




Contents lists available at ScienceDirect

Spatial Statistics

journal homepage: www.elsevier.com/locate/spasta

A spatially varying change points model for monitoring glaucoma progression using visual field data

Samuel I. Berchuck^a, Jean-Claude Mwanza^b,
Joshua L. Warren^{c,*}

^a Department of Statistical Science and Forge, Duke University, NC, USA

^b Department of Ophthalmology, University of North Carolina-Chapel Hill, NC, USA

^c Department of Biostatistics, Yale University, New Haven, CT, USA

ARTICLE INFO

Article history:

Received 27 November 2018

Accepted 14 February 2019

Available online 22 February 2019

Keywords:

Bayesian hierarchical models

Boundary detection

Multivariate conditional autoregressive model

Spatially varying change points

ABSTRACT

Glaucoma disease progression, as measured by visual field (VF) data, is often defined by periods of relative stability followed by an abrupt decrease in visual ability at some point in time. Determining the transition point of the disease trajectory to a more severe state is important clinically for disease management and for avoiding irreversible vision loss. Based on this, we present a unified statistical modeling framework that permits prediction of the timing and spatial location of future vision loss and informs clinical decisions regarding disease progression. The developed method incorporates anatomical information to create a biologically plausible data-generating model. We accomplish this by introducing a spatially varying coefficients model that includes spatially varying change points to detect structural shifts in both the mean and variance process of VF data across both space and time. The VF location-specific change point represents the underlying, and potentially censored, timing of true change in disease trajectory while a multivariate spatial boundary detection structure is introduced that accounts for the complex spatial connectivity of the VF and optic disc. We show that our method improves estimation and prediction of multiple aspects of disease management in comparison to existing methods through simulation and real data application. The R package *spCP* implements the new methodology.

© 2019 Elsevier B.V. All rights reserved.

* Correspondence to: Department of Biostatistics, Yale University, 60 College Street, Ste 213, New Haven, CT, 06510, USA.

E-mail address: joshua.warren@yale.edu (J.L. Warren).

1. Introduction

Glaucoma is an optic neuropathy that is the leading cause of irreversible vision loss worldwide (Tham et al., 2014). Once a patient is diagnosed, they are at risk of disease progression and are closely monitored by clinicians who must balance the potential of further irreversible deterioration of vision with the costs (financial and physical) of medical and surgical interventions. The visual ability of a glaucoma patient is most often quantified through visual field (VF) examinations, a psychophysical procedure that assesses a patient's field of vision. VF tests are performed during clinical visits over time, producing a time series of spatial referenced data, and the results are monitored for signs of glaucomatous VF progression. Determining if the disease is progressing remains one of the most challenging aspects of glaucoma management, as it is difficult to differentiate true progression from testing variability (Vianna and Chauhan, 2015).

Although VF testing is a common technique for assessing glaucoma progression, there is currently no reference standard approach for converting its output into a treatment decision. The majority of past studies can be grouped into two categories, global- and trend-based analyses. Global methods provide a single summary of the VF data at each visit, thus aggregating and potentially ignoring key information across VF locations (Heijl et al., 1987). These global summaries are then monitored over time to investigate the potential of disease progression. Trend methods analyze the individual time series of sensitivities at each VF location (Fitzke et al., 1996). These methods often resort to separately modeling the sensitivities over time assuming a linear relationship, often through use of simple linear regression (Katz et al., 1997). The inference obtained from these methods can be a poor diagnostic of progression or have inferior prediction accuracy, in comparison to non-linear models (Jampel et al., 2011), and suffers from low power due to ignoring the spatial nature of VF data or aggregating over space (Vianna and Chauhan, 2015).

In this paper, we develop an innovative method that utilizes spatially varying change points (CPs) to accurately model the VF sensitivities over time. Patients diagnosed with glaucoma are often monitored for years with only slow changes in visual functionality (Heijl et al., 2002). It is not until disease progression that notable vision loss occurs, and the deterioration is often swift (Jay and Murdoch, 1993). This disease course inspires a modeling framework that can identify a point of functional change in the course of follow-up. By utilizing the CP framework, the concept of disease progression becomes intrinsically parameterized into the model. Not only does the specification increase flexibility in modeling the trajectory of the VF time series by relaxing the linearity assumption, but the CP represents the point of functional change and can be used to make inferential statements about disease progression. This framework is particularly rich, however, because in addition to there being a functional change in the VF sensitivities upon the start of progression, there is also an increase in testing variance in areas of decreased sensitivity. This inverse variance relationship upon progression is well known in the glaucoma literature, although is often ignored in model development (Russell et al., 2012). We use a CP to segment the variance into a stable process that has the flexibility to increase after the point of progression. Finally, we focus on the joint modeling of all available VF data for a patient, accounting for the complex underlying spatial correlation, thereby increasing power to make improved predictions and treatment decisions.

