC Ministry of Environment and Environment Canada [37]. The model was validated with $PM_{2.5}$ measurements in regions with at least one monitoring station within the grid cell, and has been shown to provide accurate estimates ($R^2 > 0.70$) for ground-level PM.

In addition to using prediction models to improve spatial resolution, forecasting $PM_{2.5}$ over time is desirable for public health surveillance. For example, the BlueSky Western Canada Smoke Forecasting System produces hourly ground-level forecasts of $PM_{2.5}$ produced by wildfires, starting at 05:00 PST for each of the following 48 hours. The BlueSky forecasting model uses data from Canadian Wildland Fire Information System (via Natural Resources Canada) and the Canadian Forest Fire Behaviour Prediction system to provide information on rate of fire spread, fuel consumption, and fire intensity, as well as satellite data from MODIS and other sensors and data on weather and topography [38]. Research evaluating BlueSky forecasting and application in public health surveillance is ongoing [39, 35, 40].

2.2.2 Health outcomes measurement: identifying data sources for surveillance

Research shows that wildfire smoke has a wide range of health effects, from symptoms such as eye irritation to an increase in all-cause mortality. The respiratory health effects have been the focus of most epidemiological studies, and the results have been consistent across many studies [4], particularly for outcomes related to asthma and chronic lung disease. As such, acute respiratory outcomes are an important focus for tracking the population health response to smoke exposures [7].

Surveillance data for real-time assessment of health impacts during wildfire smoke events may be obtained from primary data collection, existing surveillance systems, or by repurposing of data collected for other functions. Primary data collection was historically more common and has been used in many retrospective epidemiologic studies, but is generally thought to be too time delayed and cost prohibitive for modern large-scale surveillance needs. Electronic administrative data are more commonly available now, and there are a wide range of administrative data used for the purpose of surveillance; generally, the healthcare utilization data previously described are most common. Many of these datasets can be classified by chief complaint or cause of death, narrowing the focus from all-cause to more specific measures such as respiratory hospitalization or asthma-related physician visit.

Data from existing surveillance systems monitoring relevant health outcomes could feasibly be used to establish situational awareness for the population health effects of smoke events [41]. Because health outcomes responsive to smoke include respiratory syndromes, existing surveillance systems that obtain data for other purposes, such as influenza, could be shared for these purposes, provided that there is legal authority to do so. While most states or provinces do not currently have a pre-existing surveillance system in place for wildfire smoke health effects, public health departments could plausibly rapidly establish data sharing agreements during smoke events.

In BC, scientists at BCCDC Environmental Health Services established the British Columbia Asthma Medication Surveillance (BCAMS) in 2012 [42]. The BCCDC is currently leading internationally in wildfire smoke surveillance research. The BCAMS system provides near real-time surveillance of exposure and health outcomes to local health authorities to provide situational awareness for public health and emergency management decision-making. This system uses all three exposure estimates described in the previous section: measured $PM_{2.5}$ from monitoring stations, estimated $PM_{2.5}$ from OSSEM, and forecasted $PM_{2.5}$ from the BlueSky system.

The BCAMS system uses two indicators of population respiratory health for monitoring the effects of wildfire smoke: dispensations of salbutamol sulfate, and asthma-related physician visits. Salbutamol is a medication used to alleviate exacerbations of chronic lung disease, and dispensations have been shown to increase rapidly and significantly during fire smoke episodes in British Columbia [43, 44]. Asthma-related physician visits are taken from medical billings data and classified using ICD-9 codes [45]. Excursions from the expected number of daily dispensations are identified by Public Health Intelligence for Disease Outbreaks (PHIDO) using an algorithm adapted from one originally developed for infectious diseases [46]. The PHIDO algorithm uses iterative regression to identify excursions beyond the 95th (unusual), 99th (rare) and 99.5th (very rare) percentiles of the expected daily distributions. Reports are provided retrospectively. The BCCDC is actively engaged in research on exposure estimation in wildfire smoke and surveillance of environmental exposures. Improvement and evaluation of the BCAMS system is ongoing.

2.2.3 Statistical modelling

There are many approaches to modelling time series data. For example, classic forecasting models are frequently used in infectious disease surveillance, such as autoregressive integrated moving average (ARIMA) models [47]. However, EPHS offers some unique challenges because the goals and data differ slightly from other types of surveillance.

In infectious disease surveillance, regardless of how cases are defined, historical data are generally used to establish seasonal patterns, establish expected trends, and detect aberrations from these expectations. In EPHS, I want to forecast anticipated health outcomes under different exposure scenarios, where prior epidemiologic research has established a clear relationship between the exposure the health outcomes. In the case of wildfire smoke, the exposure may be considered *latent*, because $PM_{2.5}$ essentially cannot be measured without some amount of error across large geographic areas [48]. Therefore I can conceptualize the health outcomes that are caused by wildfire smoke exposure as arising from a common underlying process that I can measure imperfectly over time. These health outcomes data are aggregate counts that are correlated both with themselves and one another over time.

There has been a growing interest around use of multivariate[1] approaches in surveillance, because the utility of non-traditional data sources can often be improved by borrowing strengths from other traditional or non-traditional sources [49, 50]. Multivariate extensions to the ARIMA model (such as vector autoregressive approaches) are most popular, because software is readily available to implement them [51]. Multivariate models are still relatively rare in the environmental epidemiology literature, but Dominici *et al.* 2004 modelled cardiovascular hospitalizations and mortality counts against air pollution (PM) for ten American cities. Her bivariate approach provided more efficient estimates of the relative rates, facilitated prediction of both outcomes in cities where only one was measured, and reduced the prediction variance [52].

A hierarchical Bayesian framework allows for great flexibility in modelling data that are correlated over time and/or space. An autoregressive forecasting model can be constructed as multivariate via mutual dependency on a latent exposure, compared with more cumbersome approaches, such as placing multivariate structures on the parameters. This hierarchical framework also allows for inclusion of unique covariates for each outcome, and can be extended to include different structures in the data, such as spatio-temporal smoothing, nonlinear effects of covariates, or more sophisticated measurement error models of the exposure. I elaborate further

[1] Throughout this thesis, I use *multivariate* to refer to multiple outcome variables in a single (potentially hierarchical) model.

on the approach I have identified in Chapter 3.

2.2.4 Implementation

The hierarchical Bayesian approach provides a more realistic way of modelling the EPHS data, but the computational burden of such a model implemented using conventional Markov chain Monte Carlo (MCMC) techniques is potentially enormous. Given that the interest is to apply such a model for (at least) daily use in a public health setting, computational efficiency must be a key constraint.

Throughout this thesis I discuss the computational advantages of an inferential approach using integrated nested Laplace approximations [53]. While MCMC is an asymptotically exact method, INLA is an approximation using latent Gaussian models (LGMs) and the theory of Gaussian Markov random fields (GMRFs); the error associated with the approximation has been comparable with MCMC error in simulation and case studies [54, 55, 56, 57]. Reframing a hierarchical model into an LGM can be non-trivial, but the marginal posterior distributions can be estimated with a computational efficiency that is improved by orders of magnitude as compared with MCMC [58]. Further details on the algorithm are summarized in Chapters 3 and 4.

2.2.5 Heterogeneous populations

Wildfires can expose large urban populations to wildfire smoke, but smaller communities are more frequently affected because they are often located in more heavily forested areas that are less aggressively protected when fires occur. Because entire populations can be exposed to smoke when it occurs, even small populations can carry a significant public health burden. However, forecasting with sophisticated hierarchical models involving many parameters can become unstable in regions with smaller populations. Smaller regions could potentially be aggregated for the purposes of health monitoring and decision-making, but this approach would lose valuable information in the process, potentially obscuring effects of the exposure in smaller areas. The demographic composition of a region can change how the population responds to the smoke, making it preferable to use smaller geographic units when possible [59]. I attempt to address these challenges in Chapter 4 by extending the previously proposed model from Chapter 3 to include spatial smoothing, hypothesizing that I can reduce the prediction variance while not losing excessive accuracy.

2.2.6 Interventions and assessment

In public health, the timeliness of interventions is often critical, and decisions must be made with the best data and evidence available. Public health officials currently act during a wildfire without the aid of surveillance and forecasting systems that can integrate data from multiple sources and provide timely information. Consequently, decision-making is based on descriptive analyses, including visual integrations of air quality measures, smoke forecasts, and healthcare utilization data. Some jurisdictions report using ad-hoc methods, such as repeated telephone calls to emergency departments or pharmacies to assess the health impact of the smoke on their community [37]. Given this situation, there is a demand for a real-time surveillance and forecasting system to decrease uncertainty in decision-making and estimate the burden of disease over time.

Interventions employed by public health departments are based on toxicological and epidemiological knowledge of smoke effects, but few interventions have had their effectiveness demonstrated in large-scale studies in real-world settings. For example, the most common public health intervention is messaging to encourage individuals to stay indoors during a smoky period, with the assumption that the indoor air quality will be superior to the outdoor air quality. This may be true for individuals who close all windows and who live in dwellings with relatively low infiltration rates (more common in newer structures), or who have central air conditioning. However, staying indoors may offer minimal protection if the infiltration of outdoor air is high [60], or there is significant generation of indoor PM from cooking or other activities. One research study has investigated the effectiveness of this intervention after a large fire and smoke event in California. Individuals were retrospectively surveyed and the researchers found that the ability to recall a public service announcement was strongly associated with a reduction in self-reported symptoms [61], but these results need to be replicated in different settings. Recommending the public, and in particular those with respiratory disease, to wear masks during smoke events is another common intervention. There is strong evidence that use of an effective air cleaner, such as High Efficiency Particulate Air (HEPA) filters, either in private or public air shelter scan drastically reduce the PM in the air Barn 2016[62]. Evacuation can theoretically be an effective way to reduce the population exposure to PM from a wildfire, but has considerable logistic, economic, and social consequences, and its effectiveness is not well established in the literature [63, 64, 61].

There is currently insufficient evidence on interventions to establish clear decision rules with respect to either exposure thresholds or anticipated population health effects. In the absence of this evidence, an exploration of the utility of the surveillance model presented in the Chapter 3 is presented in Chapter 5. I propose a surveillance framework that integrates monitored air pollution levels, monitored health outcomes, and forecasted smoke. I also explore the range of plausible impacts of a theoretical set of intervention with varying effectiveness and deployed at different times.

2.3 Summary

The epidemiologic effects of PM from wildfire smoke on respiratory health are well characterized and have been replicated frequently [2, 65, 12, 3, 4]. In general, EPHS is an area of growing importance and research on public health surveillance of wildfire smoke is needed in multiple related areas, including exposure estimation, health outcomes monitoring, methodological approaches, and analysis of public health interventions. In this thesis, I have attempted to identify the major barriers to development of effective surveillance system for EPHS. I propose and evaluate a methodological framework that is sufficiently flexible to be extended as research progresses in exposure estimation, as new health outcome datasets become available, and as more evidence on cost and effectiveness of interventions arise.

3 A latent process model for forecasting multiple time series in environmental public health surveillance

3.1 Preamble

This paper outlines a latent process model for forecasting multiple health indicators arising from a common environmental exposure. Traditionally, surveillance models in environmental health do not link health indicator measures, such as morbidity or mortality counts, to measures of exposure, such as air pollution. Moreover, different measures of health indicators are treated as independent, while it is known that they are correlated with one another over time as they arise in part from a common underlying exposure.

I propose modelling an environmental exposure as a latent process and describe the implementation of such a model within a hierarchal Bayesian framework, and its efficient computation using Integrated Nested Laplace Approximations (INLA). Through a simulation study, I compare separate univariate models for each health indicator to a bivariate approach. The bivariate model outperforms the univariate models in bias and coverage of parameter estimation, in forecast accuracy, and in computational efficiency. The methods are illustrated with a case study using healthcare utilization and air pollution data from British Columbia, Canada, 2003-2011, where seasonal wildfires produce high levels of air pollution, which has a significant impact on population health.

This manuscript was published in *Statistics in Medicine*.

3.2 Title page and footnotes

Title: A latent process model for forecasting multiple time series in environmental public health surveillance

Short title: A latent process model for environmental health surveillance

Authors: Kathryn Morrison[1], Gavin Shaddick[2], Sarah B. Henderson[3,4], David L. Buckeridge[1]

Affiliations:

[1] Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, Quebec, Canada

[2] Department of Mathematical Sciences, University of Bath, Bath, United Kingdom

[3] Environmental Health Services, British Columbia Centre for Disease Control

[4] School of Population and Public Health, University of British Columbia

Corresponding author:

Kathryn Morrison

1140 Pine Avenue West

Montreal, Quebec, Canada

H3A 1A3

E-mail: kt.morrison@mail.mcgill.ca

Citation:

Morrison KT, Shaddick G, Henderson SB, Buckeridge DL. 2016. 'A latent process model for forecasting multiple time series in environmental public health surveillance.' *Statistics in Medicine*, 35(18): 3085-3100.

3.3 Abstract

Introduction: Environmental health surveillance models do not typically link health indicator measures (e.g., morbidity counts) with measures of exposure (e.g., air pollution). Moreover, while different health events may arise from a common underlying exposure they are usually treated as though they are independent. The ability to reliably forecast health indicators under different exposure scenarios would facilitate real-time public health decision making.

Methods: We propose modelling an environmental exposure as a latent process within a hierarchical Bayesian framework, implemented using integrated nested Laplace approximations (using R-INLA) for efficient computation. We evaluate the approach through simulation, comparing univariate models for each health indicator to a combined approach. We also illustrate the method with a case study from British Columbia, where seasonal wildfires affect air quality. We forecast two daily asthma-related health indicators (prescription dispensations, physician visits) using measures of fine particulate matter concentrations.

Results: The hierarchical modelling structure was formulated in terms of latent Gaussian models in order to allow implementation in R-INLA, resulting in crucial gains in computational efficiency compared to using MCMC. Overall, the bivariate approach was superior in performance compared to using separate univariate models. Summarizing results for one fire season in 2009, the bivariate model provided overall more accurate predictions, and credible intervals that more frequently contained the true value. Notably, the univariate approach incorrectly suggested a protective effect of air pollution in a substantial number of cases, which was rarely observed with the bivariate models.

Conclusion: Modelling multiple measures of health impact with a common latent exposure provides a more cohesive and flexible framework than traditionally used methods. We show that this approach can provide timely information for public health surveillance. Future work will extend this approach for data available over both space and time, allowing dependencies between administrative areas to be exploited, and explore the application of this framework for public health decision-making.

Keywords: latent processes, time series, INLA, air pollution, Bayesian hierarchical models

3.4 Introduction

Wildfire smoke is considered a globally important cause of mortality by the World Health Organization, with an estimated 339,000 deaths from landscape fire smoke each year [1]. Epidemiological research has also shown that short-term increases in ambient particulate matter lead to increased rates of healthcare utilization for respiratory conditions, including physician visits, hospitalizations, and prescription drug dispensations [66, 44]. These health effects from wildfire smoke are expected to increase as changes in land use and climate are increasing the frequency and severity of forest fires, exposing more individuals to the harmful effects of smoke each year [22, 5]. In some cases, the threat of large-scale exposures has led public health and emergency management officials to evacuate entire communities to mitigate exposure to harmful particulate matter [64, 67, 61, 68]. Public health decision-makers in areas affected by forest fire smoke who must determine when to issue evacuation orders or implement other interventions have called for timely population health surveillance to inform their decisions [42]. However, real-time information on exposure and human health during smoke events is currently fragmented and of limited use for guiding decisions about interventions. Surveillance activities have tended to focus on monitoring exposure, without linking these exposure to population health data in a manner that can guide real-time action. Similarly, measures of health effects, such as mortality and morbidity counts, are usually monitored independently without linking these data explicitly to smoke exposure. This fragmented or 'siloed' approach to surveillance persists, despite the fact that these various data used for surveillance all arise from a common underlying environmental process. Given the obvious relationships between the exposure and the health indicators, it stands that modelling the health indicators jointly along with the measurements of the underlying process should improve forecast accuracy through an integrated surveillance framework and ultimately produce better decisions about public health interventions. Methods, such as vector autoregressive approaches, exist for modelling multivariate time series with covariates, but these models are rarely applied in public health surveillance, likely due to computational and technical challenges, lack of flexibility, and strong assumptions [69].

Surveillance data for measuring the health effects of air pollution are generally in the form of aggregate counts of healthcare utilization, morbidity, or mortality, such as physician visits, emergency room visits, hospital admissions, or deaths per geographic region. Measures of exposure for air pollution, usually in the form of particulate matter concentrations, may be taken directly from nearby air quality monitors or in some cases are interpolated using predictive models [37, 70, 71]. Each health indicator can be regressed against the exposure to produce forecasts; additional covariates can be included, such as day-of-the-week, temperature, and demographic data. A Poisson generalized linear model with a log link is often used to model counts of health indicators. In many epidemiological studies, simple models like these have been used for modelling the effect of air pollution on each health indicator separately [72]. Less commonly, multiple indicators have been modelled jointly in a multivariate approach. There are multivariate methods presented in the disease mapping literature for multiple diseases, for both spatial and spatio-temporal data [73]. In the environmental health statistics literature, Dominici et al. 2004 modelled both cardiovascular hospitalizations and mortality against particulate matter for ten American cities [52]. By modelling both indicators simultaneously, they were able to obtain more efficient estimates of the relative rates, facilitate prediction of both indicators in cities where only one was measured, and reduce the prediction variance [52].

There are temporal patterns in health indicators that may not be fully explained by their relationships with measures of the exposure, such as day-of-the-week patterns and seasonal and secular trends. These patterns can be modelled as factor variables or smooth functions of time by using polynomial terms, regression or smoothing splines, or Fourier terms in generalized linear (GLM) or generalized additive (GAM) models [52, 74]. However, residual correlation will frequently remain, and therefore the residuals can be modelled as an autoregressive process to account for short-term autocorrelation in the data. Compared with accounting for autocorrelation via factor variables or flexible modelling approaches, modelling the autoregressive process can use fewer degrees of freedom while ensuring the variance is not underestimated from residual autocorrelation, and potentially provide more accurate predictions [48].

In this manuscript, we propose a novel approach to forecasting in environmental health surveillance. We use a hierarchical approach to multivariate time series modelling of health indicator counts along with continuous autoregressive exposures, with efficient computation using integrated nested Laplace approximations (INLA) [53]. The latent process model we describe can be generalized to Bayesian latent variable modelling, where we model the health indicators at the first level the hierarchy, the environmental exposure as the latent process at the second level of

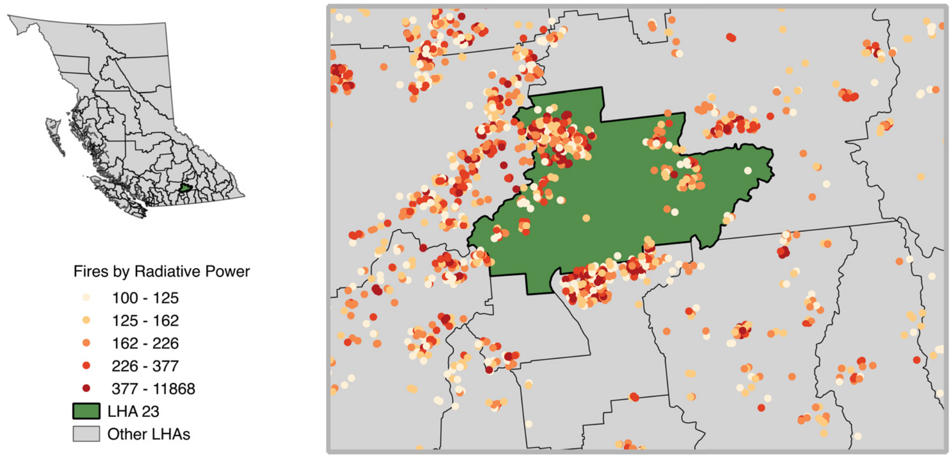

Figure 3–1: Location and severity of fires near the Central Okanagan Local Health Area (LHA #23). Fire severity is classified by radiative power, which is strongly associated with smoky conditions.

the hierarchy, and priors at the third level of the hierarchy, to obtain marginal posterior predictive distributions [48]. This method of modelling can be extended to include different structures in the data, such as spatio-temporal correlation, or non-linear effects in the exposure or other covariates. We explore our approach via a simulation study and present results from a case study using air pollution and healthcare utilization data from British Columbia.

British Columbia is the western most province in Canada. It is heavily forested, with an average of 2000 wildfires per season over the past ten years [66]. These fires can produce heavy smoke, introducing high levels of fine particulate matter ($PM_{2.5}$) into the air over the summer months [17]. The province is divided into 89 Local Health Areas (LHAs), and LHA 23 in the Central Okanagan Valley was used for our case study (Figure 3–1). This region includes the city of Kelowna with population approximately 110,000, and the entire region had population 187,058 in 2014 [75]. Figure 3–1 shows LHA 23 and the location of fires (categorized by severity) from 2003 to 2011 (see Section 3.8).

This paper is organized as follows: in Section 3.5, we present our proposed modelling approach, and then describe how we evaluate our approach via simulation and demonstrate the models for use in real-time surveillance and forecasting in the case study. We describe how the length of the time series used to forecast future values can be optimized to maximize forecast accuracy and avoid the computational costs of modelling a longer time series. In Section 3.6, we describe estimation using INLA, including how predictions (forecasts) are estimated in this framework and how the models proposed in Section 3.5 can be fit using the `R-INLA` package. Section 3.7 describes the stimulation study; in Section 3.8, we describe the case study using real surveillance and air pollution data. In Section 3.9, we offer some discussion and concluding thoughts.

3.5 Statistical Modelling

3.5.1 Hierarchical time series models

In this section we describe univariate and multivariate hierarchical models for modelling environmental exposure data with multiple correlated health indicators. The first level describes the health indicator data: the count for health indicator i at time t is modelled as a Poisson random variable, conditionally independent given their expected value μ_{it}.

$$\begin{aligned} &y_{it}|\mu_{it} \sim Poisson(\mu_{it}) \\ &log(\mu_{it}) = \alpha_i + \beta_i z_t + \sum_{j=1}^{J} \gamma_{ji} w_{jit} \end{aligned} \tag{3.1}$$

Where $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are the intercepts and effect estimates respectively, and $\boldsymbol{\gamma}$ are the regression parameters on any other covariates $\boldsymbol{w}$. It is assumed that z_t is a stationary autoregressive process of order p, common across all health indicators $i = 1, 2, ...I$.

$$z_t = \rho_1 z_{t-1} + \rho_2 z_{t-2} + ... + \rho_p z_{t-p} + \varepsilon_{zt} \tag{3.2}$$

The true underlying (latent) process z_t is unknown and cannot be directly obtained, but environmental exposure x_t (an imperfect measure of the true z_t) can be obtained as:

$$x_t = z_t + \varepsilon_{xt}$$

In a Bayesian analysis there is a third level for the prior distributions: $p(\varepsilon_{zt})$, $p(\varepsilon_{xt})$, $p(\alpha_i)$, $p(\beta_i)$, $p(\gamma_{ji})$. The two error terms, ε_{xt} and ε_{zt}, would not be identifiable under these conditions, therefore we impose a sum-to-zero constraint on the ε_{zt} term, which is reasonable given the assumed iid error assumed for the x_t measurement of z_t. If $i > 1$, then the multiple indicators arise from the same underlying process, $\boldsymbol{z}$, where $\rho_1, \ldots, \rho_p$, along with ε_{xt} and ε_{zt}, are estimated from the multiple Poisson time series $\boldsymbol{y_i}$, $i = 1, \ldots, I$ within a multivariate model. Conversely, each indicator can be fit independently, with separate estimates for ρ_p and the other lower level parameters. In either model, separate intercept and coefficient ($\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$) parameters are estimated for each health indicator, as each will likely have a different relationship with the exposure over time. In this paper, we explore and compare these two implementations, which we refer to as the univariate and bivariate models. Our goal was to present an efficient model where the multivariate structure is obtained via mutual dependency on the latent process, avoiding the considerably more cumbersome approach of placing multivariate structure on the intercept or effect parameters.

3.5.2 Assessing the performance of forecast models

There is a large body of research on applied forecasting, for example in the econometrics and machine learning literatures [76, 77]. The goal of developing and evaluating forecasting models is generally to minimize out-of-sample prediction error and methods similar to cross-validation are often employed to summarize prediction error. However, care must be used with time series data to ensure that the temporal ordering is preserved, unless there is interest in back-casting [78]. k-step ahead forecasting is often used, where up to k days in the future are predicted based on currently available data. The forecast accuracy and precision are usually measured with metrics like mean square error and standard errors to assess the performance of forecasting models [79].

In this manuscript we are proposing methods for use in public health surveillance of acute environmental exposures, where short-term prediction is of interest. We explore one-step-ahead forecasting, where the forecasting horizon y_{h_1} to y_{h_t} is explored via one-step-ahead predictions, with models estimated from data within the forecast window, y_{w_1} to y_{w_t}. For all analyses, $w_1 < w_t < h_1 < h_t$. For each parameter in our simulation study, we present bias and coverage estimates. We estimate bias as the median estimated posterior median across the simulations, and coverage of the credible intervals as the proportion of times the interval contains the true parameter. Additionally, we evaluate the performance of our forecasting models for one-step-ahead prediction both in the simulation study and in the case study.

3.5.3 Optimization of the moving forecast window

When evaluating forecast models for a time series, the length of the time series can be sequentially increased until the entire time series is used, or a moving-window approach can be employed, where the previous $w_t - w_1$ days are used to produce the k step ahead forecast for each day in the time series. For simulated data with no secular trends (stationary data), increasing length of the time series used in running the model to generate the forecast does not improve model performance (forecast accuracy), assuming the time series is sufficiently long to estimate ρ and the other parameters to capture the temporal patterns. The precision of the forecast will increase with additional data, but the improvement may be negligible after the time series reaches a certain length. However, for real data the optimal number of days supporting the forecast may not be equal to the maximum available days. Although it is rarely done, the width of the forecast window can be treated as a parameter and optimized using forecast accuracy measures. Modelling longer time series increases computation time considerably and may not significantly increase forecast accuracy and, in some cases, may actually decrease it.

3.6 Inference

The models presented in this paper could be estimated using Markov chain Monte Carlo (MCMC) methods; however, computation for the models described in Section 3.5.1, particularly the multivariate models, is very slow for even moderately sized datasets using MCMC. Slow computation may not be a serious deterrent for a single data analysis, but we have proposed these models for use in real-time public health surveillance, where timely information is critical. In considering the efficacy of using different forecast windows there is a requirement to produce implementations of the models and it proved to be computationally impractical to perform inference using such packages or bespoke MCMC. The evaluation via extensive simulation studies was also not feasible using MCMC given the computation times.

MCMC is often implemented using software such as WinBUGS/OpenBUGS [80] or JAGS, and is widely used in the applied statistical literature [81, 82, 83]. As noted by [84], care may have to be taken when using WinBUGS when not using built-in functions for auto-regressive processes where cyclical graphs may arise. If the temporal process here was a random walk then this can be formulated as a CAR process [85] and the WinBUGS `car.normal` function used, which is designed for such a purpose [86, 87]. No such function exists in JAGS and for both WinBUGS and JAGs the AR process will have to be coded, noting WinBUGS may construct full conditionals that lead to 'double counting' of the information from the neighbours, however this can negated by overriding the automatic construction of the full conditionals [84].

Here we demonstrate the use of techniques that perform 'approximate' Bayesian inference based on integrated nested Laplace approximations (INLA) and thus do not require full MCMC sampling to be performed [53, 88]. INLA has been developed as a computationally attractive alternative to MCMC for LGMs.

3.6.1 Integrated nested Laplace approximations

MCMC is an asymptotically exact method, while INLA relies on analytical approximations; however, the accuracy to which INLA can compute approximations to the posterior marginal distributions is well documented, see for example [53, 56, 89, 90]. While INLA is usually very accurate, [91] have shown some cases with binomial and Poisson data where a correction may be required, although it is noted that these are very extreme cases. A number of authors have recently reported favourable comparisons with results obtained using MCMC including [92], who fit a bivariate meta-analysis of diagnostic test accuracy studies and [93], who perform comparison for a variety of examples. When comparing the results from OpenBUGS and R-INLA in a disease mapping setting, [94] found some differences in the estimates of the random effects and their precisions when using R-INLA with the default priors, but found that exact replicates could be found by specifying alternatives.

This class of models available in INLA are those that can be recast as latent Gaussian models, such as generalized linear or additive (mixed) models, survival models, spatio-temporal models, and others where the following structure is assumed:

$$
\begin{aligned}
&\text{Likelihood: } \boldsymbol{y}|\boldsymbol{z},\boldsymbol{\theta_2} \sim \prod p(y_i|\boldsymbol{z},\boldsymbol{\theta_2}) \\
&\text{Latent field: } \boldsymbol{z}|\boldsymbol{\theta_1} \sim p(\boldsymbol{z}|\boldsymbol{\theta_1}) = \mathcal{N}(\mathbf{0},\boldsymbol{\Sigma}) \\
&\text{Hyperpriors: } \boldsymbol{\theta} = [\boldsymbol{\theta_1},\boldsymbol{\theta_2}]^T \sim p(\boldsymbol{\theta})
\end{aligned} \tag{3.3}
$$

where $\boldsymbol{y}$ are the observed data, $\boldsymbol{z}$ is the latent field, which is the joint distribution of all parameters in the linear predictor (such as the intercept, regression coefficients, and non-linear effects), $\boldsymbol{\theta_2}$ are the hyperparameters of the likelihood (such as the variance for a Gaussian outcome), and $\boldsymbol{\theta_1}$ are the hyperparameters for the latent field (such as the variance for a random effect). The latent field is assumed to be multivariate Gaussian with a conditional independence structure that results in a sparse precision matrix [58]. Estimates of the posterior marginals for the latent field can be derived from a Bayesian hierarchical model as:

$$p(z_j|\boldsymbol{y}) = \int p(z_j, \boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta} = \int p(z_j|\boldsymbol{\theta}, \boldsymbol{y})p(\boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta} \tag{3.4}$$

which are approximated by INLA as [95]:

$$\sum_k \tilde{p}(z_i|\boldsymbol{\theta}_k, \boldsymbol{y})\tilde{p}(\boldsymbol{\theta}_k|\boldsymbol{y})\Delta_k \tag{3.5}$$

where $\tilde{p}(\cdot)$ refers to an approximation. The first term in (4.5) refers to a Laplace approximation, while the second term depends on the implementation of the INLA program (a full Laplace approximation, a Gaussian approximation, or a simplified Laplace approximation); see [54]. The simplified Laplace approximation is the default and was used for all analyses in this manuscript. For a detailed explanation of the INLA approximation algorithm, see [53]. In our analyses, we use the implementation of INLA within the R package `R-INLA`, see [88].

3.6.2 Forecasting

Forecasting in a Bayesian framework is a form of out-of-bounds prediction, where the posterior distribution of the unobserved quantity, $\hat{y}_{t+1}$, given the observed data, $\boldsymbol{y}$, is desired. This marginal posterior distribution for the one-step-ahead forecast is obtained by integrating over the parameters in the latent field:

$$p(\hat{y}_{t+1}|\boldsymbol{y}, \boldsymbol{\theta}) = \int_{\boldsymbol{z}} p(\hat{y}_{t+1}|\boldsymbol{z}, \boldsymbol{\theta})p(\boldsymbol{z}|\boldsymbol{y}, \boldsymbol{\theta})d\boldsymbol{z} \tag{3.6}$$

and then numerical methods as outlined in equation (6) are used to approximate the integral over the hyperparameters $\boldsymbol{\theta}$; for more information, see [53]. The desired out-of-bounds predictions are entered into the INLA software as NA values, which forces the prediction when the model is fit

[53].

3.6.3 Implementation in `R-INLA`

The `R-INLA` syntax is similar to the R `glm` formula syntax. The form of the random effect is specified, for example, as `f(id,model='ar1')` for a first-order autoregressive prior as used in this paper, where `id` is a placeholder for the latent variable. These indicator variables, such as `id.alpha` shown in the vectors below, are used to translate the hierarchical model into the acceptable format of a latent Gaussian model. To construct the previously described hierarchical models in INLA, they must be entered as a joint model of the health indicator(s) and the latent exposure variable. In R-INLA, this is implemented by constructing a matrix for the response variable and for the latent exposure variable, and vectors for the intercept and covariates in R, for example [57]:

$$\underbrace{\begin{bmatrix} y_1 & \text{NA} \\ y_2 & \text{NA} \\ \vdots & \vdots \\ y_n & \text{NA} \\ \text{NA} & x_1 \\ \text{NA} & x_2 \\ \vdots & \vdots \\ \text{NA} & x_n \end{bmatrix}}_{\text{response.univariate}} = \underbrace{\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.alpha}} + \underbrace{\begin{bmatrix} 1 \\ 2 \\ \vdots \\ n \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.beta}} + \underbrace{\begin{bmatrix} \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ 1 \\ 2 \\ \vdots \\ n \end{bmatrix}}_{\text{id.rho}}$$

where the terms `response.univariate`, `id.alpha`, `id.rho`, and `id.beta` refer to matrix or vector names in R with the corresponding elements as specified above. The `id.rho` element is *copied* to the element `id.beta` in order to estimate β for the autoregressive covariate $\boldsymbol{x}$. The INLA copy command allows more than one element from a latent model to contribute to the linear predictor, as is common in hierarchical models and especially in multivariate models [57]. The `NA` values in the indicator variable(s) refer to an unobserved value and do not contribute to the likelihood. The `NA` values in the covariates remove it from the linear predictor. Therefore, different combinations of null values can be used to communicate the hierarchical structure of the model to the `R-INLA` program. An additional covariate W can be entered on the right side as

$\left[w_1, w_2, \ldots, w_n, \text{NA}, \text{NA}, \ldots, \text{NA}\right]^T$. It is possible to include covariates for one outcome and not for the other. For the bivariate model, the syntax is similar but the second indicator is included in the joint model formulation:

$$\underbrace{\begin{bmatrix} y_{1,1} & \text{NA} & \text{NA} \\ y_{2,1} & \text{NA} & \text{NA} \\ \vdots & \vdots & \vdots \\ y_{n,1} & \text{NA} & \text{NA} \\ \text{NA} & y_{1,2} & \text{NA} \\ \text{NA} & y_{2,2} & \text{NA} \\ \vdots & \vdots & \vdots \\ \text{NA} & y_{n,2} & \text{NA} \\ \text{NA} & \text{NA} & x_1 \\ \text{NA} & \text{NA} & x_2 \\ \vdots & \vdots & \vdots \\ \text{NA} & \text{NA} & x_n \end{bmatrix}}_{\text{response.bivariate}} = \underbrace{\begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.alpha1}} + \underbrace{\begin{bmatrix} \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ 1 \\ 1 \\ \vdots \\ 1 \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.alpha2}} + \underbrace{\begin{bmatrix} 1 \\ 2 \\ \vdots \\ n \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.beta1}} + \underbrace{\begin{bmatrix} \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ 1 \\ 2 \\ \vdots \\ n \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \end{bmatrix}}_{\text{id.beta2}} + \underbrace{\begin{bmatrix} \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ \text{NA} \\ \text{NA} \\ \vdots \\ \text{NA} \\ 1 \\ 2 \\ \vdots \\ n \end{bmatrix}}_{\text{id.rho}}$$

where `id.rho` is copied to both `id.beta1` and `id.beta2` to estimate the β_1 and β_2. Additional covariates can be entered as (for example):

$$\underbrace{\begin{bmatrix} w_1 & \text{NA} \\ w_2 & \text{NA} \\ \vdots & \vdots \\ w_n & \text{NA} \\ \text{NA} & w_1 \\ \text{NA} & w_2 \\ \vdots & \vdots \\ \text{NA} & w_n \\ \text{NA} & \text{NA} \\ \vdots & \vdots \\ \text{NA} & \text{NA} \end{bmatrix}}_{\text{id.w}}$$

3.6.4 Hyperprior formulation

The priors for each fixed effect in our analyses, including the intercepts, were $N(0, 1/\tau)$, $\tau = 0.001$, and this was also the prior on ε_{xt}. The prior on ε_{zt} is defined by default within the `ar1` INLA latent process model as $N(0, \tau)$ where the log precision is assigned a loggamma prior with shape and scale parameters 1, and 5e^{-5} respectively. We also assume a first-order autoregressive process throughout, i.e., $\rho_2, ..., \rho_p = 0$ in Equation (3.2). The default R-INLA prior for the hyperparameter ρ in a first-order autoregressive process is defined as:

$$log\frac{(1+\rho)}{(1-\rho)} \sim N(0, 1/\tau_\rho) \qquad \text{where } \tau_\rho = 0.15 \tag{3.7}$$

This translates to:

$$\rho = 2\frac{exp(\theta)}{1+exp(\theta)} - 1 \qquad \text{where } \theta \text{ is the logit of } \rho \tag{3.8}$$

which can result in a bimodal distribution (notably when using the default values). Using `beta.correlation` to define the prior in R-INLA defines a prior on θ such that the correlation parameter, ρ, has a Beta(a,b) distribution scaled to have domain in (-1, 1). Here we use Beta(5,5), which will result in a uni-modal distribution.

3.7 Simulation Study

In this section, we describe the implementation and results of a simulation study to assess the performance of the models described in 3.1. We estimate coverage of model parameters and assess forecasting accuracy with mean absolute percentage error (MAPE) and forecast coverage.

3.7.1 Data generation

We consider a range of simulation scenarios based on the following data generation process. We begin with a first-order autoregressive process z, and the log-expected values of the two health outcomes $\boldsymbol{y_1}$ and $\boldsymbol{y_2}$ are a linear function of the underlying process z. We measure z with iid error as x. We did not include any additional covariates or seasonal trends.

$$
\begin{aligned}
&e_t \sim N(0, 1/\tau_z) \text{ and } z_1 = e_1 \\
&z_t = \rho z_{t-1} + e_t, \quad \text{for } t = 2, 3, ..., T \\
&x_t \sim N(z_t, 1/\tau_x) \\
&y_{1t} \sim \text{Poisson}(\exp(\alpha_1 + \beta_1 x_t)) \\
&y_{2t} \sim \text{Poisson}(\exp(\alpha_2 + \beta_2 x_t))
\end{aligned} \tag{3.9}
$$

The parameters were defined as shown below in Table 3–1. The parameters of primary interest were the effect sizes between the exposure and the health indicators (β_1 and β_2), and the autoregressive parameter ρ. We varied the ρ parameter from 0.2, to 0.8 in increments of 0.1, and varied the β_1 and β_2 parameters between log(1.7) and log(2.0) in increments of log(0.1). The effect size parameters (β_1, β_2) were inflated to make differences in parameter retrieval more easily detectable [96, 17, 9]. The ρ parameter contributes to the strength of the bivariate correlation between the two health indicators and is expected to affect the performance of the forecasting models; were were primarily interested in the impact of this correlation parameter. The priors for each parameter are as specified in section 4.4.

Table 3–1: Simulation data generation scenarios

Parameter	α_1	α_2	β_1	β_2	ρ	τ_x	τ_z
Scenarios 1-7	3	2.8	log(2)	log(2)	0.2 to 0.8	1/0.5	1/0.1
Scenarios 8-14	3	2.8	log(2)	log(1.8)	0.2 to 0.8	1/0.5	1/0.1
Scenarios 15-21	3	2.8	log(1.8)	log(1.8)	0.2 to 0.8	1/0.5	1/0.1
Scenarios 22-28	3	2.8	log(1.8)	log(1.7)	0.2 to 0.8	1/0.5	1/0.1

We simulated 50 days of data, and ran 200 simulations for each of the 28 scenarios. To estimate bias and coverage for the parameters, we simulated and fit the models using the full 50 days. Then to evaluate forecasting performance, we removed the data for the 50th day, re-fit the models with the first 49 days, and forecasted one day ahead. The posterior predictive distribution for $\hat{y}_{i,t=50}$ is:

$$
p(\hat{y}_{i,t=50}|\boldsymbol{y}, \boldsymbol{x}) = \int_{\boldsymbol{\theta}} p(\hat{y}_{i,t=50}|\boldsymbol{\theta}) p(\boldsymbol{\theta}|\boldsymbol{y}, \boldsymbol{x}) d\boldsymbol{\theta} \tag{3.10}
$$

where $\boldsymbol{\theta}$ is a vector of parameters in the given model ($\alpha, \beta, \tau_x, \tau_z$, and ρ), and (3.10) is estimated via INLA as described in Section 3.6.1.

3.7.2 Evaluation

We summarize the performance of each model (univariate and bivariate) for each health outcome series ($\boldsymbol{y_1}$ and $\boldsymbol{y_2}$) using:

(i) The bias, estimated by taking the median of the marginal posterior distribution for a given parameter θ, and finding the median of these point-estimates across the 200 simulations: $(\theta - \hat{\theta})/\theta$ for the intercepts α_1 and α_2, β_1, β_2 autoregressive parameter ρ, and precisions τ_x and τ_z.

(ii) The coverage of the 95% credible intervals for each parameter.

(iii) The mean absolute percentage error of the forecasts: $\frac{1}{n}\sum_{i=1}^{n}(|\hat{y_i} - y_i|/y_i)$.

(iv) The coverage of the 95% credible intervals for the forecasts.

3.7.3 Simulation results

Varying β_1 and β_2 had little effect on the coverage, bias, and forecast performance of the models, so the results in this section are based on the 200 simulations where $\beta_1 = log(1.8)$ and $\beta_2 = log(1.7)$. Conversely, the ρ parameter had an impact on the model performance, and we focus results for different values of ρ.

Table 3–2 shows the bias and coverage for α_1, α_2, β_1, and β_2 by values of ρ = 0.2, 0.4, 0.6, and 0.8. When the value of the ρ parameter was higher for data generation, the coverage was lower for the intercepts (α's) and, to a lesser extent, for the exposure effects (β's). Coverage ranged from 22% for α_1 when ρ = 0.8 to 72% when ρ = 0.2. The estimates for β_1 and β_2 had overall better coverage, ranging from 60% when ρ = 0.8 to over 90% for when ρ was lower. The overall bias (based on the median of 200 posterior medians) for the intercepts and exposure effects parameter estimates was mostly below 5% and fairly consistent as ρ varied in the different simulation scenarios (Table 3–2). The univariate and bivariate models performed relatively similarly for retrieving the intercepts, but the bivariate model had higher coverage than the univariate model when retrieving the effect (β) parameters.

Table 3–2: Summary[1] of simulation results for intercept and effect size parameters by ρ parameter value.

		$\alpha_1 = 3.0$, $\alpha_2 = 2.8$			$e^{\beta_1} = 1.8$, $e^{\beta_2} = 1.7$		
	Model[2]	**Median**[3]	**Bias**	**Coverage**	**Median**[3]	**Bias**	**Coverage**
ρ=0.8	U1	2.958	0.014	0.220	1.843	**-0.024**	0.600
	B1	2.972	**0.009**	**0.240**	1.888	-0.049	**0.880**
	U2	2.758	**0.014**	**0.240**	1.769	**-0.041**	0.600
	B2	2.746	0.018	0.220	1.781	-0.048	**0.940**
ρ=0.6	U1	2.982	0.006	0.400	1.930	-0.072	0.820
	B1	2.981	0.006	0.400	1.911	**-0.062**	**0.900**
	U2	2.809	-0.003	0.400	1.779	**-0.047**	0.700
	B2	2.793	**0.002**	**0.380**	1.828	-0.075	**0.900**
ρ=0.4	U1	3.005	-0.002	0.480	1.869	-0.038	0.880
	B1	2.999	**0.000**	**0.500**	1.866	**-0.037**	**0.920**
	U2	2.801	**0.000**	0.620	1.772	-0.042	0.880
	B2	2.803	-0.001	**0.660**	1.770	**-0.041**	**0.920**
ρ=0.2	U1	3.003	**-0.001**	**0.520**	1.892	-0.051	0.780
	B1	2.992	0.003	0.500	1.850	**-0.028**	**0.920**
	U2	2.804	-0.001	**0.720**	1.758	**-0.034**	0.860
	B2	2.800	**0.000**	0.680	1.786	-0.050	**0.940**

[1] Bold values indicate better performance of a given model (univariate versus bivariate)
[2] U1 = univariate from series $\boldsymbol{y_1}$, B1 = bivariate from series $\boldsymbol{y_1}$, likewise for series $\boldsymbol{y_2}$
[3] Median across the 200 simulations (point estimates are the 50^{th} percentile of the posterior)

The precision for the measurement error (τ_x) and for the autoregressive process (τ_z) exhibited similar trends, with lower coverage when the data were simulated with a larger ρ parameter; the autoregressive variance (τ_z) was less effected (Table 3–3). The univariate and bivariate models performed relatively similarly for retrieving the autoregressive variance but the bivariate strongly outperformed the univariate in coverage and bias for the measurement error precision τ_x. Tables 3–4 and 3–5 show the simulation results for retrieving the ρ parameters from the simulation scenarios. Higher values for ρ led to lower coverage and higher bias between the true value and the posterior median, as the data approached non-stationarity. Across the different scenarios, the bivariate model had more consistent coverage and overall lower bias than estimating ρ from a single univariate model using one of the data series alone (Table 3–4).

Forecasting using the bivariate model also had overall lower forecast error and higher coverage for the predictions than the univariate models (Table 3–5). As expected, the forecast accuracy was lower when the data were simulated with lower ρ parameters, where the

Table 3–3: Summary[1] of simulation results for variances by ρ parameter value

		$\tau_x = 0.1$			$\tau_z = 0.5$		
		Median[3]	**Bias**	**Coverage**	**Median**[3]	**Bias**	**Coverage**
Data	**Model**[2]						
ρ=0.8	U1	0.919	-8.194	0.140	0.489	**0.023**	0.700
	U2	0.834	-7.341	0.160	0.485	0.030	0.700
	B	0.384	**-2.843**	**0.240**	0.745	-0.490	0.540
ρ=0.6	U1	0.241	-1.408	0.360	0.489	**0.022**	0.800
	U2	0.269	-1.690	0.340	0.446	0.108	0.820
	B	0.165	**-0.652**	**0.380**	0.559	-0.117	**0.840**
ρ=0.4	U1	0.121	-0.211	0.700	0.491	0.017	0.900
	U2	0.149	-0.492	0.500	0.458	0.085	0.820
	B	0.111	**-0.113**	**0.760**	0.497	**0.006**	**0.920**
ρ=0.2	U1	0.147	-0.474	0.500	0.445	0.111	0.760
	U2	0.135	-0.350	0.480	0.480	0.039	0.880
	B	0.123	**-0.226**	**0.700**	0.500	**-0.001**	**0.920**

[1] Bold values indicate better performance of a given model (univariate versus bivariate)
[2] U1 = univariate from series $\boldsymbol{y_1}$, U2 = univariate from series $\boldsymbol{y_2}$, B = single parameter estimated from joint model
[3] Median across the 200 simulations (point estimates are the 50^{th} percentile of the posterior)

dependency across time was weaker.

3.8 Case Study

3.8.1 Data

In this section we present results from a case study in British Columbia, Canada. We obtained daily time series data from May 1 through October 31, for each year of 2003 through 2011. We used time series of two different health indicators: series $\boldsymbol{y_1}$ refers to daily pharmaceutical dispensations of a medication used to control acute exacerbations of asthma, and series $\boldsymbol{y_2}$ refers to daily visits to primary care physicians for asthma-related complaints. Both series are correlated with each other and with increased $PM_{2.5}$ concentrations from wildfire smoke over time [44, 17, 66, 37, 35]. The pharmaceutical data include all dispensations of salbutamol sulphate, a prescription drug commonly used to treat the acute symptoms of asthma and other obstructive respiratory diseases. The BC health system requires that every prescription filled be logged into the provincial PharmaNet database in real-time [66].The physician visits data include all asthma-specific outpatient billings to the provincial Medical Services Plan (code 493 in the International

Table 3–4: Summary[1] of marginal posteriors for ρ parameter, $\beta_1 = log1.8$, $\beta_2 = log1.7$

Data	Model[2]	Median[3]	Bias	Coverage
ρ=0.8	U1	0.660	0.175	0.56
	U2	0.667	0.166	0.56
	B	**0.689**	**0.139**	**0.62**
ρ=0.7	U1	0.620	0.115	0.84
	U2	0.605	0.136	0.84
	B	**0.627**	**0.104**	**0.86**
ρ=0.6	U1	0.506	0.156	0.88
	U2	0.513	0.144	**0.92**
	B	**0.516**	**0.140**	0.88
ρ=0.5	U1	0.459	0.082	0.98
	U2	0.457	0.087	0.98
	B	**0.466**	**0.067**	0.98
ρ=0.4	U1	**0.387**	**0.034**	0.9
	U2	0.380	0.051	0.94
	B	0.384	0.040	**0.96**
ρ=0.3	U1	0.332	-0.106	0.98
	U2	**0.308**	**-0.028**	0.98
	B	0.315	-0.049	0.98
ρ=0.2	U1	0.236	-0.178	0.96
	U2	0.247	-0.234	**0.98**
	B	**0.224**	**-0.121**	**0.98**

[1] Bold values indicate better performance of a given model (univariate versus bivariate)
[2] U1 = univariate from series $\boldsymbol{y_1}$, U2 = univariate from series $\boldsymbol{y_2}$, B = single parameter estimated from joint model
[3] Median across the 200 simulations (point estimates are the 50^{th} percentile of the posterior)

Classification of Diseases, 9th revision). Increases in each of these health indicators have been observed during forest fire smoke events [66, 44]. The daily $PM_{2.5}$ values were estimated by the OSSEM model developed and validated for surveillance purposes by the BCCDC, as explained in Chapter 2 [37, 35].