The VF exhibits a complex spatial structure that is best modeled considering the anatomy of the eye. Naive spatial/neighborhood structures that assume neighboring VF locations share information similarly ignore this anatomy. Incorporating this information can aid in properly modeling this spatial dependency and in uncovering important features of disease progression across time. Berchuck et al. (2018) introduced a framework for inducing localized neighbors as a function of ocular anatomy using a spatiotemporal boundary detection method. Their method utilized a measure developed by Garway-Heath et al. (2000) that represents an estimate of the angle that each VF test location's underlying retinal ganglion cells (RGC) enter the optic disc. The Garway-Heath angles are used as a dissimilarity metric, a covariate that explains localized spatial structure. In the present paper, the methodology from Berchuck et al. (2018) is extended to the multivariate setting, such that the Garway-Heath angles dictate the neighborhood structure on the VF for multiple spatial parameters, including the spatially varying coefficients and CPs. To our knowledge, this is the first extension of boundary detection using a dissimilarity metric to the multivariate setting.

Our newly developed spatial CPs model with multivariate boundary detection structure has the ability to improve prediction of future VF sensitivities and assess progression for a glaucoma patient. Having the ability to accurately predict future deterioration of the VF will improve a clinician's ability to make timely, data-driven treatment adjustments to preserve vision. Furthermore, a latent CP process is introduced that allows us to spatially interpolate CP values at VF locations where visual damage has yet to occur (i.e., censored locations). This will allow clinicians to determine the location and timing of future vision loss.

This paper is outlined as follows. Section 2 introduces the VF data while Section 3 details past CP methods for spatially referenced data. In Section 4, our newly developed spatial CPs method is detailed and we develop a multivariate boundary detection spatial process for localized smoothing. We apply our method to a dataset of VF tests from glaucoma patients in Section 5 and present a simulation study in Section 6. We conclude in Section 7 with a discussion.

2. Visual field data

The visual ability of a glaucoma patient is most often measured through standard automated perimetry (SAP), an interactive technology that estimates the threshold sensitivity to light across a patient's field of vision. In this study, we analyze data acquired with the Humphrey Field Analyzer-II (HFA-II) (Carl Zeiss Meditec Inc., Dublin, CA), a machine that produces a visual acuity map across 54 testing locations on the VF. Two of these locations correspond to a natural blind spot corresponding to the optic disc, resulting in 52 informative points. During a test, the light is generated with increasing brightness ranging from approximately 40 decibels (dB) (excellent vision) to 0 dB (near blindness). The patient uses a hand-held trigger to indicate whether the light stimulus was detected and the machine reports the dimmest stimulus detected at each location. All stimuli that are not detected by 0 dB are censored due to the constraints of the machine, as values < 0 dB are theoretically possible. The resulting gridded spatial layout represents an estimate of a patient's visual ability across the entire field of vision at that point in time.

In the course of monitoring a glaucoma patient, VF testing is performed on a regular basis and a longitudinal series of VFs is obtained. In Fig. 1, the time series of VF observations is presented for a selected patient from our data as a motivating example for the CP framework. The majority of past modeling attempts have assumed a linear relationship between VF sensitivities and time. However, from the data in Fig. 1, it is clear that CPs may provide an improved description of the trajectory of VF sensitivities over time, and in particular for areas of decreased sensitivities. Furthermore, the figure illuminates the inverse relationship between VF sensitivities and testing variability that is well known but rarely modeled in practice (Russell et al., 2012).

In order to fully motivate the modeling framework, we provide some details on the anatomy underlying the data generating mechanism for VF data. The light stimulus from the HFA-II is absorbed by RGCs, which make up the retinal nerve fiber layer (RNFL) of the retina, that use photoreceptors to transmit information to the brain along their axons (Davson, 2012). The RGCs disperse across the RNFL and converge at the optic disc, where the RGC axons bundle to form the optic nerve. For glaucoma patients, an increase in intraocular pressure creates stress on the optic disc, which in turn causes RGC death, interrupting the visual pathway to the brain. In particular, damage to specific regions of the optic disc corresponds to loss of RGCs whose axons enter the damaged region.

This anatomy is important to understand when analyzing VF data, because vision is the brain's projection of information encoded by RGCs. Therefore, deterioration of the VF has a direct correspondence to damage of RGCs that originate from the same region of the retina. Because of this anatomy, correlation between two test points on the VF is dependent on the spatial proximity that their underlying RGCs enter the optic disc. This anatomical relationship was studied by Garway-Heath et al. (2000) who quantified the relationship between the VF and optic disc by estimating the angle that each test location's underlying RGC axons enters the optic disc, measured in degrees ($^\circ$). The measure ranges from 0 – 360° , where 0° is designated at the 9-o'clock position (right eye) and angles are counted counter clockwise. Throughout this analysis, we utilize these angles to properly inform the spatial neighborhood and correlation structure of VF data, in which there is precedent (Betz-Stablein et al., 2013; Zhu et al., 2014; Warren et al., 2016; Berchuck et al., 2018).




Fig. 1. Example visual field (VF) time series data for a patient with glaucoma. At each VF location, the time series of the observed threshold sensitivities is plotted. The trends and clustering of the time series motivate the use of spatially varying change points.

Our newly developed method is applied to data from the Vein Pulsation Study Trial in Glaucoma and the Lions Eye Institute trial registry, Perth, Western Australia. The dataset contains 1448 VF tests from 194 distinct eyes (98 patients in total). All of the subjects have some form of primary open angle glaucoma. The mean follow-up time for participants is 934 days (2.5 years) with an average of 7.4 tests per subject. Each VF series is diagnosed as progressing based on the clinical judgment of two independent clinicians. In the case that the two clinicians disagree, a third clinician is consulted (occurred for only 13 VF series). In our study, we have 141 (74%) stable and 50 (26%) progressing patient eyes. Three patient eyes had a missing clinical progression determination. These data have been previously analyzed (Betz-Stablein et al., 2013; Warren et al., 2016; Berchuck et al., 2018), in which additional details of the data can be found.

3. Change point modeling for spatially referenced data

VF data represent spatially referenced time series that are monitored for abnormalities that may indicate disease progression. Identifying abruptly changing variability in time series data was first characterized as a CP problem by Page (1954) and was used as a quality control method to automatically detect faults in industrial processes. Although the CP literature is extensive (Quandt, 1958; Carlin et al., 1992; Julious, 2001), the problem of detecting the timing of an unknown change in spatially referenced data is relatively sparse. The core of the literature involves extending the original CP model to account for spatial heterogeneity where the original CP model is given as

$$Y_t = \begin{cases} \beta_0 + \beta_1 x_t + \epsilon_t & t = 1, \dots, \theta, \\ \beta_0 + \beta_1 x_\theta + \beta_2(x_t - x_\theta) + \epsilon_t & t = \theta + 1, \dots, v \end{cases} \quad (1)$$

where Y_t , $t = 1, \dots, v$, are observations of the time series, x_t is a time varying regressor ($x_1 < \dots < x_v$), and $\epsilon_t \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma^2)$. The regression parameters β_0 , β_1 , and β_2 define a mean process that dictates a piece-wise linear regression separated at the CP, $\theta \in \{1, 2, \dots, v - 1\}$, an index of the time series.

With spatiotemporal data, the time series is now observed at a unique spatial location \mathbf{s} , $\{Y_t(\mathbf{s}), t = 1, \dots, v\}$, across the interval $[x_1(\mathbf{s}), x_v(\mathbf{s})]$. Maximum likelihood theory has been developed for the standard CP method of Eq. (1) (Hinkley, 1970), and its many extensions (e.g., multiple CPs), but more recently, and for spatial methods in particular, the Bayesian paradigm is preferred (Menzefricke, 1981; Carlin et al., 1992). There have been various Bayesian CP formulations that account for spatial or spatiotemporal dependencies through adjustments to the intercept (Beckage et al., 2007; Minin et al., 2007; Yu et al., 2008; Smith et al., 2015; Cai et al., 2016) and extensions that allow for independent spatial processes before and after the CP (Majumdar et al., 2005). These methods introduce spatial dependency using conditional autoregressive (CAR) priors, a Gaussian Markov random field that induces neighborhood dependencies across the region (Besag, 1974; Geman and Geman, 1984).