In addition to the $PM_{2.5}$ estimates, we included covariates for daily maximum apparent temperature from Environment Canada, a factor variable indicating the day-of-week, and a binary variable indicating whether the provincial sum of fire radiative power was over the 90th percentile (see Figure 2). Fire radiative power is a satellite measurement of fire intensity, which is proportional to the amount of smoke that active fires generate [37, 97]. A provincial 'fire-day' is defined in this manuscript when the FRP exceeds the 90^{th} percentile for the data series. The years 2003 and 2009 were severe wildfire seasons, where short-term increases in $PM_{2.5}$ exposure

Table 3–5: Summary[1] of marginal posteriors for forecasts, $\beta_1 = log1.8$, $\beta_2 = log1.7$

Data	Model[2]	MAPE	Coverage
ρ=0.8	U1	0.653	0.820
	B1	**0.635**	**0.860**
	U2	0.463	**0.900**
	B2	**0.456**	0.82
ρ=0.7	U1	2.128	0.840
	B1	**0.814**	**0.860**
	U2	0.702	0.820
	B2	**0.653**	**0.860**
ρ=0.6	U1	0.562	0.860
	B1	**0.579**	**0.880**
	U2	0.515	**0.860**
	B2	**0.507**	0.840
ρ=0.5	U1	0.566	0.760
	B1	**0.521**	**0.840**
	U2	0.502	0.780
	B2	**0.451**	**0.800**
ρ=0.4	U1	0.541	0.860
	B1	**0.477**	**0.920**
	U2	0.614	0.780
	B2	**0.573**	**0.86**
ρ=0.3	U1	0.659	0.840
	B1	**0.643**	**0.860**
	U2	**0.535**	0.800
	B2	0.544	**0.880**
ρ=0.2	U1	0.630	0.700
	B1	**0.618**	**0.720**
	U2	0.505	0.700
	B2	**0.465**	**0.800**

[1] Bold values indicate better performance of a given model (univariate versus bivariate)
[2] U1 = univariate from series $\boldsymbol{y_1}$, B1 = bivariate from series $\boldsymbol{y_1}$, likewise for series $\boldsymbol{y_2}$

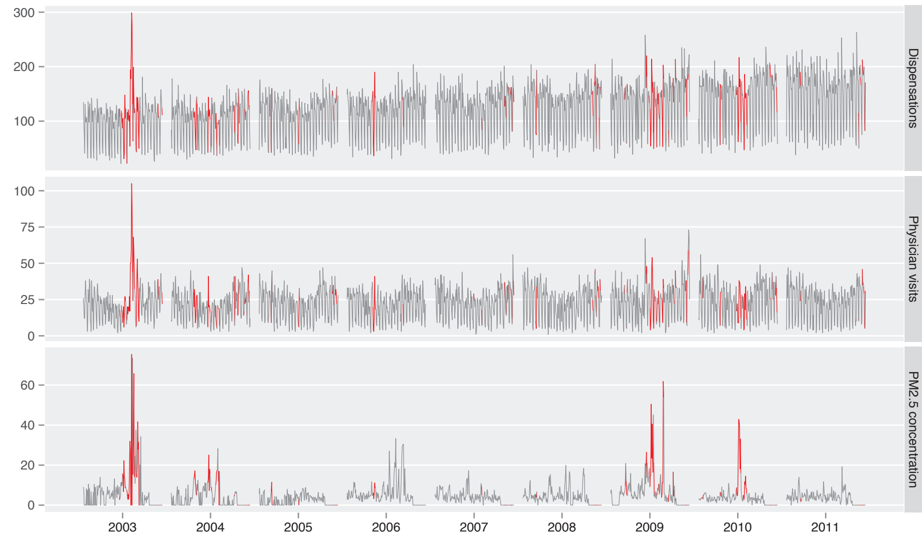

Figure 3–2: Time series for Central Okanagan local health area showing dispensations, physician visits and particulate matter concentrations by fire season (1 May to 1 October), from 2003 to 2011. Red indicates a provincial fire day based on fire severity.

coincided with large spikes in series $\boldsymbol{y_1}$ and $\boldsymbol{y_2}$ health indicators, as shown in Figure 3–2. The days of the time series that appear red are 'fire days' as defined by fire radiative power over the 90^{th} percentile [97].

Because 2003 was at the beginning of the time series, there were insufficient historical data to support model building for that period. We therefore used the 2009 fire season (May 1 through October 31, 2009) as the forecast horizon. Models will be fit using the optimal time window, as described below (see 3.8.2). The series $\boldsymbol{y_1}$ and $\boldsymbol{y_2}$ daily count data were made available by the BCCDC.

3.8.2 Moving window optimization

For the case study, we explored moving windows (w_1 to w_t) ranging from 15 to 500 days. The smallest reasonable window size requires at least one observation per day of the seven day week, given the day-of-week factor variables used. To ensure the identifiability of the model, fifteen days was used as the smallest potential window size. The overall measure of forecast accuracy (MAPE) and precision (width of credible intervals) were compared across $\boldsymbol{y_1}$, $\boldsymbol{y_2}$, and the two models (univariate, bivariate).We used a forecast horizon of 184 days (one wildfire season), y_{h_1} = May 1, 2009 to y_{h_t} = October 31, 2009. Based on these results we selected a single forecast window

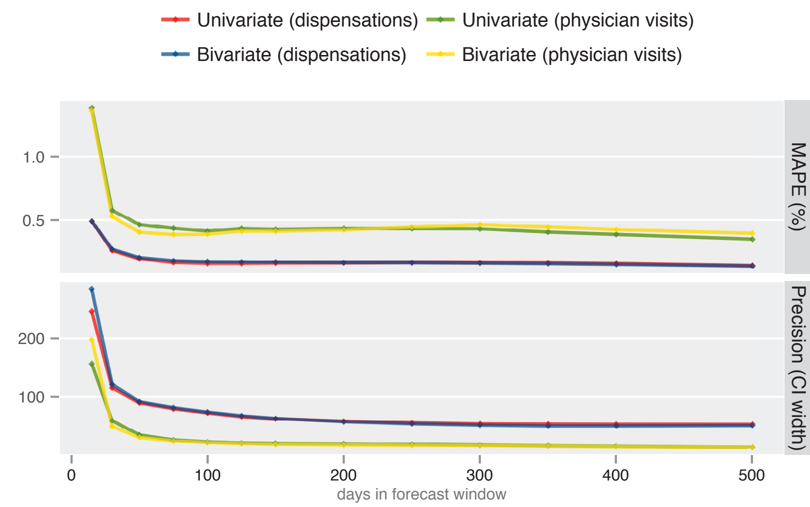

Figure 3–3: Moving window optimization results for case study forecasting, showing accuracy (top) and precision (bottom) by number of days in the forecast window.

size to use for the case study analysis. As shown in Figure 3–3, both forecast accuracy and precision improved when additional days were added to the forecast window until the window size reached 100 days, where it essentially converged. The benefit of including more data was negligible and, in some cases, actually decreased the accuracy. The decrease in accuracy could be due to secular trends not fully modelled in the data. The precision of the forecasts continued to improve with more data, but gains were small after 100 days, and even smaller after 200 days. The results were consistent across y_1 and y_2, and across the two models (univariate and bivariate). A window size of 100 days was therefore used for the analysis of the case study data.

3.8.3 Forecast accuracy performance

The health indicators data contain strong day-of-the-week effects, as can be seen in Figure 3–2. The forecasts for the two models for each health indicator are shown in Figure 3–4. Models are re-fit for each step in the forecast horizon and new parameters were estimated with slightly different data (as the forecast window captures different data). This results in 184 separate models. However, general trends emerged; the day-of-the-week factor variables were very influential in the models, as were the identification of fire-days. Temperature was less influential, possibly because the fire seasons occur primarily during the summer months, and the seasonal trends may be similar for

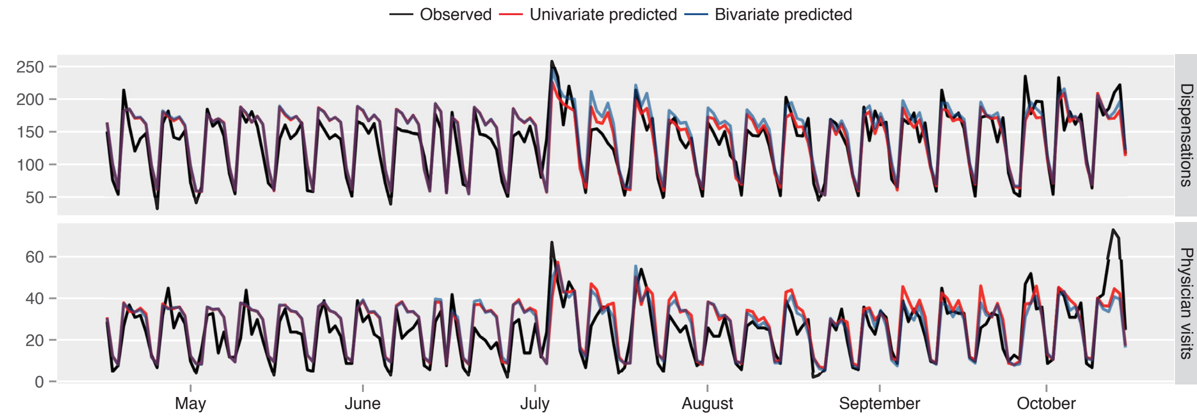

Figure 3–4: Case study health indicator data with forecasts (predicted values) from univariate and bivariate models for study area, May-October 2009 (the forecast horizon).

temperature and particulate matter. Estimates for the ρ parameters were between 0.25 and 0.8, for each univariate model as well as the bivariate models, depending on the dataset (data within the forecast window). In the univariate models, one or both of the estimated β parameters would be negative approximately one third of the time, such that $e^{\beta} < 1$ (suggesting a protective effect of the particulate matter), but the bivariate model estimated more consistent parameters that were close to the known relationship between air pollution and these health outcomes, suggesting better model stability.

The forecast accuracy measured as the mean absolute error was approximately 14% for $\boldsymbol{y_1}$ and approximately 25% for $\boldsymbol{y_2}$. The forecast coverage (percent of the forecasted values where the 95% credible interval included the true value) was approximately 80% for $\boldsymbol{y_1}$ and approximately 63% for $\boldsymbol{y_2}$. The model performance was slightly different for series 1 and series 2; the bivariate model was nearly 3% more accurate (using median MAPE) for series 2, but slightly (0.8%) lower in accuracy for series 1. The forecast coverage was higher in the bivariate model for series 2, but slightly lower for the bivariate model in series 1. However, overall, the bivariate model slightly outperformed the univariate model in forecast accuracy and forecast coverage using the mid-point (mean) for the two series (Table 3–6).

Table 3–6: Summary[1] of case study analysis

Model[2]	**U1**	**U2**	**B1**	**B2**
Median MAPE	**0.132**	0.265	0.140	**0.231**
Percent in forecast CI[3]	**0.8048**	0.614	0.778	**0.658**
	Univariate mid-point		**Bivariate mid-point**	
Median MAPE	0.198		**0.186**	
Percent in forecast CI[3]	0.709		**0.717**	

[1] Bold values indicate better performance of a given model (univariate versus bivariate). The midpoint is the mean (which is also equal to the median) of the two measures
[2] U1 = univariate from series $\boldsymbol{y_1}$, B1 = bivariate from series $\boldsymbol{y_1}$, likewise for series $\boldsymbol{y_2}$
[3] Percentage of the forecasts in the horizon where the 95% forecast credible interval contained the true value

3.9 Discussion

In this paper, we describe a latent process model that jointly models multiple correlated health indicators with a common autoregressive exposure. Most forecasting approaches assume that measurements of the health effects of an exposure on a population, such as morbidity and mortality counts arising from air pollution exposure, are independent from one another. We explored the impact of this assumption by comparing separate univariate models with a bivariate approach that includes two correlated health indicators. We evaluated the performance of the models in a simulation study and explored their application in a case study.

The latent process models proposed have a hierarchical structure common in the applied Bayesian statistics literature [98]. The health indicators of interest are seemingly independent at the first level of the hierarchy, but arise from a common underlying autoregressive process (or processes); this process is measured with (non-differential) error. Similar models have been employed in epidemiological research, but have not been previously described for public health surveillance [52]. Such a model can be implemented using MCMC; however, given the proposed application of public health surveillance, models would need to be efficiently run to aid real-time decision making. The computational costs of the proposed models using MCMC, requiring hours or even days to produce a single forecast, was prohibitive. The computational efficiency of the INLA method as compared to MCMC is well established [55] and it offers a computationally efficient approach to inference, providing forecasts orders of magnitude faster that MCMC. In order to implement the proposed models using the INLA approach, they had to be recast as latent

Gaussian models, and a non-intuitive syntax was needed to specify the joint model of the health indicator(s) and the latent exposure in the package `R-INLA`. The framework we provided could be easily expanded to include more than two health indicators, additional latent exposures measured with error, spatio-temporal correlation, or non-differential measurement error.

The simulation study demonstrated that the bivariate model outperformed the univariate models, with performance measured as bias and coverage of the simulation parameters and forecast accuracy. The results appear robust to minor variations in the magnitude of the effect of the latent process on the indicator. The coverage of the effect sizes (β's), intercepts (α's), and precisions (τ_X and τ_Z) was lower when the simulated data were generated with a larger ρ value, as the data approached non-stationarity. The parameter least affected by the changes in ρ was the autoregressive error term ε_Z, which is reasonable given the structure of the hierarchy. Across the different simulation scenarios, the bivariate models frequently provided higher coverage, lower bias, and more accurate forecasts than the separate univariate models.

We illustrated the use of the proposed models by forecasting healthcare utilization data from BC, Canada. Summarizing results for a forecast horizon of one fire season (184 days in 2009), the bivariate model provided overall more accurate predictions (as measured by MAPE), and provided credible intervals that more frequently contained the true value. The bivariate model also resulted in more stable effect estimates. The β parameter estimates were negative in approximately one third of the univariate models, suggesting that air pollution had a protective effect, which is known not to be true for these data. This result was observed rarely for the bivariate model. The bivariate model is also more computationally efficient; while it took longer to run than a single univariate model, it was faster than running the two univariate models (one for each indicator) separately.

We explored a first-order autoregressive latent process, based on exploratory data analysis of the health and air pollution data from BC, which frequently exhibited this trend. Future work could extend this approach for data that exhibit higher-order autoregressive trends, and also explore misspecification. For example, simulation studies could explore the impact of generating a second-order autoregressive process in the latent field, and modelling it as a first-order. Future developments of this approach could also incorporate data from multiple regions and, by incorporating spatio-temporal structure into the models, allow strength to be borrowed across regions [99]. This strategy may aid in obtaining both more accurate forecasts but also potentially in reducing the uncertainty associated with those forecasts.

Simultaneously modelling multiple health indicators with common exposures can provide accurate and timely information for environmental public health surveillance. The findings here are relevant for related fields (heat surveillance, other environmental exposures). Future applied work could directly evaluate the public health impact of improved accuracy by implementing the proposed models in a real-time surveillance system.

Acknowledgements

The authors would like to thank Finn Lindgren for his help and valuable advice in implementing the models. The first author is funded by the Canadian Institutes of Health Research (CIHR). The second author worked on this project whilst he was an International Visiting Research Scholar funded by the Peter Wall Institute for Advanced Studies at the University of British Columbia. The third author has contributed to this work with support from Health Canada. The final author is supported by a CIHR Chair in Applied Public Health Research. Surveillance data for the case study were made available to the BCCDC by the BC Ministry of Health.

4 Improving forecast precision in environmental public health surveillance: a spatio-temporal approach

4.1 Preamble

In the previous chapter, I presented a statistical modelling framework for multiple time series of health indicators and their common latent environmental exposure and explored at the performance of the proposed model through a simulation study and a case study with one administrative area in BC.

In this chapter, I apply the proposed model in all the administrative health areas across the province. The prediction variance can sometimes be unstable when forecasting health effects in communities with very small populations. Wildfire smoke can affect large geographic areas and many of the communities may not have the population size of an urban centre, but the public health concern is significant given that entire populations can be exposed to very poor air quality. These less populated regions can be aggregated for the sake of a forecasting approach to surveillance, but information will be lost in the process.

To facilitate the use of the proposed framework in these circumstances, the temporal model was expanded to include spatio-temporal smoothing. Smoothing between neighbouring local health areas allows borrowing of strength between regions, reducing the prediction variance in exchange for an increase in bias. In the following chapter, I explore the performance of such a bias-variance tradeoff for the sake of wildfire smoke surveillance.

This manuscript will be submitted to *Spatial and Spatio-Temporal Epidemiology*.

4.2 Title page and footnotes

Title: Improving forecast precision in environmental public health surveillance: a spatio-temporal approach

Short title: Spatio-temporal approach to environmental health surveillance

Authors: Kathryn T Morrison[1], Gavin Shaddick[2], Matthew L. Thomas[2], Sarah B. Henderson[3,4], David L. Buckeridge[1]

Affiliations:

[1] Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, Quebec, Canada

[2] Department of Mathematical Sciences, University of Bath, Bath, United Kingdom

[3] Environmental Health Services, British Columbia Centre for Disease Control

[4] School of Population and Public Health, University of British Columbia

Corresponding author:

Kathryn Morrison

1140 Pine Avenue West

Montreal, Quebec, Canada

H3A 1A3

E-mail: kt.morrison@mail.mcgill.ca

4.3 Abstract

Introduction: Wildfires are an important source of episodes of high concentrations of air pollution in many parts of the world. There is clear epidemiological evidence linking short-term increases in particulate air pollution to adverse health effects, notably on the respiratory system. In general, public health officials in smoke affected areas lack the tools required to perform surveillance and assess the health impacts of smoke events. Communities affected by wildfire smoke are heterogeneous in population size and makeup, resulting in low statistical power in some less populated regions.

Methods: Recently, a latent process approach has been proposed for public health surveillance of wildfire smoke, monitoring multiple health indicators simultaneously. However, that approach considered information from a single region and here we extend this model to use information from multiple regions, modelling both spatial and temporal dependence, within a Bayesian framework. This allows smoothing between regions, exploiting spatial dependence to stabilize the prediction variance in regions with smaller health counts. Inference using MCMC was found to be computationally infeasible where multiple models need to be fit, in this case on a daily basis, and we describe implementation using integrated nested Laplace approximations (INLA).

Results: Our results show that for the regions with smaller populations, the spatio-temporal approach smooths the forecasts, substantially improving the prediction variance at the cost of a modest reduction in accuracy.

Conclusion: Wildfire smoke can affect communities across large geographic areas. Administrative data for health monitoring are usually disseminated within geographic units. To use a forecasting approach to environmental public health surveillance and to avoid aggregating smaller regions, spatial smoothing may be beneficial to avoid issues of unstable prediction variance.

Keywords: Bayesian analysis, environmental epidemiology, hierarchical models, spatial statistics

4.4 Introduction

In many areas, wildfires are an important source of short term increases in particulate air pollution, which can reach very high concentrations. Wildfire seasons have become longer and more severe in recent decades, due to changes in land use and climate [22, 9]. Wildfires smoke is a complex mixture of gases and solids, including fine particulate matter ($PM_{2.5}$), which is known to be harmful to human respiratory health [100, 66, 17]. Smoke from wildfires can travel far from the original fire, potentially exposing millions of individuals in large urban centres hundreds of kilometres away [39]. North America has been severely affected, including British Columbia (BC), the westernmost Canadian province. A severe wildfire season in 2015 resulted in the population of greater Vancouver being exposed to extremely smoky conditions and record levels of air pollution in early July [101].

There is strong epidemiologic evidence linking increased rates of morbidity and mortality with exposure to acute $PM_{2.5}$ [102, 103, 104]. In spite of these well-established links, public health officials throughout smoke affected areas lack surveillance tools to assist in decision making during smoke events. Surveillance of acute environmental exposures is an emerging research area, as smoke pollution and extreme heat are becoming increasingly important public health events. Researchers and scientists working in public health have called for improved surveillance methods and tools [105, 106, 107, 108, 109].

Most of the methods and models commonly used in public health surveillance were developed for infectious disease [33]. Environmental public health surveillance shares many objectives with infectious disease surveillance, such as generating information to guide public health action, but seeks to detect the real-time *health impact of an exposure*, an objective that presents some different methodological challenges [6, 34]. We have identified six major challenges for timely environmental public health surveillance:

1. *Exposure estimation:* Wildfire smoke is a complex mixture of potentially harmful pollutants, but $PM_{2.5}$ has been studied most frequently in relation to respiratory health. Measuring air pollution across large geographic areas is logistically and technically challenging, and assumptions must be made to extrapolate large-scale measurements to estimate likely population-level exposure in

an individual or a community. Measurement error is unavoidable with the currently available methods.

2. *Health indicators:* $PM_{2.5}$, can have a range of different health impacts on populations, such that any single health indicator provides only partial information with differing timeliness, sensitivity, and data artifacts (e.g., day of the week effects for healthcare utilization data). While some epidemiologic research has used multivariate approaches to improve parameter estimation [52], most statistical models in epidemiologic studies of air pollution and respiratory health view each health measure as independent.

3. *Statistical modelling:* Currently, air pollution data and health surveillance data are monitored separately in public health settings. Modelling the relationship between the exposure and the health indicator(s) could provide additional information, such as the ability to forecast likely health indicators under different exposure scenarios.

4. *Implementation:* A suitable modelling approach must be flexible enough to address the previously stated challenges but be compatible with an efficiency that is sufficient to thoroughly evaluate via simulation and retrospective analyses. Ultimately, the modelling approach must also be appropriate in a public health setting for real-time surveillance.

5. *Heterogeneous populations:* Many communities most severely affected by wildfire smoke have small or moderate population sizes, making statistical predictions less stable. Surveillance data are usually monitored separately for each community, while spatial dependence could be used to reduce prediction variance.

6. *Assessment:* There are no established thresholds that define when a population will experience a public health effect from smoke as measured by $PM_{2.5}$. An integrated surveillance framework should be developed, evaluated, and calibrated to meet the needs of public health and support decision making.

Developing better estimates of population level exposure to air pollution and wildfire smoke are active research areas [110, 111]. As explained in previous chapteres, the OSSEM model developed by researchers at the BCCDC) estimates $PM_{2.5}$ for each administrative health region across the province [37, 35]. The BCCDC has also identified two health indicators that are strongly associated with acute increases in $PM_{2.5}$: (1) asthma-related physician visits, and (2) asthma-related pharmaceutical dispensations [66, 44]. These data are currently collected and monitored for wildfire smoke surveillance in the BC Asthma Monitoring System (BCAMS) [42].

Here we focus on short-term forecasting of two health indicators associated with increases in air pollution due to wildfire smoke in all BC LHAs. Short-term forecasting can be used within a comprehensive surveillance framework to detect health impacts of an exposure and to estimate public health effects under different exposure scenarios, guiding public health action in real-time. We define two periods that are fundamental to producing and evaluating forecasts: (i) the forecast window is used to build the models and produce the incremental forecasts in each region, and (ii) the forecast horizon is the period over which we obtain one-day-ahead forecasts and summarize prediction performance (Figure 4–1).

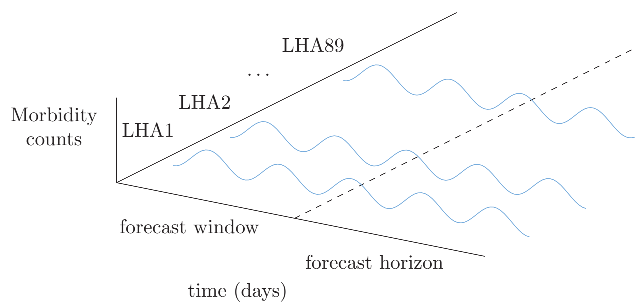

Figure 4–1: Schematic showing forecast windows and forecast horizons for each region

This work builds upon previous research addressing challenges 3 and 4 [112]. In this manuscript, we aim to address challenges 4 and 5 by expanding a previously developed temporal latent process model to a spatio-temporal approach. While the previous model performed well in a moderately populated region (approximately 200,000 residents), many fire-affected communities have considerably smaller populations (e.g., fewer than 50,000 residents). We hypothesize that by adding spatial smoothing to the temporal approach, we can reduce the prediction variance in areas with smaller populations. We also address some of the computational challenges posed by this more complex model. This work is an important step towards assessing the utility of these forecasting models within an integrated surveillance system (challenge 6).

In Section 4.8.1 we describe the study area in BC, and the data from the wildfire season of 2015. In Section 4.6, we briefly summarize a previously proposed bivariate latent process model [112] for forecasting health indicators in a single region based on estimated levels of air pollution. In that study, we summarized results for one LHA in BC for the wildfire season of 2009. The bivariate approach outperformed the univariate, where each indicator was modelled independently. The bivariate approach provided more accurate predictions, had credible intervals that more frequently contained the true value, and more consistently estimated realistic effect estimates for the relationship between air pollution and respiratory health indicators. The remainder of the Methods section describes how this bivariate model can be expanded to incorporate both spatial and temporal dependence. Section 4.7 briefly summarizes the implementation of this model using integrated nested Laplace approximations (INLA) [113] and the estimation of prediction intervals. Section 4.8 presents a comparison between the purely temporal specification for each region separately to a combined spatio-temporal approach. Forecast performance is assessed in terms of accuracy and precision. In Section 4.9 we discuss the benefits and limitations of the proposed modelling approach and suggest next steps for research to move towards an integrated surveillance framework for wildfire smoke surveillance.

4.5 Data

British Columbia is a heavily forested province in Western Canada, with a population of nearly 5 million residents [114]. The province is divided into 89 health regions known as local health areas (LHAs) for the purpose of healthcare administration and research. The average number, size, and cost of wildfires in BC have risen dramatically in the last decade. Figure 4–2 shows the locations of more severe provincial fires of 2015 based on fire radiative power (FRP), which is a value derived from satellite imagery to estimate fire intensity, and is proportional to the amount of smoke that active wildfires generate [37, 97, 115].

The forecast window includes the wildfire seasons (May to September) from May 1, 2013 to June 30, 2015. The amount of data included in the forecast window can be determined in different ways, from including all data available to choices bases on exploratory data analysis. In our previous work, we optimized the length of the forecast window using time series cross-validation [112]. Including more data in the forecast window will decrease the forecast precision, although the benefit will likely become negligible after some threshold. The accuracy will increase initially and then may continue to increase, remain unchanged, or even decrease in some cases.

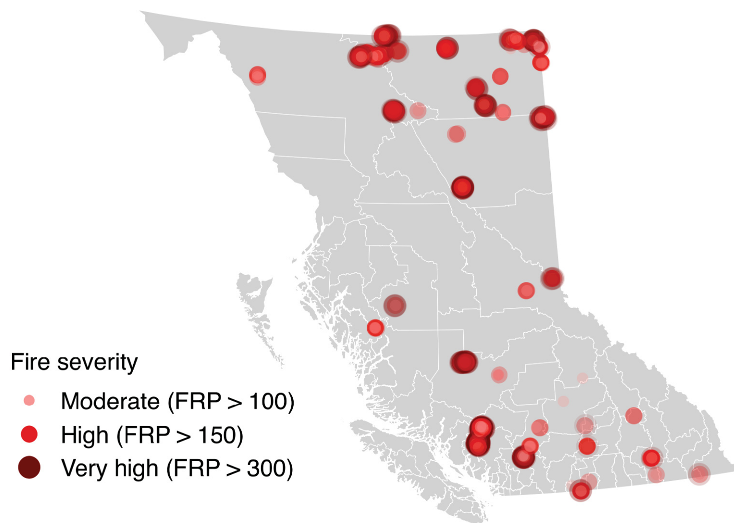

Figure 4–2: Location and severity of wildfires across British Columbia Local Health Areas, May - September 2015.

Using less data will ease the computational burden. We found previously that using a forecast window of 100 days was optimal for a single, moderately populated LHA in BC. Here, we add additional data to be conservative rather than attempting to optimize across all regions.

The forecast horizon is the 31 days of July 2015, across which the prediction performance will be evaluated. We summarize the forecast accuracy with mean absolute error (MAE) $\frac{1}{N}\sum_{i=1}^{N}|y_i - \hat{y}_i|$, where N = days in forecast horizon. We summarize the forecast variance with mean credible interval width (MCIW) $\frac{1}{N}\sum_{i=1}^{N}(\hat{y_i}^{0.975} - \hat{y_i}^{0.025})$. We use MAE, an absolute rather than relative measure of forecast accuracy, as we are interested in the comparison between the temporal and spatio-temporal models, rather than between regions with data on different scales. We defined the *forecast coverage* as the proportion of true observed values a given (credible or prediction) interval contains across a forecast horizon of N days.

4.5.1 Health surveillance data

Here we consider two health indicators identified and monitored by the BCCDC for the purpose of air pollution surveillance [42]. The first is daily dispensations of the prescription drug salbutamol sulphate, which is used to treat the acute symptoms of asthma and other chronic obstructive lung diseases [116]. BC provides a provincial drug plan that requires by law every prescription filled or refilled to be logged by the provincial PharmaNet database, regardless of the recipient or the payer, and these data are made available to the BCCDC for research and surveillance purposes [44, 117].

The second health indicator dataset is daily asthma-related visits to primary care physicians, such as walk-in clinics and family practices. Nearly all residents of BC are registered in the provincial Medical Services Plan (MSP) [118]. Primary care physician visits are billed to the provincial MSP, and the reason for the visit is coded using International Disease Classification codes (ICD-9). Asthma-related codes (group 493) are identified by the BCCDC and used for the purposes of wildfire smoke surveillance.

4.5.2 $PM_{2.5}$ exposure

The BCCDC has developed the OSSEM predictive model to estimate average population exposure to smoke-related $PM_{2.5}$ in each LHA in BC for the purpose of public health surveillance [37, 35]. Their model integrates air pollution monitor data, climate and weather data, and satellite data on smoke and fires to produce a population-weighted prediction for each LHA in BC on a daily basis. They evaluated their model by spatially predicting $PM_{2.5}$ concentration levels in areas with air quality monitors, and average accuracy was over 80%.

4.5.3 Additional data

In addition to the health surveillance and air pollution data, we use values of daily maximum apparent temperature [119], the annual population estimated for each LHA [114], and an indicator variable for provincial 'fire days' as based on FRP [37]. The healthcare utilization data exhibit strong day-of-week variation, likely due to differing levels in access to care by day by week, and so a factor variable for each day of the week, with statutory holidays coded as Sundays, is also used.

4.6 Methods

We propose a three-level hierarchical spatio-temporal model for environmental exposure data modelled as a latent process and multiple correlated health indicators. The health indicator i, for time t, in region s, is modelled as a Poisson random variable, with rate μ_{ist} as described previously in Chapter 3:

$$\begin{aligned}
Y_{ist}|\mu_{ist} &\sim Pois(\mu_{ist}) \\
log\mu_{ist} &= \alpha_i + \beta_i Z_{st} + \sum_{j=1}^{J} \gamma_{ji} W_{jist} + V_{is} + U_{is} \\
X_{ts} &= Z_{ts} + \varepsilon_{X_{ts}} \\
Z_{ts} &= \rho_1 Z_{t-1,s} + \rho_2 Z_{t-2,s} + ... + \rho_p Z_{t-p,s} + \varepsilon_{Z_t}
\end{aligned} \tag{4.1}$$

where X_t is the measured $PM_{2.5}$ of the latent exposure Z_t; $\boldsymbol{\alpha}$ are the intercept terms and $\boldsymbol{\beta}$ are effects of the daily $PM_{2.5}$, allowing a different relationship with each health indicator. The measurement error on the air pollution, ε_{X_t}, and the random variation on the autoregressive latent process, ε_{Z_t}, are each normally distributed with zero mean; ε_{Z_t} is also constrained to sum to zero to allow for identifiability of each parameter. The variances for ε_{X_t} and ε_{Z_t} (τ_{X_t} and τ_{Z_t}, respectively) are defined using the marginal precision, see [88]. The effect parameters of other covariates **W**, which can be common or different for each indicator, are given by $\boldsymbol{\gamma}$. The two random effects, V and U, are uncorrelated and spatially correlated random effects, respectively.

The second level of the model describes the latent process, $\boldsymbol{Z}$, the true underlying concentration of air pollution, which is modelled an autoregressive process of order p. The underlying air pollution concentration is common across all health indicators. The third and final level specifies the prior distributions for each parameter.

A temporal version of this model is given in Chapter 3, where a single region, s, is considered [112]. The temporal specification can be replicated for each region in a larger geographic study area, but models in regions with small counts of health indicators may be vulnerable to low power. The smaller regions could be aggregated, but information would be lost. Here we use the temporal model in each region to independently predict each health indicator on each day in the forecast horizon, and we also implement a spatio-temporal approach and compare the performance of the two specifications. The spatio-temporal model estimates the first level model parameters across the regions, smoothing towards the mean and allowing these smaller areas to borrow strength from their neighbours, therefore stabilizing the prediction variance.

The spatio-temporal model has two additional random effects at the second level of the hierarchy, $\boldsymbol{U}$, and $\boldsymbol{V}$, each assigned multivariate Gaussian priors. The random effect $\boldsymbol{V}$ represents uncorrelated random error, $\boldsymbol{V} \sim \mathcal{N}(\mathbf{0}, \tau_v^{-1}\boldsymbol{I})$. The random effect $\boldsymbol{U}$ represents spatially correlated random error, $\boldsymbol{U} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma})$, $\boldsymbol{\Sigma}^{-1} = \boldsymbol{Q} = \boldsymbol{D}^{-1}(\boldsymbol{I} - \boldsymbol{W})$, where $\boldsymbol{D} = diag(m_s^{-1})$, $W_{ij} = m_s^{-1}$ if regions i and j are neigbhours and $W_{ij} = 0$ otherwise; m_s is the number of neighbouring regions for region s using first-order adjacency.

The precision matrix of the intrinsic CAR model $\boldsymbol{Q}$ is of rank $(n-1)$ and therefore the joint distribution of the random effects does not exist [120]. However, the conditionals (with conditional variance τ_u^{-1}) are proper distributions, and can be estimated as:

$$u_{is}|\boldsymbol{u}_{i,\text{-}\boldsymbol{s}}, \text{-}\boldsymbol{s} \in \delta_s \qquad \sim \mathcal{N}(\bar{u}_{i\delta_s}, \frac{1}{m_s\tau_u}) \tag{4.2}$$

where $\bar{u}_{i\delta_s}$ is the mean of the random effect s in its neighbourhood region δ_s. The conditional precision τ_u is proportional to the strength of the spatial association in a given region [121] (Figure 4–3).

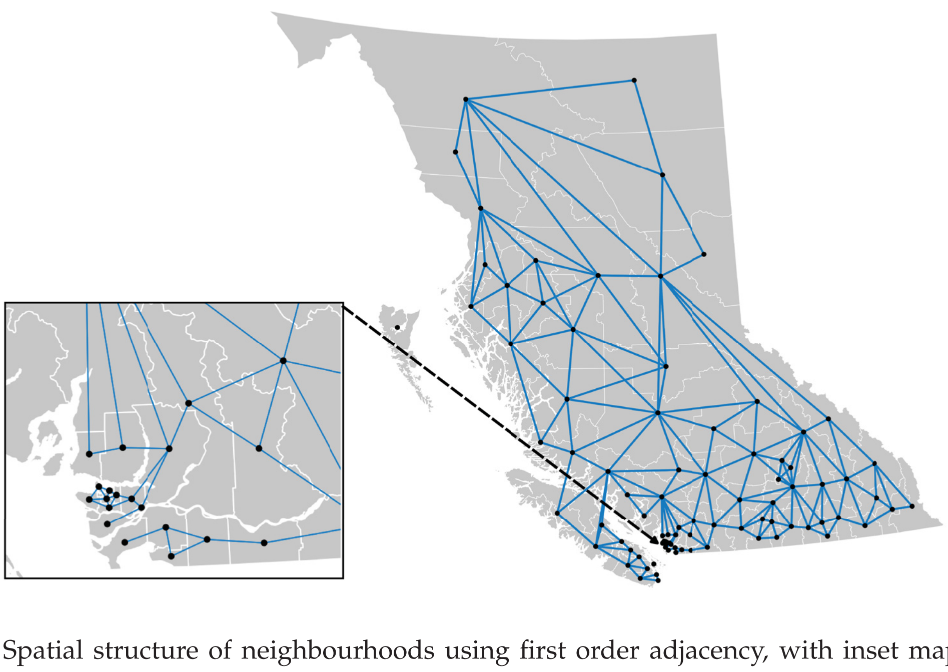

Figure 4–3: Spatial structure of neighbourhoods using first order adjacency, with inset map for greater Vancouver

Priors for the other parameters in the model are as follows, $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma} \sim \mathcal{N}(\mathbf{0}, \tau\boldsymbol{I})$, where $\tau = 0.001$ [88]. INLA defines the prior for the ρ parameter on the logit scale [88]. As in [112], the prior is changed here to provide a unimodal distribution: $log\frac{(1+\rho_p)}{(1-\rho_p)} \sim \beta(5,5)$.

4.7 Inference and computation

To address many of the methodological challenges in environmental public health surveillance, the complexity of the models proposed here require considerable computational costs. As these models are intended to be integrated into daily use in a public health setting, they must be able to produce forecasts on a daily basis. We found that for the temporal version of this model, simulation-based approaches such as Markov chain Monte Carlo (MCMC) were impractical given long convergence times [112]. Therefore, for the spatio-temporal extension, we again use integrated nested Laplace approximations (INLA) for computationally efficient approximate Bayesian inference [53].

The INLA algorithm is described in detail in the seminal paper [53] and outline briefly in Chapter 3. Essentially, many hierarchical models can be expressed as latent Gaussian models (LGMs) and used within the INLA framework. LGMs share a common structure:

$$\begin{aligned}
&\text{Likelihood: } \boldsymbol{Y}|\boldsymbol{Z},\boldsymbol{\theta_2} \sim \prod p(Y_i|\boldsymbol{z},\boldsymbol{\theta_2}) \\
&\text{Latent field: } \boldsymbol{Z}|\boldsymbol{\theta_1} \sim p(\boldsymbol{z}|\boldsymbol{\theta_1}) = \mathcal{N}(\boldsymbol{0},\boldsymbol{\Sigma}) \\
&\text{Hyperpriors: } \boldsymbol{\theta} = [\boldsymbol{\theta_1},\boldsymbol{\theta_2}]^T \sim p(\boldsymbol{\theta})
\end{aligned} \tag{4.3}$$

where $\boldsymbol{Y}$ are the observed data; $\boldsymbol{Z}$ is the latent field, which is the joint distribution of the linear predictor and parameters within (such as the regression coefficients); $\boldsymbol{\theta_2}$ are the hyperparameters of the likelihood; $\boldsymbol{\theta_1}$ are the hyperparameters for the latent field. The latent field is assumed to be multivariate Gaussian with a conditional independence structure that results in a sparse precision matrix; this structure works naturally for autoregressive models such as those common in spatio-temporal statistics [58, 121]. Estimates of the posterior marginals for the latent field can be derived from a Bayesian hierarchical model as:

$$p(z_j|\boldsymbol{y}) = \int p(z_j,\boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta} = \int p(z_j|\boldsymbol{\theta},\boldsymbol{y})p(\boldsymbol{\theta}|\boldsymbol{y})d\boldsymbol{\theta} \tag{4.4}$$

which are approximated by INLA as [95]:

$$\sum_k \tilde{p}(z_i|\boldsymbol{\theta}_k, \boldsymbol{y})\tilde{p}(\boldsymbol{\theta}_k|\boldsymbol{y})\Delta_k \tag{4.5}$$

where $\tilde{p}(\cdot)$ refers to an approximation. The first term in (4.5) refers to a Laplace approximation, while the second term depends on the implementation of the INLA program (a full Laplace approximation, a Gaussian approximation, or a simplified Laplace approximation); see [54]. The simplified Laplace approximation is the default and was used for all final analyses in this manuscript. For a detailed explanation of the INLA approximation algorithm, see [53]. In the analyses we use the implementation of INLA within the R package `R-INLA`, see [88].

Although INLA provides an efficient approach to performing approximate Bayesian inference, there may still be computational issues when repeatedly fitting large datasets using complex models with many parameters, as is the case here. The models described in this paper were implemented using a high performance computing system, making use of a number of high memory nodes (including two nodes with 512GB of memory). Even so, an initial approach using the default method for selecting initial values for the optimization for each day independently proved to be computationally demanding, as the Newton-Raphson optimizer had to restart to produce the forecast each day when a new model was fit. Using the previous days posterior values to initialize the optimization for the following forecast day provided a solution for days where convergence was not achieved and, overall, allowing learning between days resulted in marked improvements in efficiency.

4.7.1 Prediction intervals

A confidence interval around the expectation $E(Y|X)$ is straightforward to obtain for generalized linear models, with the assumption that the linear predictor is asymptotically normal [122]. However, any CI around $E(Y|X)$ will only capture the variability in the mean of the linear predictor, as used in model evaluations in Chapter 3, and not the additional variability in future individual observations.

Prediction intervals are straightforward to compute with linear models but may be more complex with non-Gaussian indicators, such as Poisson counts used here. In such cases, empirical prediction intervals may be constructed using bootstrapping or simulation based approaches [123].

In a Bayesian setting, empirical prediction intervals can be constructed for a new observation using samples from the posterior predictive distribution. Such samples are readily available using MCMC as well as in R-INLA. The following algorithm shows the process used here to construct prediction intervals around the forecasted health indicators on day t for given fit R-INLA model M using k posterior samples. Here, we used k=1000 to generate prediction intervals.

Algorithm: Constructing Bayesian prediction intervals for a Poisson hierarchical model

for i =1 to k

Generate a sample value P_{S_i} from marginal posterior predictive distribution of M

Generate random variable: $P_{rv_i} \sim \texttt{Poisson}(P_{S_i})$

Construct bounds from empirical distribution: (1-α)

PI $= P_{rv}^{(\alpha/2)}, P_{rv}^{(1-\alpha/2)}$

4.8 Results

4.8.1 Data summary

The summer of 2015 was a severe wildfire season in BC, with over 280,000 hectares of wildfires burning, and more days of smoke exposure in populated regions of the province than previously on record [124]. Within the month of July and across the 89 LHAs, there were 180 region-days of high air pollution exposure from wildfire smoke (PM$_{2.5} > 25\mu g/m^3$). Because some highly populated regions were affected, including greater Vancouver, this resulted in 13,869,985 person-days of high PM exposure in the province. In particular, large fires burned in the western part of the province from July 5 to 10, resulting in elevated air pollution concentrations in many urban communities (Figure 4–3).

Daily average counts per LHA for salbutamol dispensations and asthma-related physician visits for the month of July 2015 varied by population size. Dispensations occur at much greater frequency than physician visits (Figure 4–4), resulting in larger average daily counts per region. Across the entire province, mean dispensations per region were 15 per day and mean physician visits were six per day. During periods of high PM exposure (PM$_{2.5} > 25\mu g/m^3$), the means were

35 per day and 15 per day, respectively.

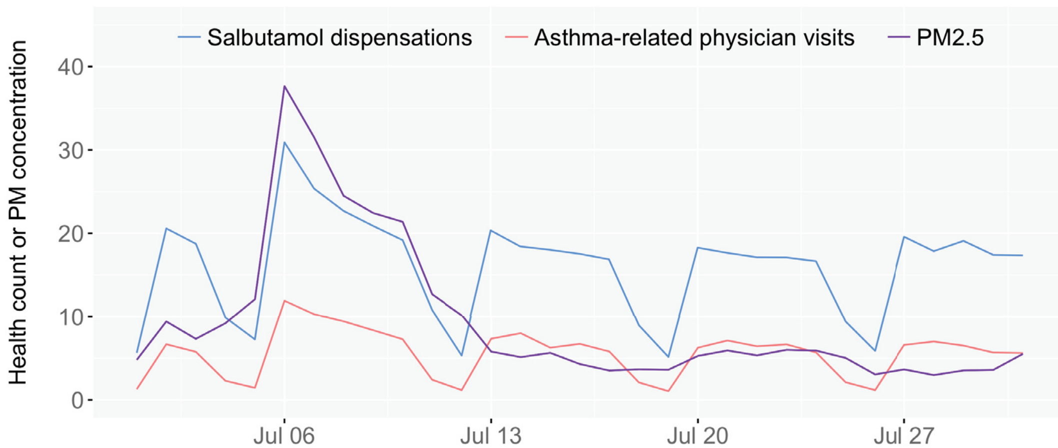

Figure 4–4: Provincial mean daily salbutamol dispensations, asthma-related physician visits, and estimated $PM_{2.5}$, July 2015. Note that the y-axis refers to health indicator counts for dispensations and physician visits, while it refers to concentration in $\mu g/m^3$ for $PM_{2.5}$.

4.8.2 Forecast accuracy and precision

In this study, we compared the forecast performance of a temporal latent process model [112] with a spatio-temporal extension. We evaluated performance empirically using MAE and MCIW for a 31 day forecast horizon during the wildfire season of 2015.

Overall, the accuracy between the two models was similar, though slightly lower for the spatio-temporal model. This reduction in accuracy is expected, as the spatial smoothing used here is a form of bias-variance tradeoff. The forecast variance was substantially reduced in the spatio-temporal models on average, generally at the expense of a modest reduction in accuracy (Table 4–1).

We place a special focus on the regions with the smallest populations, as these regions tend to have much smaller daily counts and less precise prediction intervals. In these regions, and especially for the physician visits data, the accuracy is similar between the temporal and spatio-temporal models, but the forecast variance is drastically improved with the spatio-temporal model (Figure 4–6). Because more than half of the LHAs in the province have fewer than 20,000 residents, model performance is of particular importance in these regions.

Table 4–1: Forecast horizon accuracy (MAE) and precision (MCIW) by region/day subgroup

Indicator	Dispensations		Physician visits		Notes
Model	Temporal	Spatio-temporal	Temporal	Spatio-temporal	
All regions and days					Full horizon
MAE	121.02	286.13	59.33	62.91	
MCIW	406.44	25.44	117.59	0.47	
Region-days where $PM_{2.5}$ >25 $\mu g/m^3$					180 region-days
MAE	1491.38	2830.03	873.23	906.34	
MCIW	394.57	36.19	3.99	0.76	
Regions with population >200,000					4 regions
MAE	1667.38	5277.77	1049.92	1148.85	
MCIW	50.73	101.83	7.09	1.63	
Regions with population <20,000					45 regions
MAE	2197.49	3432.75	978.92	1001.58	
MCIW	781.33	5.55	229.03	0.18	
Regions with population <10,000					29 regions
MAE	1073.93	1522.99	453.72	467.6	
MCIW	1209.55	3.51	354.58	0.13	
Regions with population <5,000					15 regions
MAE	352.8	472.63	151.59	158.56	
MCIW	2336.08	1.79	684.82	0.09	

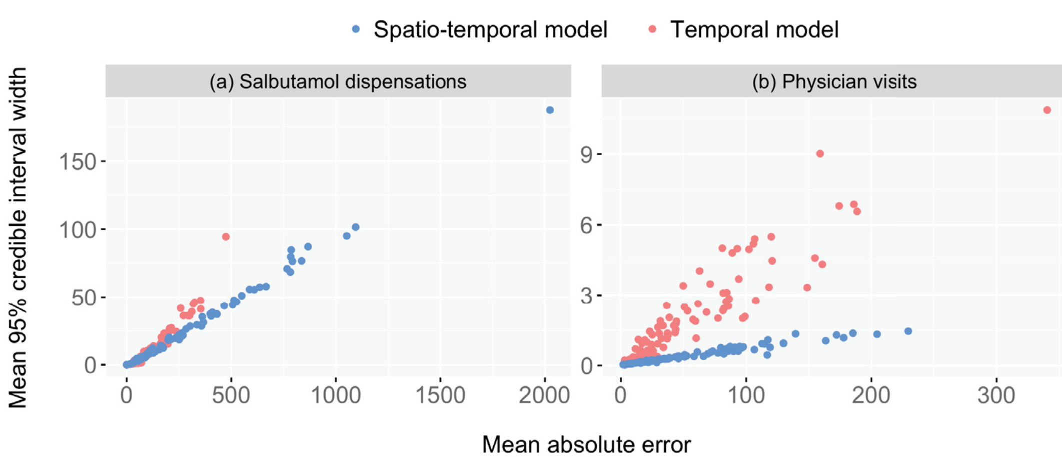

Figure 4–5: Mean forecast 95% credible interval width across horizon by mean absolute error per LHA in BC, for (a) salbutamol dispensations and (b) asthma-related physician visits. Note that some extreme outliers for the 95% credible interval width for the temporal model have been removed for both indicators.

During periods of elevated air pollution it is particularly important to forecast accurately and precisely because public health may intervene to reduce exposure during these smokey periods. The performance of the models in the region-days of high exposure will be largely driven by the demographics of the regions that happen to be exposed to a given wildfire smoke event.

In July of 2015, highly populated regions in the lower mainland were heavily exposed, therefore performance of the spatio-temporal model in the high exposure regions mirrored the performance in the highly populated regions.

4.8.3 Prediction interval coverage

The variance for a predicted mean indicator is proportional to the variance for a new observation; the latter is larger as it accounts for additional variation. Using credible intervals around the mean expected values provided average forecast coverage of 47.0% for predicted dispensations and 5.5% for predicted physician visits, using the spatio-temporal model. Using the empirical prediction interval method described in Section 4.7.1, the average forecast coverage was 75.5% for dispensations, and 92.5% for physician visits. Figure 4–6 shows the time series of true and forecasted physician visits for the July 2015 forecast horizon in a moderately populated region of the province (Ladysmith, population approximately 18,500), including the credible and prediction intervals. The forecast coverage is nearly zero using the credible interval of the mean, and 90.3% using the prediction interval.

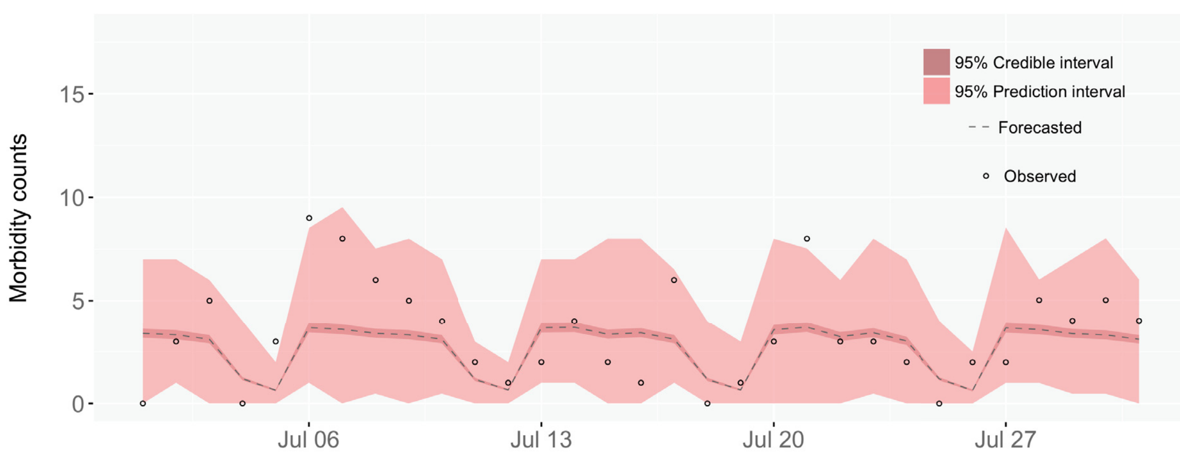

Figure 4–6: True and forecasted values for physician visits with credible and prediction intervals a single LHA (Ladysmith). Forecasts are the median of the posterior predictive distribution

4.9 Discussion

In this manuscript we extended our previously proposed latent process surveillance model [112] to include spatial smoothing. This extension was proposed to address the large prediction variance in regions with small counts of health indicators under surveillance.