In the context of Eq. (1), this implies that spatial dependencies are most often modeled through adjustments to the intercepts, now spatially referenced as $\beta_0 + \beta_0(\mathbf{s})$. In particular, this implies the following conditional prior form, $\beta_0(\mathbf{s}_i) | \beta_0(\mathbf{s}_{j \neq i}), \tau^2 \stackrel{\text{ind}}{\sim} N(\sum_{j=1}^m w_{ij} \beta_0(\mathbf{s}_j) / w_{i+}, \tau^2 / w_{i+})$, where $w_{ij} = 1(i \sim j)$ is an indicator equal to one if locations i and j are neighbors and $w_{i+} = \sum_{j=1}^m w_{ij}$ is the number of neighbors for location i . This conditional prior distribution has an intuitive interpretation as the mean is an average of the neighboring values and the variance decreases as the number of neighbors increases (Banerjee et al., 2003). Although the CAR prior is an effective method for modeling spatial variability for areal data, these previous methods ignore the spatial variability in the CPs themselves.

In a recent study, Wagner and Midway (2014) introduced a model with spatially varying intercepts, slopes, and CPs. In their method, there is a unique CP at each spatial location and additionally, the process is re-parameterized so that the CP parameter is continuous. The general form of the mean structure (after applying the appropriate link function) is given as

$$\mu_t(\mathbf{s}) = \begin{cases} \beta_0(\mathbf{s}) + \beta_1(\mathbf{s})x_t(\mathbf{s}) & x_t(\mathbf{s}) \leq \theta(\mathbf{s}), \\ \beta_0(\mathbf{s}) + \beta_1(\mathbf{s})\theta(\mathbf{s}) + \beta_2(\mathbf{s})\{x_t(\mathbf{s}) - \theta(\mathbf{s})\} & x_t(\mathbf{s}) > \theta(\mathbf{s}). \end{cases} \quad (2)$$

In this specification, the CP $\theta(\mathbf{s})$ is no longer an index of the time series and instead ranges in $[x_1(\mathbf{s}), x_v(\mathbf{s})]$. In Wagner and Midway (2014), after rescaling the CPs to the real line, the regression parameters and CPs are modeled jointly using a multivariate Gaussian distribution centered at a vector of fixed effects with an unstructured cross-covariance matrix. The authors assumed independence across space in their model specification. Subsequently, this method has been extended to account for spatial dependency using a multivariate CAR (MCAR) process (Warren et al., 2017). The MCAR allowed for an improved modeling framework by accounting for both spatial variability and cross-covariance in the location specific intercepts, slopes, and CPs.

4. Methods

Following the approach of Warren et al. (2017), we propose extending the general CP framework by introducing a set of spatially varying intercepts, slopes, and CPs that are modeled with an MCAR process. This increased flexibility allows us to accurately describe the individual time series behavior at each unique spatial location while accounting for spatial correlation and cross-covariance between the parameters that may result in increased power to detect important changes.

However, we extend the existing framework in a number of ways. First, we also allow for a CP (and spatially varying regression coefficients) in the model for the nuisance (or variance) parameter to reflect the situation where the observed data become more or less variable after a CP occurs. Next, we define the observed CP at a specific spatial location as a function of a potentially censored latent CP process (also spatially varying). This flexible specification permits spatial interpolation of currently censored or unobserved CPs across the spatial domain. This becomes particularly useful when working with data where the CPs describe the timing of treatment failure or an adverse event. Finally, a novel boundary detection MCAR process is developed that allows for a dissimilarity metric in the spatial covariance definition. Inference for this model is based on Markov chain Monte Carlo (MCMC) simulation, and a description of the algorithm is given in Appendix A. The newly introduced model is computationally intensive, so the MCMC algorithm is implemented using Rcpp (Eddelbuettel et al., 2011) and is available from the R package spCP (R Core Team, 2016).