Our hierarchical approach builds a temporal prediction model based on historic data by modelling two respiratory health indicators as Poisson random variables jointly by regressing them against a common exposure. The exposure is then modelled as a latent process at the second level of the hierarchy, to account for the measurement error inherent in estimating exposure to $PM_{2.5}$. Spatial smoothing is incorporated to allow regions with smaller populations and/or smaller health indicator(s) counts to borrow strength from spatial neighbours, which are likely to share common exposure given the spatial distribution of air pollution, as well as similar demographics.

Daily counts of health indicators will be smaller in less populated regions, which is common for many communities affected by wildfire smoke. Additionally, small counts are lower on average for more severe health effects. The health impacts of wildfire smoke range from mild symptoms through to death [8].

Each health indicator contains information on the impact of the environmental exposure and the data used to measure each health effect may have different characteristics, such as timeliness, day-of-the-week effects, or types of censoring. In general, monitoring multiple indicators will provide public health with more information, and useful datasets should not be excluded from surveillance due to low expected counts and the associated statistical challenges.

We found that spatial smoothing was beneficial for predicting the physician visits data, which have a lower incidence than the medication dispensations data. The substantial reduction in prediction variance for the spatio-temporal model results in only a modest decrease in the overall accuracy from the smoothing. The spatial smoothing was also beneficial for the dispensations data in regions with smaller populations, particularly in regions with fewer than 5,000 individuals. In regions with large populations and high mean counts, forecasts of medication dispensations were less accurate when smoothed. In [112] we found that the temporal model performed well for a moderately populated region, but more than half of the regions in BC have small populations (less than 20,000 residents). This borrowing of strength between regions makes it possible to accurately forecast indicators for surveillance throughout the province, when previously it was not feasible to do so.

Prediction intervals should capture the true observed health indicators with probability close to the given confidence level (e.g., 95%). Using the credible interval around the mean expected count neglects the additional variability in predicting individual observations. However, estimating prediction intervals empirically can be computationally intensive [125]. The empirical prediction intervals obtained here fit well into the inferential framework using INLA, provide reasonable coverage even in regions with sparse data, and maintain the computational feasibility of the approach. The availability of these prediction intervals increases the utility of this forecasting approach as a timely surveillance tool. It is clear that there is useful information in the posterior distributions that could be used for surveillance decision making. Calibration and evaluation will need to be based on a specific public health application and given different tolerance for sensitivity and timeliness of alerts. This will be explored in future work. For example, thresholds can be set using historical data and rather than basing the threshold decision making on the median, other quantiles of the posterior, or the upper bounds of the PI, could be used to trigger an alert. The performance in terms of specificity and sensitivity of the models can be made once thresholds for decision making are available for a specific setting and context.

There are several other key next steps for this research. First, for regions with large populations and health counts where the smoothing is less beneficial, the spatio-temporal model could be adjusted by explicitly weighting the smoothing inversely to population size or mean expected count. Ideally, this weighting would be optimized for prediction accuracy and variance trade-off. Such smoothing can be made to suit the needs of the public health organizing using the models, based on the demographics of the communities under surveillance, and on the nature of the morbidity or mortality data used for monitoring public health response to smoke.

Second, future work should evaluate the utility of the proposed models for use in a real public health setting, and fine-tune the approach to maximize practical public health benefits. The proposed models can be used to estimate the expected health impacts under different exposure scenarios and this information can be used to guide decision making. These evaluations should include external sources of information such as scheduled outdoor events, wildfire behaviour, relevant climate factors, and, in particular, any available forecasts of wildfire smoke or air pollution, such as those available from BlueSky and FireWork [39, 101]. Reliable future estimates of air pollution would allow for meaningful forecasts beyond 24 hours in the future, providing public health with more time to plan and implement interventions, such as disseminating public health messaging. Estimation of population-level exposure to air pollution has a long history in exposure science, and considerable progress has been made in the last decade [126, 39, 110, 37, 127]. Advances in this field will have a profound effect on air pollution public health surveillance.

Finally, the exposure measurement error model used here is simple, based on the assumption of non-differential measurement error. The uncertainty associated with the predictive LHA-level $PM_{2.5}$ model [37] used to estimate exposure was not incorporated explicitly into these forecasts. A more cohesive approach to environmental public health surveillance would be to explicitly incorporate the smoke or air pollution modelling into this hierarchical model and to propagate the uncertainty of these estimates through to the forecasting of the associated health indicators.

Wildfire smoke is a serious public health concern with the potential to expose entire populations across vast geographic areas to extremely high levels of air pollution, including $PM_{2.5}$ and other harmful contaminants. Vulnerable populations, such as those with chronic lung disease, will be particularly at risk. Severe smoke events are expected to continue in the coming decades as climate change and land use changes increase the duration and severity of wildfire seasons [108, 22]. Practical, well-evaluated, flexible surveillance tools are needed to assist public health with decision-making during smoke events.

5 One step ahead of the epidemiologic curve: a forecasting approach for public health surveillance of wildfire smoke

5.1 Preamble

In the previous two chapters, my objective was to identify a statistical approach that would be appropriate for the characteristics of the exposure and outcome data encountered in EPHS, be flexible so that changes to data characteristics could be accommodated (e.g., additional outcomes or changes in exposure metrics), and address the computational constraints of requiring an 'online' daily surveillance system. In the following chapter, I evaluate how the proposed system could be used in surveillance.

To do this, I retrospectively analyze an urban wildfire smoke event by integrating the previously used air pollution estimates and health outcome indicators with smoke forecasts to assess our ability to quickly detect a public health impact. In July 2015, severe wildfires burning in southwestern British Columbia caused a prolonged period of smoky conditions across the highly populated regions of Greater Vancouver and Vancouver Island. The potential impact of this event and the lack of real-time surveillance tools to facilitate decision making has further motivated research in wildfire smoke surveillance.

This manuscript will be submitted to *Environmental Health Perspectives*.

5.2 Title page and footnotes

Title: One step ahead of the epidemiologic curve: a forecasting approach for public health surveillance of wildfire smoke

Short title: Wildfire smoke public health surveillance

Authors: Kathryn T Morrison[1], David L. Buckeridge[1], Gavin Shaddick[2], Sarah B. Henderson[3,4]

Affiliations:

[1] Department of Epidemiology, Biostatistics, and Occupational Health, McGill University, Montreal, Quebec, Canada

[2] Department of Mathematical Sciences, University of Bath, Bath, United Kingdom

[3] Environmental Health Services, British Columbia Centre for Disease Control

[4] School of Population and Public Health, University of British Columbia

Corresponding author:

Kathryn Morrison

1140 Pine Avenue West

Montreal, Quebec, Canada

H3A 1A3

E-mail: kt.morrison@mail.mcgill.ca

5.3 Abstract

Introduction: Severe wildfires burned near densely populated regions of British Columbia, Canada in July 2015. Millions of individuals were exposed to the smoke, and such exposures can cause increases in asthma-related physician visits and dispensations of asthma relief medications. Public health authorities did not have access to a surveillance system capable of predicting the health impacts to guide decision-making and no public health interventions were initiated during the smoke exposure.

Objectives: To evaluate the ability of an integrated surveillance system to predict the health effects of wildfire smoke exposure and to estimate the morbidity potentially preventable through public health interventions triggered by such a system.

Methods: We retrospectively emulated the use of a wildfire smoke surveillance system during the severe wildfires that occurred in 2015 in BC. Daily estimated $PM_{2.5}$ concentrations and daily counts of two indicators of population respiratory health were obtained for 41 local health areas in BC during the 2014 and 2015 wildfire seasons. Smoke forecasts from the BlueSky system were obtained for the smoky period. A computationally efficient hierarchical time series model was used to predict counts of the health indicators 24 and 48 hours into the future using forecasted smoke concentrations from BlueSky. Potentially preventable morbidity was also estimated, based on the scenario of the surveillance system triggering, at different time points during the smoke exposure, a public health messaging campaign targeted at vulnerable individuals.

Results: Smoky conditions resulted in $PM_{2.5}$ concentrations up to ten times higher than background levels between July 5 and July 10, 2015. Asthma-related physician visits and dispensations of relief medications were increased by 34% and 27%, respectively, during this period. On July 6, the smokiest day, the increases were 72% and 78%, respectively. The surveillance system provided prediction intervals for the 24-hour and 48-hour health forecasts for each day of the July 3 to July 10, 2015 period. The 95% prediction intervals contained the true observed values in 12 of the 14 dispensation forecasts, and 9 of the 14 physician visits forecasts. The surveillance system did not require exact smoke forecasts to produce reasonable health indicator forecasts, but the accuracy of health forecasts was affected by large errors in smoke forecasts. Simple, early interventions with lower effectiveness, such as messaging about steps to reduce exposure, would likely have prevented more morbidity than delayed interventions with higher effectiveness, such as evacuation.

Conclusion: After retrospectively emulating the application of a wildfire smoke surveillance system to the severe wildfires in BC in 2015, we conclude that integrating data from multiple sources into a wildfire smoke forecasting model can accurately predict the health impacts of smoke exposure in a timely manner. Forecasts from such a system could be used to trigger public health action and guide the selection of interventions to prevent morbidity. Prospective evaluation of wildfire smoke surveillance systems in practice settings is now needed to generate evidence that can guide their adoption and effective use.

Keywords: Wildfire smoke, air pollution, forecasting, public health surveillance

5.4 Introduction

The summer of 2015 was one of the smokiest wildfire seasons on record in the province of British Columbia (BC), Canada, especially in the wetter coastal areas where fire risk was unusually elevated due to a protracted drought. Most notably, the Elaho and Boulder Creek fires were discovered burning near Pemberton (Figure 5–1) on June 14, and they grew to 12,523 and 6,735 hectares, respectively, before they were extinguished [128]. From July 5 to July 10 smoke from these two fires severely affected the densely populated regions of greater Vancouver and Vancouver Island. Using the World Health Organization (WHO) daily fine particulate matter ($PM_{2.5}$) threshold of 25 $\mu g/m^3$, these wildfires caused approximately 14,000,000 person-days of unhealthy exposure. Previous research has demonstrated a strong association between smoke exposure and multiple indicators of respiratory health, including (1) asthma-related physician visits, and (2) dispensations of the asthma relief medication, salbutamol sulphate [35]. After the extreme fire season of 2010 the BC Centre for Disease Control (BCCDC) used these two respiratory health indicators to develop the BC Asthma Monitoring System (BCAMS) [42]. Retrospective surveillance reports from BCAMS showed sharp increases in both indicators and observed air pollution concentrations during the 2015 smoke episode, but the reports were not timely enough to facilitate intervention.

Conventionally, public health surveillance has been used to monitor infectious diseases by analyzing the incidence of confirmed or suspected cases over time, detecting changes from what is expected, and triggering timely public health interventions to prevent disease transmission and mitigate the impact of an epidemic. Environmental public health surveillance (EPHS) is analogous to infectious disease surveillance, but there are important differences such as a focus on exposure monitoring and the nature of the interventions. The relevance of EPHS is growing as the global climate changes [5], resulting in more frequent and severe environmental exposures to hazards that have population health impacts, such as heat and wildfire smoke [129]. Similar to infectious disease surveillance, an objective of EPHS is to rapidly detect changes in the health status of the population, and to provide timely information to guide public health interventions to reduce hazardous exposures. Most EPHS systems differ from infections disease surveillance systems in that they detect aberrant changes in the exposure, either independently or linked to the changes in population health attributable to that exposure.

During a severe smoke event with very high $PM_{2.5}$ concentrations, the ideal public health intervention would prevent excess health outcomes by maintaining population exposure at background levels. There are no perfect interventions that could achieve this ideal, but there are practical interventions that can reduce exposure and subsequent morbidity, especially for the most susceptible populations [130]. Research in this area is limited, and stakeholders have expressed the need for clear evidence about the efficacy and timeliness of different interventions. In one of the few existing studies, Mott *et al.* 2002 found that over 80% of individuals in Hoopa, California recalled at least one public safety advisory following a severe smoke event [61]. Furthermore, those who recalled the advisories were less likely to report worsening symptoms, especially people with asthma and the parents of younger children. A recent opinion article has also called for stronger recommendations regarding use of home air cleaners during smoke episodes, based on evidence of improvements in both indoor air quality and health endpoints [131]. Simulation-based research has shown that interventions based on smoke forecasts would significantly reduce the associated morbidity and its costs, particularly for vulnerable populations such as asthmatics [132, 133].

We know from previous studies in BC that smoke forecasts can provide public health authorities with valuable information about expected exposures [42, 35]. However, the utility of smoke forecasts is limited because they are neither readily accessible nor interpretable by many public health practitioners. First, there are multiple modelling uncertainties in smoke forecasting that cannot be easily understood without sophisticated knowledge about forest fires and atmospheric processes. Second, forecasts indicate the expected distribution of smoke without considering the underlying distributions and vulnerabilities of the potentially exposed populations. Finally, smoke forecasts are generally available on web platforms, but most public health professionals do not have the skills needed to systematically download, process, and use the raw data files. Here, we attempt to address these challenges by integrating smoke forecasts and health outcomes in a single surveillance system that can guide public health practice proactively by predicting future asthma-related morbidity based on anticipated smoke levels.

The current implementation of the BCAMS system retrospectively gathers data about smoke from smoke forecasting frameworks, air quality monitoring stations, and satellites. These data are spatially integrated with health indicators, and reports are pushed out to regional medical health officers [42]. The reports are always retrospective, because the system uses observed rather

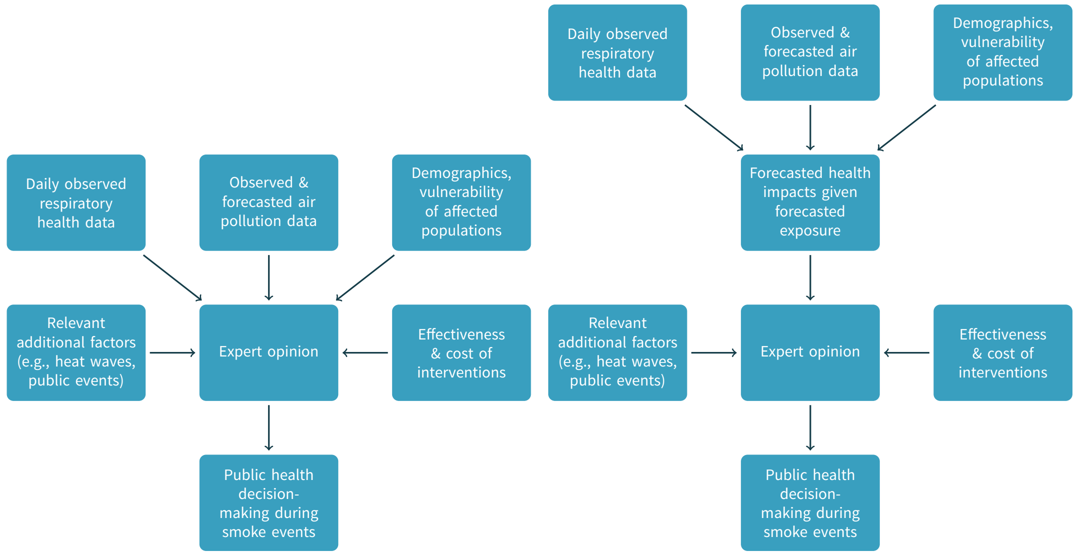

Figure 5–1: Schematic showing the existing (left) and proposed (right) decision-making process in wildfire smoke surveillance.

than predicted information about the respiratory health indicators. For example, a BCAMS report supplied on July 7, 2015 would have included health indicator counts up to July 6, a delay of at least 24 hours. Here we propose the prospective BC Asthma Prediction System (BCAPS) as an alternate public health framework in which smoke forecasts are routinely downloaded, spatially matched to population data, and simplified before being fed into a health forecasting model based on the prior research presented in Chapters 3 and 4. Under this scenario, public health professionals would receive daily health forecasts in the early morning providing predictions for the following 24 and 48 hours, including measures of uncertainty, rather than being expected to infer such predictions based on their knowledge and the smoke forecasts alone. Such reports could be used to inform proactive, evidence-based decisions (Figure 5–1) in the context of other information contributing to situational awareness. We demonstrate the applications of BCAPS by retrospectively emulating its real-time use to forecast the 2015 wildfire smoke event in southern BC. We also estimate the potential exposure reduction and public health benefits that could have been achieved if BCPAS had been used to guide intervention decisions during that period, as opposed to waiting until health impacts had already begun to occur to intervene, or failing to intervene altogether.

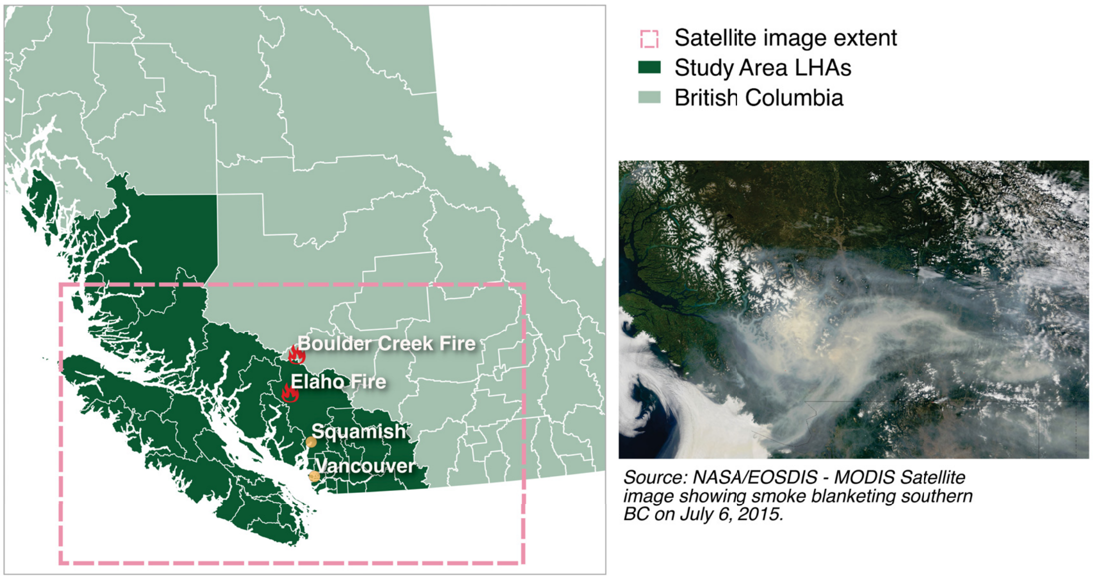

Figure 5–2: Map of BC showing the 41 study area LHAs and a satellite image blanketed in wildfire smoke on July 6, 2015.

5.5 Methods

5.5.1 Study context

The heavily forested province of BC is located on the west coast of Canada, with much of its population living in the south-western corner of the province. Local Health Areas (LHAs) are the smallest geographic health units defined by the provincial government for administration and data dissemination purposes. Our study area is defined by a subset of 41 LHAs that were highly exposed to $PM_{2.5}$ during the wildfires of July 2015 (Figure 5–2). This area includes greater Vancouver and Vancouver Island as well as surrounding areas, with a total population of 3,658,523 persons in 2015 [114], approximately 80% of the provincial population. We defined the study period as July 3 to July 10, 2015, where wildfire smoke caused increased $PM_{2.5}$ exposures from July 5 to July 10, which we defined as the high exposure period. We include July 3 and 4 in the study period to demonstrate how the smoke event unfolded.

5.5.2 Environmental data

The BCAPS approach requires both historical $PM_{2.5}$ concentrations and $PM_{2.5}$ forecasts. We used the BCCDC Optimized Statistical Smoke Exposure Model (OSSEM) for the historical estimates. This model integrates $PM_{2.5}$ measurements from the provincial air quality monitoring network with remotely sensed aerosols, remotely sensed fires, tracings of smoke plumes from multiple satellite images, and the atmospheric venting index [35]. The model estimates $PM_{2.5}$ concentrations at a spatial resolution of 5x5 km, and has been used in real-time surveillance [42] and epidemiologic research [35] since 2012. Every morning OSSEM produces $PM_{2.5}$ estimates for the previous day, and we obtained all estimates for the 2014-2015 period.

We used the BlueSky Canada smoke forecasting framework for the $PM_{2.5}$ forecasts. This system uses current data to generate hourly forecasts of $PM_{2.5}$ concentrations from forest fires up to 60 hours in advance with a spatial resolution of 4x4 km. It is a complex forecasting system, combining meteorological forecasts, fire locations, fuel consumption estimates, and smoke emissions estimates in a dispersion model to predict ground-level $PM_{2.5}$ concentrations [39]. Unlike some other forecasting systems, it reflects only the $PM_{2.5}$ contributions from forest fire smoke, and not the contributions from all sources. BlueSky forecasts have been used in epidemiologic research [35] and public health surveillance [42], and research improving the forecasts is ongoing [40]. BlueSky generates two forecasts per day, and we used the 05:00 smoke forecast to generate health forecasts for the next 24 and 48 hours.

To assign daily OSSEM and BlueSky values to the study area we took a population-weighted average using data from the 2006 census. Disseminations areas (DAs) are the second-smallest geographic census unit in Canada, with most DAs having a population of 400-700 residents [134]. We mapped the geographic centre of each of the 7469 DAs in the study area, and then assigned daily estimates using the grid cell that contained the centroid. In addition, we assigned daily values for maximum apparent temperature at the nearest Environment Canada weather station, and a daily indicator of fire activity based on satellite measurements of radiative power, as done in Chapters 3 and 4 [35].

5.5.3 Respiratory health indicators

The current implementation of BCAMS uses two indicators of population respiratory health: (1) dispensations of salbutamol sulphate, and (2) asthma-related physician visits. Salbutamol is a prescription drug used to treat the acute symptoms of asthma and chronic obstructive lung diseases such as COPD [116]. Research has shown that smoky conditions are associated with increased salbutamol dispensations in BC [135]. The BC government requires every prescription dispensed to be logged by the provincial PharmaNet database, regardless of the recipient or the payer, and these data are made available to the BCCDC for surveillance purposes [44, 117, 35]. Research has also shown increased asthma-related physician visits during smoky periods in BC [66, 35]. These visit are paid for by the provincial health insurance program, and physician billings are logged to the central Medical Services Plan (MSP) database [117]. Visits are coded by reason using the International Classification of Diseases, 9th Revision (ICD-9). Only visits for the code specific to asthma (group 493) are used in BCAMS [42]. Daily counts of both indicators were retrieved for each LHA in the study area for the period of 2014 through August 2015.

5.5.4 Forecasting respiratory health indicators with a bivariate latent process model

Our previous work proposed and evaluated a statistical approach for forecasting counts of multiple health indicators that are associated with a common environmental exposure. Because estimating population exposure to forest fire smoke across large geographic areas is challenging and prone to error, our statistical approach was to consider the true exposure as latent, or unmeasurable, and to assume that $PM_{2.5}$ estimates contain some error [112, 48]. By considering the exposure as latent we incorporate this uncertainty directly into the BCAPS health indicator forecasting system.

Because wildfire smoke can have many health impacts, multiple different outcomes can be used to assess the relationship between smoke exposure and population health. Our previously developed approach allows for multiple outcomes to be simultaneously forecasted against the common latent exposure. The multivariate approach puts less emphasis on a single, potentially noisy data source and improves overall forecast performance in terms of accuracy and precision [112]. Computational processing is a significant constraint in public health, as forecasts must be timely. To implement the modelling approach and ensure efficient computation, we used a hierarchical Bayesian framework with model fitting performed using integrated nested Laplace approximations (INLA) [53]. This approach also provides posterior prediction intervals for each forecast, giving the likely bounds around the mean or median estimates. When making health forecasts using this framework, the model is trained on historical, observed data and it uses any available information about future conditions to generate its predictions. More specifically, the historical model is trained using the following daily observations: counts of salbutamol dispensations; counts of physician visits; $PM_{2.5}$ concentrations from OSSEM; maximum apparent temperature; fire activity indicator; and day-of-week. As such, the health forecasts for July 6, 2015 would use all historical data up to July 5, the day-of-week for July 6, and $PM_{2.5}$ forecasts from BlueSky for July 6 to make the prediction. Similarly, the health forecasts for July 7 would include all historical data up to July 6, and so forth.

5.5.5 Day-by-day, retrospective emulation of BCAPS use in a public health scenario

We emulated the use of the proposed BCAPS surveillance system as if it had been operational in July 2015. For each day in the study period, we used all available observed data from May 1, 2014 forward to fit the health forecasting model and to produce estimates for the following 24 and 48 hours. The proposed system takes the input data and the modelling results and summarizes all of the information graphically for interpretation. The system could be modified to work with different input data and at different temporal and geographic scales. Here, we simulated completion of the following steps every morning:

1. At 05:00 PST the BlueSky smoke forecasts became available. We extracted the population-weighted mean forecasted $PM_{2.5}$ concentrations for each hour in each of the 41 LHAs over the next 48 hours. Then, the hourly data were averaged across the first 0-24 hour period and the second 25-48 hour period to generate daily mean smoke forecasts for each LHA.
2. The daily mean smoke forecasts for each LHA were classified as low, moderate, or high smoke to improve visual interpretation; the following $PM_{2.5}$ thresholds were used: low was less than 25 $\mu g/m^3$; moderate was between 25 and 60 $\mu g/m^3$; and high was greater than 60 $\mu g/m^3$. These values were chosen because they are consistent with thresholds used for the Air Quality Health Index in BC.
3. For each of the 0-24 hour and 25-48 hour forecast periods, the maximum daily mean smoke forecast was identified from the 41 LHAs in the study area. These two values were fed into the BCAPS system to produce forecasted health outcomes for the entire region, expressed as prediction intervals. The maximum values were chosen because smoke forecasts are inherently uncertain, and health forecasts should be designed to include the worst case scenario.

Once the necessary historical data have been collated, this entire process can be automated, including retrieval of the BlueSky forecasts from the server, generation of the health forecasts and prediction intervals, and production of a summary report that could be disseminated to the relevant public health professionals early each morning.

5.5.6 Estimating the potential impacts of public health interventions

Both of the respiratory health indicators suggested impacts on chronic lung diseases such as asthma and COPD, which affect at least 6.2% and 2.3% of the BC population, respectively [20]. Given a scenario in which public health professionals receive a daily report about the potential population health impacts of expected smoke exposures, we assume that they would take action to start mitigating the risks in advance of the exposures, because they would have information on expected morbidities ahead of time. As such, we estimated the impacts of hypothetical public health interventions based on the timing of their initiation, their estimated exposure reduction, and their effectiveness. The latter is defined as the actual proportion of a population that experiences risk reduction, which is the product of adherence to the intervention and the efficacy of the intervention [136]. Specifically, we estimated the decrease in excess morbidity counts for each health outcome using the following steps for each day of the study period:

1. The expected count for each day was estimated as the mean count for that outcome on that day-of -week (e.g. all Tuesdays) during the selected low exposure period of May 1 to July 1, 2015.
2. The number of excess salbutamol dispensations and physician visits was calculated by subtracting the expected counts from the observed counts for each day.
3. The number of potentially preventable excess dispensations and visits was calculated for different levels of effectiveness, ranging from 10% to 90%, and for different timings of initiation.
4. The estimated potentially preventable excess dispensations and visits were compared for three intervention scenarios:
   (a) No intervention;

(b) A retrospective surveillance system (BCAMS) triggering public health intervention when an increase in morbidity is observed in conjunction with increased air pollution;

(c) A prospective surveillance and prediction system (BCAPS) triggering public health interventions when a increase in morbidity is forecasted based on historical air pollution and morbidity and forecasted air pollution.

Research on the effectiveness of interventions is sparse, but the available evidence suggests that adherence to public health recommendations is higher among vulnerable groups, such as asthmatics [61]. Furthermore, there is some evidence that recommendations to stay indoors, limit physical activity, use air conditioners and air filters are associated with decreased negative health effects [137], whereas recommendations to use masks are generally unsupported [138]. Evacuation for vulnerable populations or entire populations may be recommended in extreme cases, either by mandate or financial incentives[68, 39, 61]. The small amount of research on evacuations has shown limited effectiveness, and is associated with considerable logistic and economic challenges [68, 64]. Simple recommendations, such as remaining indoors and limiting physical activity, may have higher adherence and faster uptake than more complex recommendations, such a purchasing an air cleaner. However, personal air cleaners have been recognized and recommended as one of the best interventions that public health professional can recommend for vulnerable people sheltering in place [131].

5.6 Results

5.6.1 Summary of smoke exposure

Smoke had a negligible impact on air quality in the entire study area From May 1 to July 2 2015. The mean daily OSSEM $PM_{2.5}$ concentration across the 41 LHAs during this low-exposure period was 6.3 $\mu g/m^3$ and the maximum was 11.6 $\mu g/m^3$. During the July 3 to 10 study period, the concentrations remained low on July 3 and 4 at 6.8 $\mu g/m^3$ and 8.3 $\mu g/m^3$ respectively, but began to increase between July 5 to 10 with a mean of 33.9 $\mu g/m^3$ and a maximum of 61.7 $\mu g/m^3$ on July 6 (Figure 5–3).

On July 5 the observed OSSEM $PM_{2.5}$ concentration rose to 15.2 $\mu g/m^3$ as smoke from wildfires burning near Squamish began to sweep across southern BC. The BlueSky forecast

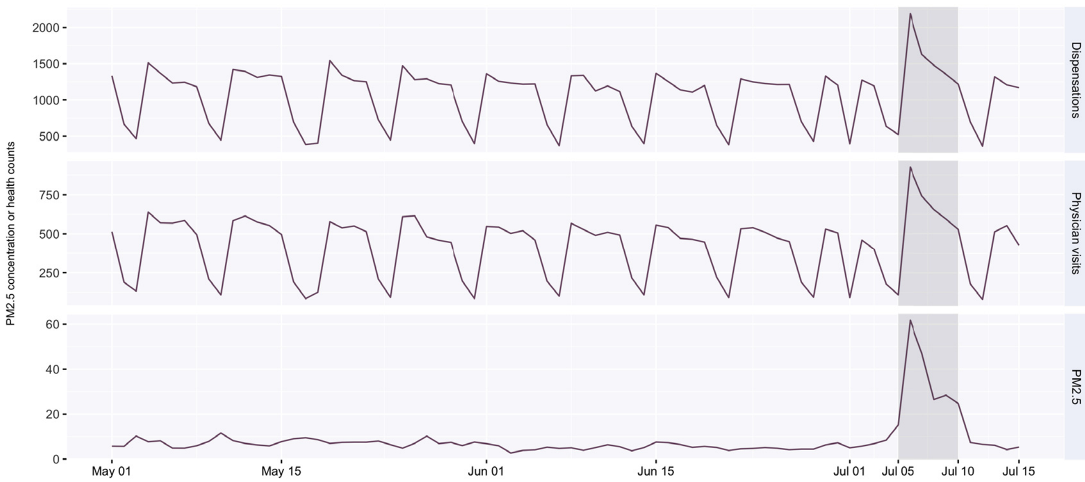

Figure 5–3: Time series showing the daily observed particulate matter ($PM_{2.5}$) concentrations from the Optimized Statistical Smoke Exposure Model (OSSEM) along with the observed cumulative counts of salbutamol dispensations and asthma-related physician visits. The highlighted portion is the period of high exposure from July 5 through 10.

from the morning of July 5 incorrectly predicted low concentrations throughout the province for the remainder of July 5 and, notably, for July 6 (Figure 5–4). However, the average observed concentration on July 6 was 61.7 $\mu g/m^3$, which was the highest for this geographic area during the entire wildfire season, at approximately ten times background concentrations. On the mornings of July 6, 7, and 8, BlueSky correctly forecasted high smoke for the following 24-hr and 48-hr periods. On the morning of July 9, BlueSky correctly forecasted high concentrations for that day, with moderate concentrations for the following day (Figure 5–4). Observed concentrations remained above 25$\mu g/m^3$ until July 10, when they decreased towards background levels (Figure 5–3).

5.6.2 Summary of respiratory health indicators

The mean daily counts of salbutamol dispensations and physician visits during the low exposure period of May 1 to July 1 2015 were 1,038 and 405, respectively. During the high exposure period the mean daily counts were 1,399 salbutamol dispensations and 592 physician visits. The highest peak of each indicator occurred on Monday July 6, coinciding with peak observed $PM_{2.5}$ concentrations (Figure 5–3). Both salbutamol dispensations and physician visits exhibit strong day-of-week patterns, with the highest counts often observed on Mondays (Figure 5–3). Even so, the counts for both indicators on Monday July 6 were considerably elevated (1.7-fold for

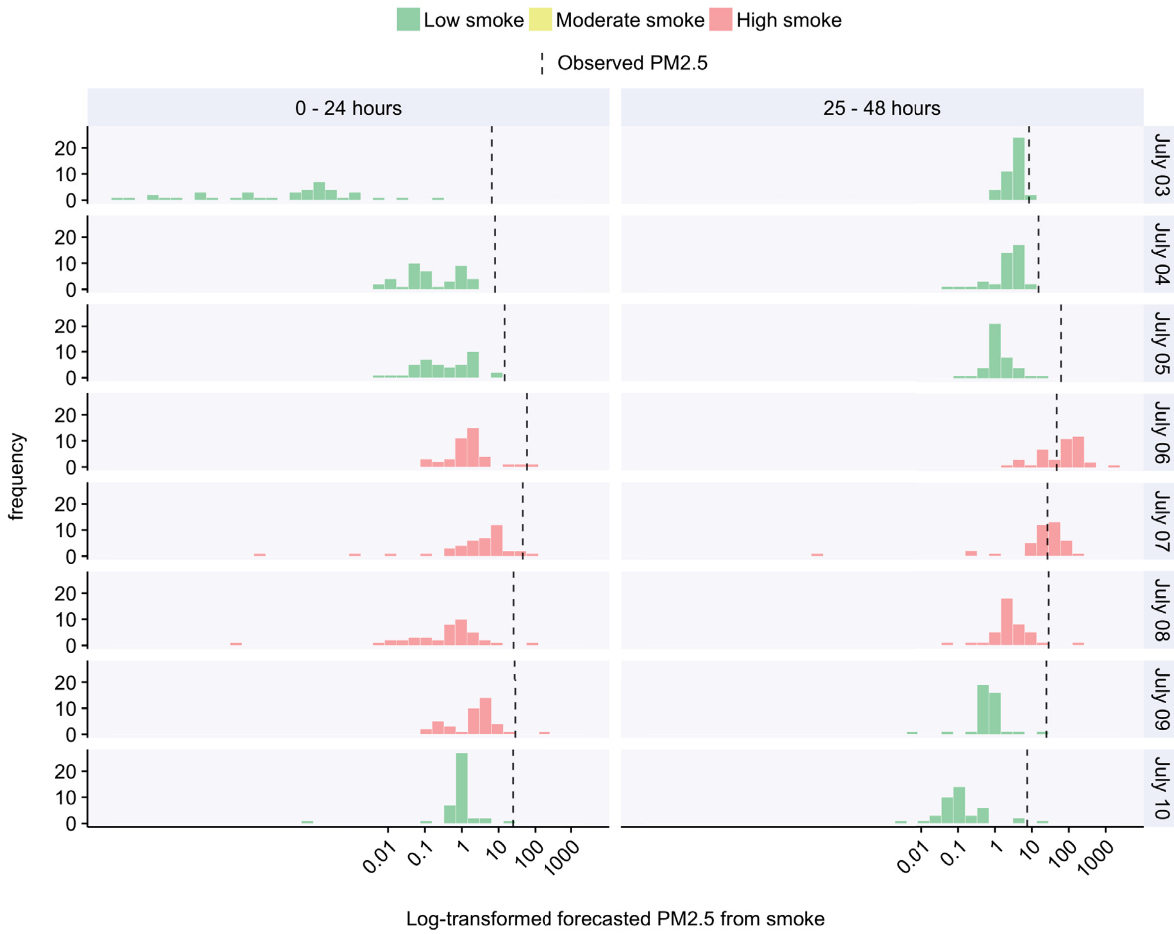

Figure 5–4: Histograms of the BlueSky 24-hour and 48-hour smoke forecasts at 05:00 each morning (right axis) for the following 0-24 and 25-48 hour periods. These are the distributions of the 24-hour concentrations for the 41 LHAs in the study area. Histogram colour indicates the severity of $PM_{2.5}$ concentrations with low hourly maximums < 25 $\mu g/m^3$, moderate hourly maximums between 25 and 60 $\mu g/m^3$, and high hourly maximums > 60 $\mu g/m^3$. The vertical line indicates the observed 24-hour mean $PM_{2.5}$ concentrations from the Optimized Statistical Smoke Exposure Model (OSSEM) for the study period on the forecast date.

dispensations, 1.8-fold for visits) above the average counts for Mondays within the low exposure period (Table 5–1).

Table 5–1: Daily counts for each respiratory health indicator during the study period, and percent difference when compared with expected counts for same day of week during the May 1 - July 1 2015 low exposure period.

	Dispensations		Physician visits	
	Observed count	**% Difference from expected count**	**Observed count**	**% Difference from expected count**
Fri, July 3	1193	-3	401	-16
Sat, July 4	637	-6	178	-11
Sun, July 5	522	27	108	9
Mon, July 6	2192	72	926	78
Tue, July 7	1628	23	741	33
Wed, July 8	1478	29	654	39
Thu, July 9	1353	10	593	15
Fri, July 10	1218	-1	527	10
Excess[1]	**1785**		**908**	

[1] Excess during high exposure period (July 5 to 10)

5.6.3 Health forecasts

All observed counts for salbutamol dispensations were contained within the 95% prediction intervals from the health outcomes forecasting model, with the exception of the forecasted values for July 6 (Figure 5–5). On the morning of July 5, the 25-48 hour health forecast interval for July 6 did not contain the true value, and neither did the 0-24 hour health forecast produced on the morning of July 6. Lower accuracy for this day was expected, given that smoke forecasts did not predict the peak $PM_{2.5}$ concentrations that occurred on July 6, when both of the health indicator counts also peaked. Generally, the observed counts were close to or in excess of the upper bounds of the 95% prediction interval on days when the smoke forecasts predicted higher exposures when compared with day when lower exposures were predicted (Figure 5–5).

Several of the 95% prediction intervals for the asthma-related physician visits on July 6, 7, and 8 did not include the observed counts, although most were borderline (Figure 5–5). Once again, the July 5 and July 6 morning health forecasts for July 6 both underestimated the impacts because smoke forecasts were low. The health forecasts performed better on July 7 and 8, when the smoke forecasts more accurately predicted an increase in $PM_{2.5}$ that correlated well with observed

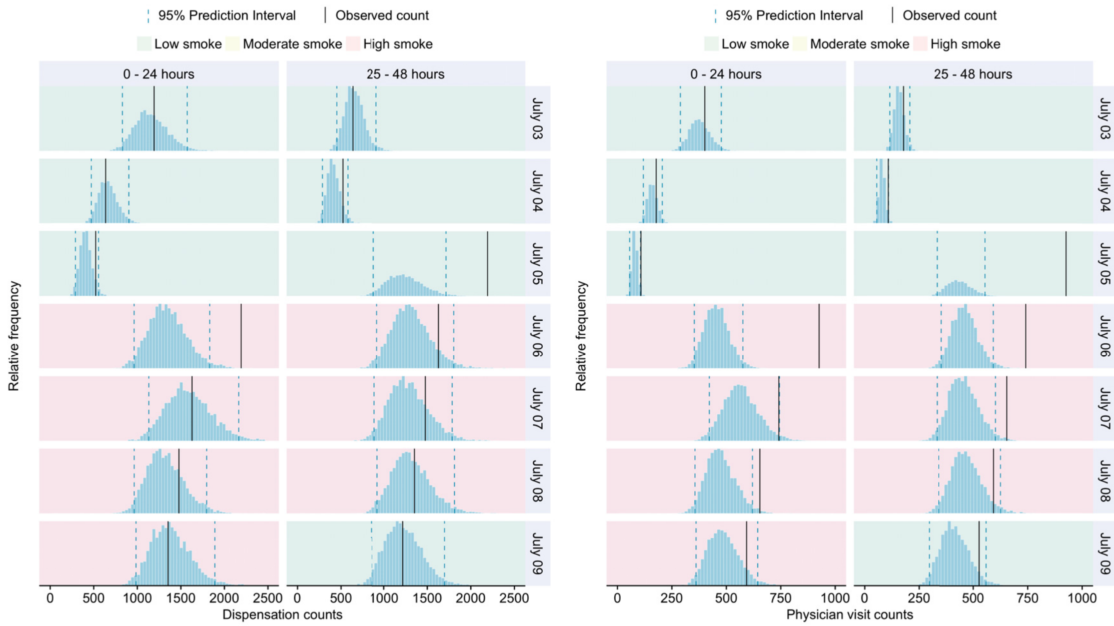

Figure 5–5: The 0-24 hour and 25-48 hour prediction intervals and observed counts for the entire study area are shown for each day of the study period. The y-axis is omitted because the absolute frequency is a function of the number of posterior samples taken to construct the prediction interval, and therefore only meaningful in terms of relative frequency. The background of each histogram is coloured to indicate the severity of the smoke forecast used to produce the health forecast.

concentrations.

5.7 Generating wildfire smoke surveillance reports

On any given day during the wildfire season, morning smoke forecasts generally become available at 05:00. An automated surveillance summary report could be produced as soon as the smoke forecasts are published. A surveillance report that shows the anticipated air quality impacts and their corresponding health impacts could reduce the volume of disparate data that public health professionals are currently required to consult, and it could transform those data into more actionable information (Figure 5–6).

5.7.1 Potential impact of interventions

We estimated the total number of excess morbidities (sum of salbutamol dispensations and physician visits) that could have been prevented under three intervention scenarios (Table 5–2):

(a) Using no surveillance, monitoring system or interventions

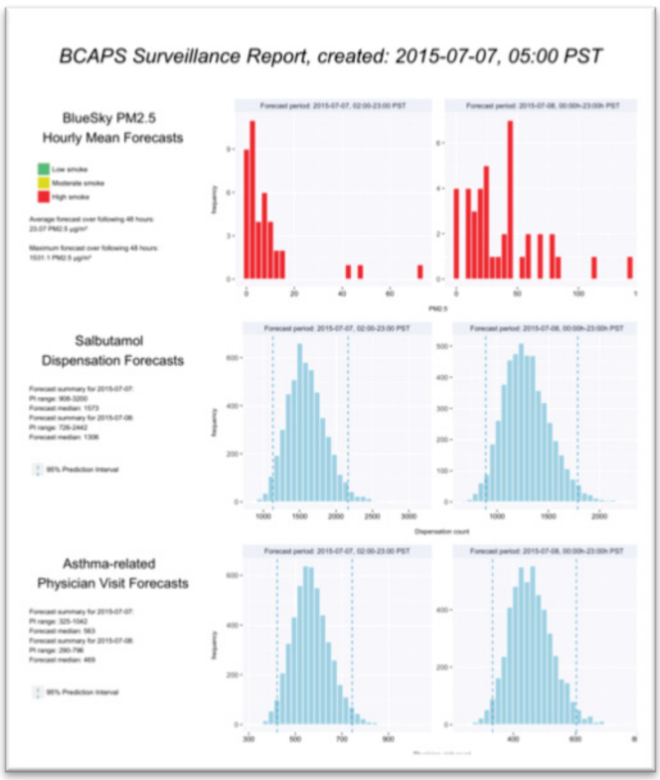

Figure 5–6: An example summary report that could be automatically generated as soon as the morning smoke forecast is published. The entire process for generating such reports can be automated.

(b) Using BCAMS to detect elevated morbidities retrospectively

(c) Using BCAPS to forecast anticipated elevated morbidities prospectively

Naturally, no excess morbidities would have been prevented if no action had been taken. The foresight provided by BCAPS could have prevented twice as many overall excess morbidity counts compared to a retrospective monitoring system like BCAMS, because interventions would have been triggered on the morning of July 6 rather than on July 7. An intervention with 50% effectiveness initiated on July 6 by BCAPS could have prevented more combined health indicator counts than even an intervention with 100% effectiveness initiated on July 7 by BCAMS.

Table 5–2: Estimated number (percent) of excess dispensations and physician visits that could have been prevented had the a surveillance system been used to trigger an intervention of a given effectiveness. There were a total of 908 excess health outcome indicators during the high exposure period of July 5 to July 10. BCAPS-based interventions could have occurred 24 hours earlier than with BCAMS.

		100% effective	75% effective	50% effective	25% effective
Dispensations	None	0 (0%)	0 (0%)	0 (0%)	0 (0%)
	BCAMS	760 (43%)	570 (32%)	380 (21%)	190 (11%)
	BCAPS	1675 (94%)	1256 (70%)	837.5 (47%)	419 (23%)
Visits	None	0 (0%)	0 (0%)	0 (0%)	0 (0%)
	BCAMS	494 (54%)	370.5 (41%)	247 (27%)	124 (14%)
	BCAPS	899 (99%)	674 (74%)	450 (50%)	225 (25%)

5.8 Discussion

In areas seasonally affected by wildfire smoke, surveillance programs are under development to monitor the population health response to smoky conditions. Data collection and analysis must be timely to provide public health authorities with actionable information about an acute environmental exposure. The data that are available and appropriate for smoke surveillance purposes are generally noisy, such as time series of healthcare usage counts, and prone to misclassifying exposure, such as air pollution estimates. Here, we have proposed and evaluated an integrated surveillance system for the public health impacts of wildfire smoke.

Using historical data, smoke forecasts, and the methodological framework proposed by Morrison *et al.*, we demonstrated that it would have been possible to accurately forecast two health indicators, counts of salbutamol dispensations and asthma-related physician visits, during a massive wildfire smoke exposure in the province of BC [112]. Although further research on interventions in different settings is needed, we explored the potential for forecasts of health indicators from the proposed BCAPS system to trigger public health action to mitigate health effects of wildfire smoke exposure. The existing surveillance system, BCAMS, would have retrospectively shown a concurrent increase in air pollution and asthma-related health indicators, but intervention after that point would have prevented less than half of the morbidity, even with an intervention that was 100% effective. The BCAPS approach proposed here builds upon the BCAMS system by integrating historic healthcare utilization data with smoke predictions to forecast the likely health impacts of smoky conditions. This foresight could have resulted in proactive intervention on July 6, and prevented nearly two thirds of health care utilization with an intervention that was only 65% effective. We found that the accuracy of the health forecasts was heavily dependent on the accuracy of the smoke forecasts. Overall, we found that the timeliness of the intervention was more important than effectiveness of the intervention. A less effective intervention with faster uptake would have reduced morbidity more than a highly effective delayed uptake. These results highlight the critical importance of reliable smoke forecasts for public health and other applications.

The statistical approach used here was an autoregressive latent process model that simultaneously forecasts two health indicators against a common exposure. The motivation for using such an approach is noisiness of the health indicator data. To avoid an excessive number of false alerts in the predictions, the models smooth towards the historical means of both indicators, which has considerable benefits when working syndromic surveillance data. For example, if one indicator rose by chance but the other did not, both of the forecasted health indicators would be smoothed towards the mean unless the variation was sustained or the other health indicator was also increased. Because healthcare utilization data can be noisy, it is beneficial to rely on multiple streams of data, rather than placing too much importance on a single health indicator [139, 140]. However, there are some disadvantages to this approach. The model may be slower to adjust to sudden aberrations in the exposure or the outcomes, which may partly explain why the first day of high exposure was less accurately predicted than the subsequent days. Although we used the worst case scenario from the smoke forecasts, the training data did not contain any exposure estimates beyond a concentration of $65\mu g/m^3$ because this event was unprecedented in the study area. For this reason, the prediction intervals for both health indicators were shifted, but not to the extreme observed values. However, if there had been periods in the training data with very high observed $PM_{2.5}$ concentrations and corresponding high health outcome counts, the subsequent prediction intervals would likely have been more accurate. The model also generates prediction interval distributions, which provide more information than a single predicted value. Given the model performance described here, decision makers should consider the upper bounds of the prediction interval as the most likely estimate when $PM_{2.5}$ values are predicted to be extreme. If the methodological framework proposed here were integrated into an explicit decision-making framework, an upper quantile of the prediction intervals could be used to calibrate the signals based on desired levels of sensitivity and specificity.

We made several assumptions, particularly regarding data availability and interventions. Scientists at the BCCDC are engaged in research to design a proactive surveillance system based on the two health indicators used here. However, the salbutamol dispensations data are not currently available in real time, and further administrative work is needed to facilitate a daily feed. Similarly, the physician visits data are populated as clinicians submit their billings to MSP, meaning that the daily feed only includes complete information for those who bill on a daily basis. The current implementation of BCAMS typically uses all billings submitted within the past five days, though it can be run for daily billings alone. Regardless, any single health indicator has strengths and limitations, such as the availability or timeliness of the data and the severity of the symptoms it is used to assess. The use of multiple health indicators can be used to overcome the limitations of any single indicator, and adding additional data could be beneficial to the performance of BCAPS. For example, ambulance dispatch data are theoretically available in real-time, and others have shown an association with smoke exposures [141]. The methodological approach here would easily extend to include such additional health indicators. There may be computational challenges as the number indicators grows, but they should be surmountable given the efficient Bayesian approach using INLA approximations and modern parallel computing.

Improving smoke forecasts for the purposes epidemiologic exposure assessment and public health surveillance is an active area of research [40]. The latent process model used here assumes some non-differential misclassification of exposure, but a more sophisticated approach could be used to perform spatial interpolations of the historical OSSEM $PM_{2.5}$ estimates and smoke forecasts within the BCAPS system. Given the hierarchical framework of the health forecasting model, this is methodologically feasible [112]. The benefit would be to propagate the exposure uncertainty more realistically through to the final health indicator predictions, while still providing estimated $PM_{2.5}$ for external analyses. However, the computational challenges may not be easily addressed. In the medium-term, improved smoke forecasts can easily be integrated into the existing system, which would improve the performance of the health forecasts. This could include hybrid estimates that combine smoke forecasts with observed data [40], or ensemble estimates from multiple forecast models. For example, the FireWork smoke forecasting system is also operational across Canada [101], but it was excluded from these analyses because of technical concerns during the 2015 event in BC.

The research and evidence on interventions to mitigate the exposure of wildfire smoke are still very limited. The estimates shown are intended to provide relativistic information on the potential utility and impact of different scenarios when informed by the BCAPS surveillance system. In reality, no intervention would be so systematically adopted by a large population across space and time. More research on interventions will be key in this area, and having established surveillance systems that monitor health indicators, such as BCAMS, will facilitate the retrospective evaluation of interventions. When better information on the efficacy, uptake, and timeliness of interventions is available, the prospective BCAPS framework proposed here could be expanded to include information on the expected impacts of different scenarios. Future work should also more carefully explore the thresholds at which interventions would be triggered and the sensitivity and specificity of possible thresholds. Different public health settings may have different tolerances for false alerts (higher sensitivity) or for failure to alert (higher specificity), and these decisions should be made within the context of the cost and effectiveness of available interventions.