4.1. Spatially varying change points model

We begin by describing our new methodology generally before applying it to VF data in Section 5. The general framework is flexible enough to accommodate many different spatially referenced time series data settings. Let $Y_t(\mathbf{s})$ be an observation from spatial location \mathbf{s} at time $x_t(\mathbf{s})$, for $\{\mathbf{s}_i : i = 1, \dots, m\}$, $t = 1, \dots, v$, where the times are strictly increasing. As an extension of Eq. (2), we introduce a model that allows for location specific CPs, intercepts, and slopes in the mean and nuisance parameter processes of the data such that

$$\begin{aligned} Y_t(\mathbf{s}) | \mu_t(\mathbf{s}), \zeta_t(\mathbf{s}) &\stackrel{\text{ind}}{\sim} f(\mu_t(\mathbf{s}), \zeta_t(\mathbf{s})), \quad g_1\{\mu_t(\mathbf{s})\} = \gamma_t(\mathbf{s}), \quad g_2\{\zeta_t(\mathbf{s})\} = \omega_t(\mathbf{s}) \\ \gamma_t(\mathbf{s}) &= \begin{cases} \beta_0(\mathbf{s}) + \beta_1(\mathbf{s})x_t(\mathbf{s}) & x_t(\mathbf{s}) \leq \theta_1(\mathbf{s}) \\ \beta_0(\mathbf{s}) + \beta_1(\mathbf{s})\theta_1(\mathbf{s}) + \beta_2(\mathbf{s})\{x_t(\mathbf{s}) - \theta_1(\mathbf{s})\} & x_t(\mathbf{s}) > \theta_1(\mathbf{s}) \end{cases} \\ \omega_t(\mathbf{s}) &= \begin{cases} \lambda_0(\mathbf{s}) + \lambda_1(\mathbf{s})x_t(\mathbf{s}) & x_t(\mathbf{s}) \leq \theta_2(\mathbf{s}) \\ \lambda_0(\mathbf{s}) + \lambda_1(\mathbf{s})\theta_2(\mathbf{s}) + \lambda_2(\mathbf{s})\{x_t(\mathbf{s}) - \theta_2(\mathbf{s})\} & x_t(\mathbf{s}) > \theta_2(\mathbf{s}). \end{cases} \end{aligned} \quad (3)$$

The form of Eq. (3) is an extension of Eq. (2) with a general likelihood, f , and a CP in the nuisance process. The mean process is modeled through a link function, $g_1(\cdot)$, and thus can be non-linear. The nuisance parameters are partitioned into constant terms and a location and time specific component, such that $\zeta_t(\mathbf{s}) = \{\gamma_t(\mathbf{s}), \omega_t(\mathbf{s})\}$. The CP for the nuisance parameter is introduced similarly as for the mean process, through the link function $g_2(\cdot)$. At each spatial location, the CP for the mean process, $\theta_1(\mathbf{s})$, and nuisance parameter process, $\theta_2(\mathbf{s})$, are unique.

Allowing for a continuous valued CP is important because we generally do not expect the change to occur discretely during one of the observed time points, particularly when there is a large amount of time between observations. The specification can be improved however, through a natural assumption about the censoring of CPs at the edges of the observed time period. In particular, when $\theta_j(\mathbf{s}) = x_1(\mathbf{s})$, $j = 1, 2$, the CP has occurred before the time series at location \mathbf{s} was monitored, while $\theta_j(\mathbf{s}) = x_v(\mathbf{s})$ indicates that the CP has not yet occurred by the end of the follow-up period. As such, both the lower and upper bounds of the CP are censored. To account for this, we model the observed CP as a function of the underlying and potentially unobserved true latent CP process, $\eta_j(\mathbf{s})$, such that $\theta_j(\mathbf{s}) = \max\{\min\{\eta_j(\mathbf{s}), x_v(\mathbf{s})\}, x_1(\mathbf{s})\}$ where $\eta_j(\mathbf{s})$ has support on the real line and is allowed to vary spatially. This framework gives us the ability to spatially interpolate latent CP parameters that have yet to occur which could be important in some applications for determining when and where future damage can be expected.

4.2. Multivariate dissimilarity metric spatial process

In order to account for the potentially complex spatial correlation and cross-covariance between the introduced spatially varying intercepts, slopes, and CPs, we develop a novel boundary detection MCAR model that allows for a dissimilarity metric within the framework introduced by Mardia (1988). Incorporating the dissimilarity metric into the framework is important for correctly modeling areal spatial data where the neighborhood structure of the region is uncertain and controlled by a set of underlying factors (e.g., VF data). Similar to the CAR process, the MCAR introduces spatial dependencies through a Markov random field technique, and is in fact a multivariate generalization of the CAR with a separable covariance. Since Mardia (1988), the MCAR has seen numerous developments throughout the years that have generalized the separable assumption (Gelfand and Vounatsou, 2003; Jin et al., 2005, 2007). Although there may be motivation for a non-separable process, in this paper we introduce our boundary detection technique within the standard MCAR. This could easily be generalized to a non-separable process using the linear model of coregionalization.