Wildfire smoke is an increasingly important public health challenge, and environmental surveillance research must continue to facilitate evidence-based decision making to mitigate the harmful effects of smoke on populations. A multi-pronged approach is needed to improve the science on all sides, from exposure assessment and epidemiologic studies to methodological advances in surveillance across space and time and research assessing the effectiveness of potential interventions.

5.9 Appendix: Sensitivity analyses of interventions

There was an estimated excess of 1785 salbutamol dispensations during the high exposure period of July 5 to 10. If a moderately effective (40%) intervention had been implemented on the morning of July 6, approximately 38% of excess dispensations could have been prevented. To achieve the same reduction in excess dispensations for an intervention that did not begin until the July 7, the intervention would have to be extremely effective (90%). If action had not been taken until July 8, an intervention with 90% effectiveness could have reduced excess dispensations by only 23%. Similarly, there was an estimated excess of 908 physician visits during high exposure period. Because physician visits had a higher proportion of excess counts occurring on Sunday July 5, the potential impact of intervening on or after July 6 was lower than for dispensations. If a moderately effective (40%) intervention had been implemented on the morning of July 6, approximately 22% of excess cumulative physician visits could have been prevented. To achieve the same reduction in excess dispensations for an intervention that did not begin until July 7, the intervention would have to be highly effective (60-70%). If action had not been taken until July 8, a very effective intervention (90%) could have reduced excess dispensations by only 13%. For both outcomes, intervening on or after Thursday July 9 would have had a negligible reduction in excess counts.

Table 5–3: Dispensations: Percentage of excess indicator counts by effectiveness and date of intervention.

	10%	20%	30%	40%	50%	60%	70%	80%	90%
Date of intervention									
Sun, July 5	9.9	19.8	29.7	39.6	49.6	59.4	69.3	79.2	89.1
Mon, July 6	5.4	10.9	16.3	21.8	27.2	32.6	38.1	43.5	49
Tue, July 7	3.4	6.8	10.4	13.8	17.2	20.6	24	27.5	30.9
Wed, July 8	1.4	2.8	4.2	5.6	7	8.4	9.8	11.2	12.6
Thu, July 9	0.6	1.1	1.5	2.1	2.6	3.2	3.7	4.2	4.7
Fri, July 10	0	0	0	0	0	0	0	0	0

Table 5–4: Physician visits: percentage of excess indicator counts by effectiveness and date of intervention.

	10%	20%	30%	40%	50%	60%	70%	80%	90%
Date of intervention									
Sun, July 5	10	20.1	30.1	40.2	50.3	60.3	70.4	80.4	90.5
Mon, July 6	9.4	18.9	28.3	37.8	47.2	56.6	66.1	75.5	84.9
Tue, July 7	4.3	8.6	12.9	17.3	21.6	25.8	30.1	34.5	38.8
Wed, July 8	2.6	5.2	7.8	10.4	12.9	15.5	18.1	20.7	23.3
Thu, July 9	0.7	1.4	2.1	2.9	3.6	4.3	5	5.7	6.4
Fri, July 10	0	0	0	0	0	0	0	0	0

6 Discussion

6.1 Motivation

The changing global climate is having direct and indirect impacts on wildfire seasons. In BC, longer, dryer summers have led to fires of increasing size, frequency, and severity [22]. Warmer winters in the early 1990's resulted in an epidemic of mountain pine beetles, severely damaging the provincial forests and creating a high risk environment for large wildfires [142]. While wildfires only directly threaten a small number of individuals each season, wildfire smoke can travel hundreds or even thousands of kilometres away from the fire origin. Consequently, wildfire smoke is an increasingly important issue in public health.

The epidemiologic research on wildfire smoke has shown that, like exposure to other sources of $PM_{2.5}$, acute exposures to high levels of smoke are strongly associated with increases in respiratory health effects [9, 3]. When air quality is poor due to wildfire smoke, public health systems are called upon to make recommendations for public safety. The limited evidence necessary to make such recommendations has been repeatedly identified as a challenge in the academic and government literature [42, 7, 65]. There have been calls for more research on the effectiveness of interventions and on the need for surveillance to inform evidence-based decision making.

In areas seasonally affected by wildfire smoke, such as BC, surveillance programs to monitor the public health response to smoky conditions are being developed. Monitoring population response to acute, extreme, and widespread air pollution events requires timely data collection and analysis. The data that are available and appropriate for surveillance purposes are generally noisy, such as healthcare utilization data, and error-prone, such as discrete air pollution monitors across a large geographic area.

6.2 Summary of contributions

The goal of this thesis was to address the methodological research gap for the development of a cohesive surveillance framework for acute environmental exposures. Through a series of three manuscripts, I sought to propose, apply, and evaluate a methodological framework that integrates multiple data sources and forecasts the anticipated health effects of observed or expected air pollution concentrations. Wildfire smoke surveillance was used as a motivating case study and application area throughout the work. Many of the research findings may be relevant to surveillance of other types of acute environmental exposures, such as water contamination or extreme heat events.

In the first manuscript (Chapter 3), I proposed methods to forecast two health outcomes against a common latent exposure. My objectives were to 1) identify a statistical approach that would be appropriate for the characteristics of the exposure and outcome data typically encountered in EPHS, 2) be flexible so that new exposure or outcome data could be readily accommodated, and 3) address the computational constraints of requiring an 'online' daily surveillance system. Given the difficulty of measuring air pollution across large geographic areas and using those data to assign population exposure estimates, some amount of error is expected. The latent process model I proposed assumed some non-differential measurement error for the exposure ($PM_{2.5}$). Additionally, my proposed model was bivariate, with two health indicators modelled and forecasted simultaneously against this common exposure. This allowed consideration of the correlation between the two outcomes, which improved forecast performance. In evaluating the bivariate model, I found that it was more statistically stable in terms of estimating reasonable effect sizes, and appeared to be equivalent or superior to the univariate models in terms of forecast accuracy and precision. The multivariate structure obtained by mutual dependency on the latent process was statistically simpler than placing the multivariate structure on, for example, the parameters between the exposure and the outcomes. However, the proposed model was not computationally efficient enough to use readily available Marcov chain Monte Carlo (MCMC) software for implementation. Therefore, I demonstrated how the proposed model could be implemented using integrated nested Laplace approximations (INLA), a new approach to approximate Bayesian inference. The final bivariate model could be extended to include, for example, additional outcomes, additional predictors, more sophisticated exposure models, and spatial smoothing.

In the second manuscript (Chapter 4), I explored the challenges of using an EPHS system across an entire state or province, when data will likely be aggregated by some geographic unit.

Some communities may have very small absolute populations, but the public health concern of wildfire smoke exposure may still be significant, given that entire populations can be exposed. Many forecasting models will suffer from unstable prediction variance when populations and health outcome counts are consistently small across a time series, severely limiting or even eliminating the utility of a forecasting approach to surveillance. While it is possible to group the geographic units to address this challenge, information will be lost in the aggregation, and less precise information will be gleaned from the surveillance system. As an alternative, I added spatial smoothing to the previously proposed bivariate latent model. This stabilized the prediction variance in smaller regions, at the cost of a small loss of accuracy. In regions with larger populations, the smoothing was generally not beneficial or necessary. The decision of whether to include spatial smoothing can be made within the context of the data characteristics for a set of administrative regions, such as the population sizes, the geographic sizes, the distribution of the population at risk, and the spatial precision of the exposure measurements available.

In the third manuscript (Chapter 5), I described how the model proposed in the Chapter 3 could be used as a real-time surveillance system for public health decision making. I did this in the form of a retrospective case study in southwestern British Columbia, which was affected by an extreme smoke event in July 2015. The BCAPS system was used to produce daily 24-hour and 48-hour forecasts from July 3 to July 10, as if the system were being used in real-time by public health authorities during the event. The BCAPS system could have prevented twice as many excess health indicator counts had it been used to guide public health interventions, as compared to the existing BCAMS system. I found that the performance of the health indicators forecasts was dependent on the performance of the smoke exposure forecasts (BlueSky). In brief, forecasted health indicators counts were low when the smoke forecasting system did not predict the occurrence of a smoke event. Syndromic surveillance data, such as the health indicators of salbutamol dispensations and physician visits used in this thesis, tend to be noisy and contain artefacts, such as day of the week effects. The proposed modelling framework seeks to reduce excessive false 'alerts' in the predicted health indicators by smoothing the two time series towards each other and towards the historical mean. This is beneficial to avoid sudden increases in predicted indicators due to isolated, chance variations in the exposure. However, the disadvantage of this more conservative smoothing approach is that the model is slower to adjust to sudden aberrations in the exposures or outcomes. I found that the prediction interval distributions were fairly useful for capturing the true future observed value. If the methodological framework proposed here were integrated into a decision making framework, an upper quantile of the prediction intervals could be used to calibrate the signals based on desired levels of sensitivity and specificity.

6.3 Challenges and opportunities for further research

6.3.1 Modelling and computation

There are many ways to model and forecast time series data. In this thesis, my approach was to propose a methodological framework that started with a simple model, but was flexible enough to accommodate changes in the characteristics of the exposure or outcome data. For example, the bivariate approach has the benefit of borrowing strength from two correlated outcomes, and allows additional outcomes to be included as they become available. The latent exposure approach assumes some non-differential measurement error in the air pollution estimates, but can be made more sophisticated as advances in exposure modelling occur. The spatial smoothing described in Chapter 4 is another example of the benefit of flexible modelling approaches, so that the proposed methodological framework can be used even in areas with small populations. The focus on simplicity and flexibility was an ongoing trade-off between complexity and computational constraints. In Chapter 5, providing a single set of forecasts per day only required approximately 15 minutes to run. However, even with the approximate Bayesian inference implementation, computation was a challenge in evaluation of the proposed approach in Chapters 3 and 4, because the models needed to be re-fit for each forecast day in the simulation or in the forecast horizon to obtain summary measures, such as mean absolute percentage errors (MAPE).

The computation time for producing a set of forecasts depended on the size of the forecast window, which is the amount of 'training data' included in the model. It took approximately 15 minutes to produce one day of forecasts for one geographic region in Chapter 3 and 4. In Chapter 3, there were 27 simulation scenarios each of which were run 200 times, taking approximately two days per scenario. Because of these computational constraints, the simulations from Chapter 3 were fairly simple, primarily to assess the impact of the correlation parameter on the performance, and only 200 simulations were run per scenario. In the Chapter 4, the analyses replicated for each of the 89 geographic regions in the study area, first for the temporal model and then for the spatio-temporal model. The spatio-temporal model was actually more computationally efficient, because fewer parameters were estimated, which is also a limitation of spatial smoothing using a conditional autoregressive random effect, because it assumes a common association between the exposure and health indicators throughout the entire geographic region. There are alternatives that would allow for unique relative risks between the exposure and outcome in each geographic region while still providing spatial smoothing, such as spatially varying coefficient models [143]. However, exploratory analyses found this approach to be computationally unfeasible in the British Columbia example used here.

The analyses for Chapter 4 were implemented using a high performance computing system at the University of Bath, making use of a number of high memory nodes (including two nodes with 512GB of memory). Even so, an initial approach using the default method for selecting initial values for the optimization for each day independently proved to be computationally demanding, as the Newton-Raphson optimizer had to restart to produce the forecast each day when a new model was fit. Using the posterior values from the previous day to initialize the optimization for upcoming forecast provided a solution for days on which convergence was not achieved. Overall, allowing this learning between days resulted in marked improvements in efficiency. Future work could aim to take further advantage of parallel computing and evaluate the performance of the forecasting models in greater detail, such as the impact of the correlation between the two health indicators.

6.3.2 Exposure assessment

I used modelled estimates of air pollution estimated by OSSEM throughout this thesis, as they provided better information and at a more precise spatial scale than simply assuming air quality monitors were representative of a large geographic area [37, 35]. The BCAPS system assumed that these measurements were made with non-differential measurement error but were otherwise correct. This is an overly simplistic approach, but it made a first step towards accounting for the measurement error challenges in this type of research. Ideally, the exposure and health indicator modelling would be performed within the same hierarchical model, which would more honestly propagate the error throughout the forecasts. Bayesian inference using INLA would also facilitate use of their stochastic partial differential equation (SPDE) approach to estimate a continuous latent field measured with error over a time at a large number of point-reference locations [144]. The latent field does not have the same natural Gaussian Markov random field (GMRF) structure as would be expected when aggregate spatial data are being considered. The SPDE approach instead starts with a Gaussian field and induces a GMRF over a continuous domain to which the usual INLA assumptions can be applied, as well as the assumption that the field must have a Matérn covariance structure [144]. The usual approach to computation for this type of Gaussian field model would be impossible, but making the assumptions required to model the latent field as a GMRF allows the INLA approach to be used for approximate inference. This approach has been used to estimate air pollution at a fine spatial scale across large geographic areas [145, 146], and integrating this approach into the forecasting of health indicators would be an exciting avenue for further research.

In the short term, however, integrating daily exposure modelling using SPDEs would raise additional computational challenges, and may not be practically feasible. In the meantime, research on exposure modelling is ongoing, and it was recently reported that combining available exposure measurements for wildfire smoke can lead to better overall accuracy than relying on a single exposure model [40].

6.3.3 Monitoring health indicators

Two respiratory health indicators are currently available for surveillance of wildfire smoke in BC, and this may be expanded to include additional health outcomes as they become available. The practical issues around securing timely collection of and access to health data will need to be addressed for the use of this system within the BCCDC and for any public health setting that may wish to apply such methods. Based on exploratory data analyses and the desire to reduce model complexity where possible, I used a first-order autoregressive model for the exposure, with no lag on the exposure. As additional data become available, or if my approach were to be applied in different settings, further exploration of these assumptions should be performed. Computational challenges may arise, but they should be surmountable given the approximate approach to Bayesian inference using INLA, and the availability of parallel computing systems.

6.3.4 Interventions research gap

Perhaps the most urgent research gap in public health and wildfire smoke is on the effectiveness of potentially protective interventions. Stakeholders have expressed the ongoing need for clear evidence about different interventions with respect to efficacy, uptake, and timeliness. The ongoing collection of surveillance data will facilitate the retrospective evaluation of interventions such as public health messaging, deployment of portable air cleaners [131], and evacuations under the most extreme conditions [68]. Future work could potentially take advantage of the quasi-experiment designs that would exploit the natural experiments that occur with smoky conditions, such as pre-post designs, difference-in-differences, or interrupted time series [147, 148]. Until more research clearly shows the best route forward, it will be difficult for public health authorities to guide definitive decision making.

6.4 Conclusion

As the earth warms and land use patterns shift, harmful environmental exposures such as heat waves and episodes of wildfire smoke are becoming more frequent and severe. At this time EPHS is an emerging field, and even the most innovative air pollution surveillance systems use separate data streams to track exposures and outcomes. There is a pressing need to integrate these data to better inform public health practice through more comprehensive forecasting, alerting, and decision making tools. While there have been repeated calls for improvement of EPHS systems, progress requires collaboration between those embedded in public health practice and methodologists with the necessary expertise to extract maximum benefit from large, noisy, and correlated datasets.

Given the multidisciplinary nature of this topic, my work was supported by three collaborators who each played a crucial role in the conception and execution of this dissertation. Dr. David Buckeridge is the director of the Surveillance Laboratory at McGill University, and his research program is focused on methodological development and evaluation of applied disease surveillance systems. Dr. Gavin Shaddick at the University of Bath is a Bayesian statistician who has published broadly in environmental epidemiology and spatial-temporal statistics. Dr. Sarah Henderson a senior scientist at the BCCDC, and she oversees the development and implementation of a wide range of applied EPHS systems, including the BCAMS system for wildfire smoke. For my dissertation, I sought to outline a methodological framework for wildfire smoke surveillance that is computationally feasible, exploits the characteristics of the available data, and can be easily modified and adapted to her applications as research in this area progresses. Drs. Buckeridge and Shaddick supervised and consulted on the methodological development and computational challenges. Dr. Henderson supervised and consulted on the practical aspects of applying such a framework within the applied public health setting.

Dr. Henderson intends to incorporate the framework proposed in this thesis into the existing BCAMS surveillance system for wildfire smoke in time for the 2017 wildfire season. As well, I will develop the modelling framework into an open-source R package with visualization tools, making it more user-friendly for those engaged in environmental public health surveillance.

References

[1] Fay H Johnston, Sarah B Henderson, Yang Chen, James T Randerson, Miriam Marlier, Ruth S DeFries, Patrick Kinney, David MJS Bowman, and Michael Brauer. Estimated global mortality attributable to smoke from landscape fires. *Environmental Health Perspectives*, 120(5):695–701, 2012.

[2] Josephine Malilay. A review of factors affecting the human health impacts of air pollutants from forest fires. Technical report, National Centre for Environmental Health, Centres for Disease Control and Prevention, 1998.

[3] Jia C Liu, Gavin Pereira, Sarah A Uhl, Mercedes A Bravo, and Michelle L Bell. A systematic review of the physical health impacts from non-occupational exposure to wildfire smoke. *Environmental Research*, 136:120–132, 2015.

[4] Colleen E Reid, Michael Brauer, Fay H Johnston, Michael Jerrett, John R Balmes, and Catherine T Elliott. Critical review of health impacts of wildfire smoke exposure. *Environmental Health Perspectives*, 124(9):1334–1343, 2016.

[5] Yongqiang Liu, Scott Goodrick, and Warren Heilman. Wildland fire emissions, carbon, and climate: Wildfire–climate interactions. *Forest Ecology and Management*, 317:80–96, 2014.

[6] Stephen B Thacker, Donna F Stroup, R Gibson Parrish, and Henry A Anderson. Surveillance in environmental public health: issues, systems, and sources. *American Journal of Public Health*, 86(5):633–638, 1996.

[7] Kathryn Morrison, Jeffrey Johnson, Jat Sandhu, and Karen Rideout. Evidence review: Health surveillance for wildfire smoke events. Technical report, British Columbia Centre for Disease Control, Environmental Health Services, 2014.

[8] Olorunfemi Adetona, Timothy E Reinhardt, Joe Domitrovich, George Broyles, Anna M Adetona, Michael T Kleinman, Roger D Ottmar, and Luke P Naeher. Review of the health effects of wildland fire smoke on wildland firefighters and the public. *Inhalation Toxicology*, 28(3):95–139, 2016.

[9] Luke P Naeher, Michael Brauer, Michael Lipsett, Judith T Zelikoff, Christopher D Simpson, Jane Q Koenig, and Kirk R Smith. Woodsmoke health effects: a review. *Inhalation Toxicology*, 19(1):67–106, 2007.

[10] Amir Sapkota, J Morel Symons, Jan Kleissl, Lu Wang, Marc B Parlange, John Ondov, Patrick N Breysse, Gregory B Diette, Peyton A Eggleston, and Timothy J Buckley. Impact of the 2002 Canadian forest fires on particulate matter air quality in Baltimore City. *Environmental Science & Technology*, 39(1):24–32, 2005.

[11] Martine Dennekamp and Michael J Abramson. The effects of bushfire smoke on respiratory health. *Respirology*, 16(2):198–209, 2011.

[12] Norbert Englert. Fine particles and human health - a review of epidemiological studies. *Toxicology Letters*, 149(1):235–242, 2004.

[13] Jane Q Koenig and William E Pierson. Air pollutants and the respiratory system: toxicity and pharmacologic interventions. *Clinical Toxicology*, 29(3):401–411, 1991.

[14] Darold E Ward and Colin C Hardy. Smoke emissions from wildland fires. *Environment International*, 17(2):117–134, 1991.

[15] Natalie M Suzuki. Particulate matter in BC: A report on PM_{10} and $PM_{2.5}$ mass concentrations up to 2000. Technical report, BC Ministry of Water, Land and Air Protection & Environment Canada, 2003.

[16] Guido R van der Werf, James T Randerson, Louis Giglio, GJ Collatz, M Mu, Prasad S Kasibhatla, Douglas C Morton, RS DeFries, Y van Jin, and Thijs T van Leeuwen. Global fire emissions and the contribution of deforestation, savanna, forest, agricultural, and peat fires (1997–2009). *Atmospheric Chemistry and Physics*, 10(23):11707–11735, 2010.

[17] Sarah B Henderson and Fay H Johnston. Measures of forest fire smoke exposure and their associations with respiratory health outcomes. *Current Opinion in Allergy and Clinical Immunology*, 12(3):221–227, 2012.

[18] Fabienne Reisen, Sandra M Duran, Mike Flannigan, Catherine Elliott, and Karen Rideout. Wildfire smoke and public health risk. *International Journal of Wildland Fire*, 24(8):1029–1044, 2015.

[19] Christopher McLeod, Terry Bogyo, Paul Demers, Demet Edeer, Clyde Hertzman, Susan Kennedy, Mieke Koehoorn, Kimberlyn McGrail, and Lillian Tamburic. Asthma in British Columbia. Technical report, Centre for Health Services and Policy Research, University of British Columbia, 2007.

[20] Statistics Canada. CANSIM table 105–0501. Health indicator profile, annual estimates, by age group and sex, Canada, provinces, territories, health regions (2011 boundaries) and peer groups, occasional [data file]. Ottawa (ON): Statistics Canada.

[21] World Health Organization. Vegetation fires, technical hazard sheet, natural disaster profiles. Technical report, World Health Organization, June 2014.

[22] MD Flannigan, Brian J Stocks, and BM Wotton. Climate change and forest fires. *Science of the Total Environment*, 262(3):221–229, 2000.

[23] Anthony L Westerling, Hugo G Hidalgo, Daniel R Cayan, and Thomas W Swetnam. Warming and earlier spring increase Western US forest wildfire activity. *Science*, 313(5789):940–943, 2006.

[24] Leslie A Richardson, Patricia A Champ, and John B Loomis. The hidden cost of wildfires: economic valuation of health effects of wildfire smoke exposure in southern California. *Journal of Forest Economics*, 18(1):14–35, 2012.

[25] Michael J Daniels, Francesca Dominici, Jonathan M Samet, and Scott L Zeger. Estimating particulate matter-mortality dose-response curves and threshold levels: an analysis of daily time-series for the 20 largest US cities. *American Journal of Epidemiology*, 152(5):397–406, 2000.

[26] Evaluation Directorate. Evaluation of the surveillance function at the public health agency of Canada. Technical report, Public Health Agency of Canada & Health Canada, January 2014.

[27] Andrew B Lawson. Spatial and spatio-temporal disease analysis. In Andrew B Lawson and Ken Kleinman, editors, *Spatial and Syndromic Surveillance for Public Health*, pages 55–76. Wiley, 2005.

[28] Steven F Magruder, S Happel Lewis, A Najmi, and E Florio. Progress in understanding and using over-the-counter pharmaceuticals for syndromic surveillance. *Morbidity and Mortality Weekly Report*, S53:117–122, 2004.

[29] Loïc Josseran, Anne Fouillet, Nadège Caillère, Dominique Brun-Ney, Danièle Ilef, Gilles Brucker, Helena Medeiros, and Pascal Astagneau. Assessment of a syndromic surveillance system based on morbidity data: results from the Oscour® Network during a heat wave. *PloS one*, 5(8):e11984, 2010.

[30] Herman A. Carneiro and Eleftherios Mylonakis. Google trends: a web-based tool for real-time surveillance of disease outbreaks. *Clinical Infectious Diseases*, 49(10):1557–1564, 2009.

[31] Kelly J Henning. What is syndromic surveillance? *Morbidity and Mortality Weekly Report*, S53:5–11, 2004.

[32] Dena M Bravata, Kathryn M McDonald, Wendy M Smith, Chara Rydzak, Herbert Szeto, David L Buckeridge, Corinna Haberland, and Douglas K Owens. Systematic review: surveillance systems for early detection of bioterrorism-related diseases. *Annals of Internal Medicine*, 140(11):910–922, 2004.

[33] Joseph S Lombardo and David L Buckeridge. *Disease Surveillance: A Public Health Informatics Approach*. John Wiley & Sons, 2012.

[34] Howard Frumkin, Jeremy Hess, George Luber, Josephine Malilay, and Michael McGeehin. Climate change: The public health response. *American Journal of Public Health*, 98(3):435–445, 2008.

[35] Jiayun Yao, Jeff Eyamie, and Sarah B Henderson. Evaluation of a spatially resolved forest fire smoke model for population-based epidemiologic exposure assessment. *Journal of Exposure Science and Environmental Epidemiology*, 26:233—240, 2014.

[36] Pavlos S Kanaroglou, Michael Jerrett, Jason Morrison, Bernardo Beckerman, M Altaf Arain, Nicolas L Gilbert, and Jeffrey R Brook. Establishing an air pollution monitoring network for intra-urban population exposure assessment: a location-allocation approach. *Atmospheric Environment*, 39(13):2399–2409, 2005.

[37] Jiayun Yao and Sarah B Henderson. An empirical model to estimate daily forest fire smoke exposure over a large geographic area using air quality, meteorological, and remote sensing data. *Journal of Exposure Science and Environmental Epidemiology*, 3:328–335, 2014.

[38] Dana Coe Sullivan, Sean M Raffuse, Daniel A Pryden, Kenneth J Craig, Stephen B Reid, Neil JM Wheeler, Lyle R Chinkin, Narasimhan K Larkin, Robert Solomon, and Tara Strand. Development and applications of systems for modeling emissions and smoke from fires: the BlueSky smoke modeling framework and SMARTFIRE. In *17th International Emissions Inventory Conference*, pages 2–5, 2008.

[39] Narasimhan K Larkin, Susan M O'Neill, Robert Solomon, Sean Raffuse, Tara Strand, Dana C Sullivan, Candace Krull, Miriam Rorig, Janice Peterson, and Sue A Ferguson. The BlueSky smoke modeling framework. *International Journal of Wildland Fire*, 18(8):906–920, 2010.

[40] Weiran Yuchi, Jiayun Yao, Kathleen E McLean, Roland Stull, Radenko Pavlovic, Didier Davignon, Michael D Moran, and Sarah B Henderson. Blending forest fire smoke forecasts with observed data can improve their utility for public health applications. *Atmospheric Environment*, 145:308–317, 2016.

[41] Ana G Rappold, Susan L Stone, Wayne E Cascio, Lucas M Neas, Vasu J Kilaru, Martha Sue Carraway, James J Szykman, Amy Ising, William E Cleve, and John T Meredith. Peat bog wildfire smoke exposure in rural North Carolina is associated with cardiopulmonary emergency department visits assessed through syndromic surveillance. *Environmental Health Perspectives*, 119(10):1415, 2011.

[42] Kathleen E McLean, Jiayun Yao, and Sarah B Henderson. An evaluation of the British Columbia Asthma Monitoring System (BCAMS) and $PM_{2.5}$ exposure metrics during the 2014 forest fire feason. *International Journal of Environmental Research and Public Health*, 12(6):6710–6724, 2015.

[43] Valerie A Cullum, JB Farmer, D Jack, and GP Levy. Salbutamol: a new, selective β-adrenoceptive receptor stimulant. *British Journal of Pharmacology*, 35(1):141–151, 1969.

[44] Catherine T Elliott, Sarah B Henderson, and Victoria Wan. Time series analysis of fine particulate matter and asthma reliever dispensations in populations affected by forest fires. *Environmental Health*, 12(11):1–9, 2013.

[45] World Health Organization. International classification of diseases, 9th revision., 1975.

[46] L MacDonald and M Naus. Behind the scenes in public health: Adverse events following immunization AEFI signal investigation in British Columbia. *Canada Communicable Disease Report*, 40(S3):24–30, 2014.

[47] Ben Y Reis and Kenneth D Mandl. Time series modeling for syndromic surveillance. *BMC Medical Informatics and Decision Making*, 3(1):1, 2003.

[48] Ruth Salway, Duncan Lee, Gavin Shaddick, and Stephen Walker. Bayesian latent variable modelling in studies of air pollution and health. *Statistics in Medicine*, 29(26):2732–2742, 2010.

[49] Henry Rolka, Howard Burkom, Gregory F Cooper, Martin Kulldorff, David Madigan, and Weng-Keen Wong. Issues in applied statistics for public health bioterrorism surveillance using multiple data streams: research needs. *Statistics in Medicine*, 26(8):1834–1856, 2007.

[50] Joseph Lombardo, Howard Burkom, Eugene Elbert, Steven Magruder, Sheryl Happel Lewis, Wayne Loschen, James Sari, Carol Sniegoski, Richard Wojcik, and Julie Pavlin. A systems overview of the electronic surveillance system for the early notification of community-based epidemics (ESSENCE II). *Journal of Urban Health*, 80(1):32–42, 2003.

[51] Artur Dubrawski. Detection of events in multiple streams of surveillance data. In *Infectious Disease Informatics and Biosurveillance*, pages 145–171. Springer, 2011.

[52] Francesca Dominici, Antonella Zanobetti, Scott L Zeger, Joel Schwartz, and Jonathan M Samet. Hierarchical bivariate time series models: a combined analysis of the effects of particulate matter on morbidity and mortality. *Biostatistics*, 5(3):341–360, 2004.

[53] Håvard Rue, Sara Martino, and Nicolas Chopin. Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B*, 71(2):319–392, 2009.

[54] Thiago G Martins, Daniel Simpson, Finn Lindgren, and Håvard Rue. Bayesian computing with INLA: new features. *Computational Statistics & Data Analysis*, 67:68–83, 2013.

[55] Sara Martino and Håvard Rue. Case studies in Bayesian computation using INLA. In Piercesare Secchi Pietro Mantovan, editor, *Complex Data Modeling and Computationally Intensive Statistical Methods*, pages 99–114. Springer, 2010.

[56] Leonhard Held, Birgit Schrödle, and Håvard Rue. Posterior and cross-validatory predictive checks: a comparison of MCMC and INLA. In *Statistical Modelling and Regression Structures*, pages 91–110. Springer, 2010.

[57] Marta Blangiardo, Michela Cameletti, Gianluca Baio, and Håvard Rue. Spatial and spatio-temporal models with R-INLA. *Spatial and Spatio-temporal Epidemiology*, 7:39–55, 2013.

[58] Havard Rue and Leonhard Held. *Gaussian Markov Random Fields: Theory and Applications*. CRC Press, 2005.

[59] Matthew A Cole and Eric Neumayer. Examining the impact of demographic factors on air pollution. *Population and Environment*, 26(1):5–21, 2004.

[60] Christopher M Long and Jeremy A Sarnat. Indoor-outdoor relationships and infiltration behavior of elemental components of outdoor $PM_{2.5}$ for Boston-area homes. *Aerosol Science and Technology*, 38(S2):91–104, 2004.

[61] Joshua A Mott, Pamela Meyer, David Mannino, Stephen C Redd, Eva M Smith, Carol Gotway-Crawford, and Emmett Chase. Wildland forest fire smoke: health effects and intervention evaluation, Hoopa, California, 1999. *Western Journal of Medicine*, 176(3):157–162, 2002.

[62] David E Henderson, Jana B Milford, and Shelly L Miller. Prescribed burns and wildfires in colorado: impacts of mitigation measures on indoor air particulate matter. *Journal of the Air & Waste Management Association*, 55(10):1516–1526, 2005.

[63] Jennifer L Beverly and Peter Bothwell. Wildfire evacuations in Canada 1980–2007. *Natural Hazards*, 59(1):571–596, 2011.

[64] Joanne Stares, Rick Erland, Peter Mazey, and Karen Rideout. Evidence review of use of evacuation to protect public health during wildfire smoke events. Technical report, Environmental Health Services, British Columbia Centre for Disease Control, 2014.

[65] Sandra Durán, Fabienne Reisen, and Karen Rideout. Evidence review of wildfire smoke and public health risk. Technical report, Environmental Health Services, British Columbia Centre for Disease Control, 2014.

[66] Sarah B Henderson, Michael Brauer, Ying C MacNab, and Susan M Kennedy. Three measures of forest fire smoke exposure and their associations with respiratory and cardiovascular health outcomes in a population-based cohort. *Environmental Health Perspectives*, 119(9):1266, 2011.

[67] Philippe Duclos, Lee M Sanderson, and Michael Lipsett. The 1987 forest fire disaster in California: assessment of emergency room visits. *Archives of Environmental Health*, 45(1):53–58, 1990.

[68] Nikolas Krstic and Sarah B Henderson. Use of MODIS data to assess atmospheric aerosol before, during, and after community evacuations related to wildfire smoke. *Remote Sensing of Environment*, 166:1–7, 2015.

[69] Carles Barcelo-Vidal, Lucıa Aguilar, and Josep Antoni Martín-Fernández. Compositional VARIMA time series. In *Compositional Data Analysis: Theory and Applications*, pages 87–101. John Wiley & Sons, 2011.

[70] Owen F Price, Grant J Williamson, Sarah B Henderson, Fay Johnston, and David MJS Bowman. The relationship between particulate pollution levels in Australian cities, meteorology, and landscape fire activity detected from MODIS hotspots. *PloS one*, 7(10):e47327, 2012.

[71] M Sofiev, R Vankevich, M Lotjonen, M Prank, V Petukhov, T Ermakova, J Koskinen, and J Kukkonen. An operational system for the assimilation of the satellite information on wild-land fires for the needs of air quality modelling and forecasting. *Atmospheric Chemistry and Physics*, 9(18):6833–6847, 2009.

[72] Joel Schwartz, C Spix, G Touloumi, L Bacharova, T Barumamdzadeh, A Le Tertre, T Piekarksi, A Ponce De Leon, A Pönkä, G Rossi, et al. Methodological issues in studies of air pollution and daily counts of deaths or hospital admissions. *Journal of Epidemiology and Community Health*, 50(S1):3–11, 1996.

[73] Leonhard Knorr-Held and Nicola G Best. A shared component model for detecting joint and selective clustering of two diseases. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 164(1):73–85, 2001.

[74] Francesca Dominici, Aidan McDermott, Scott L Zeger, and Jonathan M Samet. On the use of generalized additive models in time-series studies of air pollution and health. *American Journal of Epidemiology*, 156(3):193–203, 2002.

[75] BC Statistics. Population projections: LHA 23, 2006. Accessed: 2015-03-31.

[76] Robert S Pindyck and Daniel L Rubinfeld. *Econometric Models and Economic Forecasts*, volume 4. Irwin/McGraw-Hill Boston, 1998.

[77] Nesreen K Ahmed, Amir F Atiya, Neamat El Gayar, and Hisham El-Shishiny. An empirical comparison of machine learning models for time series forecasting. *Econometric Reviews*, 29(5-6):594–621, 2010.

[78] Gianluca Bontempi, Mauro Birattari, and Hugues Bersini. Local learning for iterated time series prediction. In *International Machine Learning Society*, pages 32–38, 1999.

[79] Neil R Ericsson. Parameter constancy, mean square forecast errors, and measuring forecast performance: an exposition, extensions, and illustration. *Journal of Policy Modeling*, 14(4):465–495, 1992.

[80] David J Lunn, Andrew Thomas, Nicky Best, and David Spiegelhalter. WinBUGS - a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and Computing*, 10(4):325–337, 2000.

[81] Steve Brooks, Andrew Gelman, Galin Jones, and Xiao-Li Meng. *Handbook of Markov Chain Monte Carlo*. CRC press, 2011.

[82] Andrew Gelman, John B Carlin, Hal S Stern, and Donald B Rubin. *Bayesian Data Analysis*, volume 2. Taylor & Francis, 2014.

[83] Martyn Plummer et al. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. In *Proceedings of the 3rd International Workshop on Distributed Statistical Computing*, volume 124, page 125. Vienna, 2003.

[84] David J Spiegelhalter, Andrew Thomas, and NG Best. Computation on Bayesian graphical models. In A. Dawid J. Bernardo, J. Berger and A. Smith, editors, *Fifth Valencia International Meeting on Bayesian Statistics*. Oxford University Press, 1996.

[85] Julian Besag, Jeremy York, and Annie Mollié. Bayesian image restoration, with two applications in spatial statistics. *Annals of the Institute of Statistical Mathematics*, 43(1):1–20, 1991.

[86] Gavin Shaddick and Jon Wakefield. Modelling daily multivariate pollutant data at multiple sites. *Journal of the Royal Statistical Society: Series C*, 51(3):351–372, 2002.

[87] Andrew Thomas, Nicky Best, Dave Lunn, Richard Arnold, and David Spiegelhalter. GeoBugs user manual. *Cambridge: Medical Research Council Biostatistics Unit*, 2004.

[88] H Rue, S Martino, and F Lindgren. The R-INLA project. Online: www.r-inla.org, 2012.

[89] Benjamin M Taylor and Peter J Diggle. INLA or MCMC? A tutorial and comparative evaluation for spatial prediction in log-Gaussian Cox processes. *Journal of Statistical Computation and Simulation*, 84(10):2266–2284, 2014.

[90] Birgit Schrödle and Leonhard Held. A primer on disease mapping and ecological regression using INLA. *Computational Statistics*, 26(2):241–258, 2011.

[91] Egil Ferkingstad, Håvard Rue, et al. Improving the INLA approach for approximate Bayesian inference for latent Gaussian models. *Electronic Journal of Statistics*, 9(2):2706–2731, 2015.

[92] M Paul, A Riebler, LM Bachmann, H Rue, and L Held. Bayesian bivariate meta-analysis of diagnostic test studies using integrated nested Laplace approximations. *Statistics in Medicine*, 29(12):1325–1339, 2010.

[93] Youyi Fong, Håvard Rue, and Jon Wakefield. Bayesian inference for generalized linear mixed models. *Biostatistics*, 11(3):1–16, 2010.

[94] R Carroll, AB Lawson, C Faes, RS Kirby, M Aregay, and K Watjou. Comparing INLA and OpenBUGS for hierarchical Poisson modeling in disease mapping. *Spatial and Spatio-temporal Epidemiology*, 14:45–54, 2015.

[95] Stefanie Muff, Andrea Riebler, Leonhard Held, Håvard Rue, and Philippe Saner. Bayesian analysis of measurement error models using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series C*, 64(2):231–252, 2015.

[96] Norbert Englert. Fine particles and human health - a review of epidemiological studies. *Toxicology Letters*, 149(1):235–242, 2004.

[97] Martin J Wooster, G Roberts, GLW Perry, and YJ Kaufman. Retrieval of biomass combustion rates and totals from fire radiative power observations: FRP derivation and calibration relationships between biomass consumption and fire radiative energy release. *Journal of Geophysical Research: Atmospheres*, 110(D24):1–24, 2005.

[98] Deborah Ashby. Bayesian statistics in medicine: a 25 year review. *Statistics in Medicine*, 25(21):3589–3631, 2006.

[99] Noel Cressie and Christopher K Wikle. *Statistics for Spatio-Temporal Data*. John Wiley & Sons, 2011.

[100] Joshua A Mott, David M Mannino, Clinton J Alverson, Andrew Kiyu, Jamilah Hashim, Tzesan Lee, Kenneth Falter, and Stephen C Redd. Cardiorespiratory hospitalizations associated with smoke exposure during the 1997 Southeast Asian forest fires. *International Journal of Hygiene and Environmental Health*, 208(1):75–85, 2005.

[101] Radenko Pavlovic, Jack Chen, Kerry Anderson, Michael D Moran, Paul-Andre Beaulieu, Didier Davignon, and Sophie Cousineau. The firework air quality forecast system with near-real-time biomass burning emissions: recent developments and evaluation of performance for the 2015 North American wildfire season. *Journal of the Air & Waste Management Association*, 66(9):819–841, 2016.

[102] Claudia Spix, H Ross Anderson, Joel Schwartz, Maria Angela Vigotti, Alain Letertre, Judith M Vonk, Giota Touloumi, Franck Balducci, Tomasz Piekarski, Ljuba Bacharova, et al. Short-term effects of air pollution on hospital admissions of respiratory diseases in Europe: a quantitative summary of APHEA study results. *Archives of Environmental Health*, 53(1):54–64, 1998.

[103] Annette Peters, Douglas W Dockery, James E Muller, and Murray A Mittleman. Increased particulate air pollution and the triggering of myocardial infarction. *Circulation*, 103(23):2810–2815, 2001.

[104] Gerard Hoek, Ranjini M Krishnan, Rob Beelen, Annette Peters, Bart Ostro, Bert Brunekreef, and Joel D Kaufman. Long-term air pollution exposure and cardio-respiratory mortality: a review. *Environmental Health*, 12(43):1–15, 2013.

[105] R Sari Kovats and Shakoor Hajat. Heat stress and public health: a critical review. *Annual Review of Public Health*, 29:41–55, 2008.

[106] James W Buehler, Ellen A Whitney, Donna Smith, Michael J Prietula, Sarah H Stanton, and Alexander P Isakov. Situational uses of syndromic surveillance. *Biosecurity and Bioterrorism*, 7(2):165–177, 2009.

[107] Ana G Rappold, Susan L Stone, Wayne E Cascio, Lucas M Neas, Vasu J Kilaru, Martha Sue Carraway, James J Szykman, Amy Ising, William E Cleve, John T Meredith, et al. Peat bog wildfire smoke exposure in rural North Carolina is associated with cardiopulmonary emergency department visits assessed through syndromic surveillance. *Environmental Health Perspectives*, 119(10):1415–1420, 2011.

[108] Manuela De Sario, K Katsouyanni, and P Michelozzi. Climate change, extreme weather events, air pollution and respiratory health in Europe. *European Respiratory Journal*, 42(3):826–843, 2013.

[109] Shakoor Hajat, Sotiris Vardoulakis, Clare Heaviside, and Bernd Eggen. Climate change effects on human health: projections of temperature-related mortality for the UK during the 2020s, 2050s and 2080s. *Journal of Epidemiology and Community Health*, 0:1–8, 2014.

[110] Michael Brauer, Markus Amann, Rick T Burnett, Aaron Cohen, Frank Dentener, Majid Ezzati, Sarah B Henderson, Michal Krzyzanowski, Randall V Martin, Rita Van Dingenen, et al. Exposure assessment for estimation of the global burden of disease attributable to outdoor air pollution. *Environmental Science & Technology*, 46(2):652–660, 2012.

[111] Pius Lee and Yang Liu. Preliminary evaluation of a regional atmospheric chemical data assimilation system for environmental surveillance. *International Journal of Environmental Research and Public Health*, 11(12):12795–12816, 2014.

[112] Kathryn T Morrison, Gavin Shaddick, Sarah B Henderson, and David L Buckeridge. A latent process model for forecasting multiple time series in environmental public health surveillance. *Statistics in Medicine*, 35(18):3033–3240, 2016.

[113] Sara Martino and Håvard Rue. Implementing approximate Bayesian inference using Integrated Nested Laplace Approximation: A manual for the INLA program. *Computational Statistics and Data Analysis*, 67:68–83, 2009.

[114] BC Statistics. Population estimates. Online: www.bcstats.gov.bc.ca/StatisticsBySubject/Demography, 2016.

[115] Nikisa S Jordan, Charles Ichoku, and Raymond M Hoff. Estimating smoke emissions over the US Southern Great Plains using MODIS fire radiative power and aerosol observations. *Atmospheric Environment*, 42(9):2007–2022, 2008.

[116] RC Mashru, VB Sutariya, MG Sankalia, and PP Parikh. Development and evaluation of fast-dissolving film of salbutamol sulphate. *Drug Development and Industrial Pharmacy*, 31(1):25–34, 2005.

[117] Steven G Morgan. Prescription drug expenditures and population demographics. *Health Services Research*, 41(2):411–428, 2006.

[118] David Moore, Ray Copes, Robert Fisk, Ruth Joy, Keith Chan, and Michael Brauer. Population health effects of air quality changes due to forest fires in British Columbia in 2003: Estimates from physician-visit billing data. *Canadian Journal of Public Health*, 97(2):105–108, 2006.

[119] Environment Canada. Historical climate data. Online: http://climate.weather.gc.ca/, 2016.

[120] JC Wakefield, NG Best, and L Waller. Bayesian approaches to disease mapping. In Paul Elliott, Jon C Wakefield, Nicola G Best, DJ Briggs, et al., editors, *Spatial Epidemiology: Methods and Applications*. Oxford University Press, 2000.

[121] Gavin Shaddick and James V Zidek. *Spatio-Temporal Methods in Environmental Epidemiology*. CRC Press, 2015.

[122] Peter McCullagh and John A Nelder. *Generalized Linear Models*, volume 37. CRC press, 1989.

[123] Douglas C Montgomery, Elizabeth A Peck, and G Geoffrey Vining. *Introduction to Linear Regression Analysis*. John Wiley & Sons, 2015.

[124] Government of BC Public Safety and Emergency Services. Wildfire statistics. Online: www2.gov.bc.ca/gov/content/safety/wildfire-status/wildfire-statistics, 2016.

[125] Chris Chatfield. Prediction intervals for time-series forecasting. In *Principles of Forecasting*, pages 475–494. Springer, 2001.

[126] Michael Jerrett, Altaf Arain, Pavlos Kanaroglou, Bernardo Beckerman, Dimitri Potoglou, Talar Sahsuvaroglu, Jason Morrison, and Chris Giovis. A review and evaluation of intraurban air pollution exposure models. *Journal of Exposure Science and Environmental Epidemiology*, 15(2):185–204, 2005.

[127] Paolo Zannetti. *Air Pollution Modeling: Theories, Computational Methods and Available Software*. Springer Science & Business Media, 2013.

[128] Government of British Columbia. Wildfire Season Summary, 2016.

[129] Jennifer R Marlon, Patrick J Bartlein, Megan K Walsh, Sandy P Harrison, Kendrick J Brown, Mary E Edwards, Philip E Higuera, M J Power, R S Anderson, C Briles, and Others. Wildfire responses to abrupt climate change in North America. *Proceedings of the National Academy of Sciences*, 106(8):2519–2524, 2009.

[130] Luisa V. Giles, Michael Brauer, Prabjit Barn, Nino Künzli, Isabelle Romieu, Murray A. Mittleman, Stephan van Eeden, Ryan Allen, Chris Carlsten, Dave Stieb, Curtis Noonan, Audrey Smargiassi, Joel D. Kaufman, Shakoor Hajat, Tom Kosatsky, and Michael Brauer. From good intentions to proven interventions: effectiveness of actions to reduce the health impacts of air pollution. *Environmental Health Perspectives*, 119(1):29–36, 2011.

[131] Prabjit K Barn, Catherine T Elliott, Ryan W Allen, Tom Kosatsky, Karen Rideout, and Sarah B Henderson. Portable air cleaners should be at the forefront of the public health response to landscape fire smoke. *Environmental Health*, doi.org/10.1186/s12940-016-0198-9, 2016.

[132] W J Fisk and W R Chan. Health benefits and costs of filtration interventions that reduce indoor exposure to $PM_{2.5}$ during wildfires. *Indoor Air*, 26(1):25–38, 2016.

[133] Ana G Rappold, Neal L Fann, James Crooks, Huang Jin, Wayne E Cascio, Robert B Devlin, and David Diaz-Sanchez. Forecast-Based Interventions Can Reduce the Health and Economic Burden of Wildfires. *Environmental Science & Technology*, 48(18):10571–10579, 2014.

[134] Statistics Canada. Health regions: boundaries and correspondence with census geography. Online: http://www.statcan.gc.ca/pub/82-402-x/82-402-x2013003-eng.htm, 2013.

[135] Catherine T Elliott, Sarah B Henderson, and Victoria Wan. Time series analysis of fine particulate matter and asthma reliever dispensations in populations affected by forest fires. *Environmental Health*, 12(1):1, 2013.

[136] Russell E Glasgow, Edward Lichtenstein, and Alfred C Marcus. Why don't we see more translation of health promotion research to practice? Rethinking the efficacy-to-effectiveness transition. *American Journal of Public Health*, 93(8):1261–1267, 2003.

[137] Linda Dix-Cooper, Fay Johnston, Paul Hasselback, and Karen Rideout. Evidence review of reducing time outdoors during wildfire smoke events. Technical report, Environmental Health Services, British Columbia Centre for Disease Control, 2014.

[138] Hind Sbihi, Mark Nicas, and Karen Rideout. Evidence review of using masks to protect public health during wildfire smoke events. Technical report, BC Centre for Disease Control, Vancouver BC, 2014.

[139] Seth L Foldy. Linking better surveillance to better outcomes. *Morbidity and Mortality Weekly Report*, S53:12–17, 2004.

[140] Joseph S Lombardo, Carol A Sniegoski, Wayne A Loschen, Matthew Westercamp, Michael Wade, Shandy Dearth, and Guoyan Zhang. Public health surveillance for mass gatherings. *Johns Hopkins APL Technical Digest*, 27(4):347–355, 2008.

[141] Farhad Salimi, Sarah B. Henderson, Geoffrey G. Morgan, Bin Jalaludin, and Fay H. Johnston. Ambient particulate matter, landscape fire smoke, and emergency ambulance dispatches in Sydney, Australia. *Environment International*, doi.org/10.1016/j.envint.2016.11.018, 2016.

[142] Michael A Wulder, Stephanie M Ortlepp, Joanne C White, Trisalyn Nelson, and Nicholas C Coops. A provincial and regional assessment of the mountain pine beetle epidemic in British Columbia: 1999–2008. *Journal of Environmental Informatics*, 15(1):1–13, 2010.

[143] Alan E Gelfand, Hyon-Jung Kim, CF Sirmans, and Sudipto Banerjee. Spatial modeling with spatially varying coefficient processes. *Journal of the American Statistical Association*, 98(462):387–396, 2003.

[144] Finn Lindgren, Håvard Rue, and Johan Lindström. An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498, 2011.

[145] Michela Cameletti, Finn Lindgren, Daniel Simpson, and Håvard Rue. Spatio-temporal modeling of particulate matter concentration through the SPDE approach. *Advances in Statistical Analysis*, 97(2):109–131, 2013.

[146] Gavin Shaddick, Haojie Yan, Ruth Salway, Danielle Vienneau, Daphne Kounali, and David Briggs. Large-scale Bayesian spatial modelling of air pollution for policy support. *Journal of Applied Statistics*, 40(4):777–794, 2013.

[147] Peter Craig, Cyrus Cooper, David Gunnell, Sally Haw, Kenny Lawson, Sally Macintyre, David Ogilvie, Mark Petticrew, Barney Reeves, Matt Sutton, et al. Using natural experiments

to evaluate population health interventions: new Medical Research Council guidance. *Journal of Epidemiology and Community Health*, 66(12):1182–1186, 2012.

[148] James Lopez Bernal, Steven Cummins, and Antonio Gasparrini. Interrupted time series regression for the evaluation of public health interventions: a tutorial. *International Journal of Epidemiology*, doi.org/10.1093/ije/dyw098, 2016.