Similar to the univariate case, the joint distribution of the parameters is derived from their full conditional distributions. Under the Markov random field assumption, these conditional distributions can be specified as follows,

$$f(\phi_i | \phi_{j \neq i}, \mu, \Gamma_i) = \text{MVN} \left(\mu_i + \sum_{i \sim j} \mathbf{B}_{ij}(\phi_j - \mu_j), \Gamma_i \right), \quad i, j = 1, \dots, m$$

where $\phi_i = (\phi_{i1}, \phi_{i2}, \dots, \phi_{ip})^T$ and $\mu_i = (\mu_{i1}, \mu_{i2}, \dots, \mu_{ip})^T$ are p -dimensional vectors, $\mu = (\mu_1^T, \mu_2^T, \dots, \mu_m^T)^T$, and Γ_i and \mathbf{B}_{ij} are $p \times p$ matrices. Using Brook's lemma (Brook, 1964), Mardia proved that the full conditional distributions uniquely determine the joint distribution given by $f(\phi|\mu, \{\Gamma_i\}) \propto \exp\left\{-\frac{1}{2}(\phi - \mu)^T \Gamma^{-1}(\mathbf{I}_{mp} - \tilde{\mathbf{B}})(\phi - \mu)\right\}$, where $\phi = (\phi_1^T, \phi_2^T, \dots, \phi_m^T)^T$, Γ is a block diagonal matrix with blocks Γ_i and $\tilde{\mathbf{B}}$ is an $mp \times mp$ matrix with (i, j) th block \mathbf{B}_{ij} . Different MCAR forms can be obtained through various Γ and $\tilde{\mathbf{B}}$.

In this case, we induce an MCAR structure that generalizes the univariate version of the Leroux CAR specification, thus allowing the dissimilarity metric to be used in a multivariate setting (Leroux et al., 2000; MacNab and Gustafson, 2007; Lee and Mitchell, 2011). As in the univariate case, symmetry of $\Gamma^{-1}(\mathbf{I}_{mp} - \tilde{\mathbf{B}})$ is required. A convenient special case sets $\mathbf{B}_{ij} = b_{ij}\mathbf{I}_p$. Under this condition, the multivariate version of the Leroux dissimilarity metric likelihood can be obtained with the following specifications, $b_{ij} = \rho w_{ij}(\alpha) \Delta^{-1}$ and $\Gamma_i = \Sigma \Delta^{-1}$, where $\Delta = \rho \sum_{j=1}^m w_{ij}(\alpha) + (1 - \rho)$. Spatial neighborhood structure is introduced through the adjacencies $\{w_{ij}(\alpha)\}$.

The adjacencies are defined similar to the specification in Berchuck et al. (2018). In particular, they are a function of a dissimilarity metric and regression parameter such that $w_{ij}(\alpha) = 1\{i \sim j\} \exp\{-z_{ij}\alpha\}$ where $i \sim j$ is the event that locations i and j share an edge or corner ($w_{ii} = 0$ for all i). The quantity z_{ij} is a measure of (dis-)similarity between two locations and is defined as $z_{ij} = |z_i - z_j|$, where z_i is a covariate at location \mathbf{s}_i . The parameter α is forced to be non-negative, so that the adjacencies are in the open unit interval, and this is in fact a condition for guaranteeing a valid covariance.

This specification induces a multivariate version of the conditional distributions which is analogous to the univariate Leroux likelihood and is given as $f(\phi_i|\phi_{j \neq i}, \delta, \Sigma, \alpha, \rho) =$

$$\text{MVN}\left(\frac{\rho \sum_{j=1}^m w_{ij}(\alpha) \phi_j + (1 - \rho)\delta}{\rho \sum_{j=1}^m w_{ij}(\alpha) + (1 - \rho)}, \frac{\Sigma}{\rho \sum_{j=1}^m w_{ij}(\alpha) + (1 - \rho)}\right)$$

where $\mu_i = \delta$, so that the mean is constant across spatial locations. The distribution for the joint specification can be established by noting that $\Gamma^{-1}(\mathbf{I}_{mp} - \tilde{\mathbf{B}}) = \mathbf{Q}(\alpha, \rho) \otimes \Sigma^{-1}$, with $\mathbf{Q}(\alpha, \rho) = \rho \mathbf{W}^*(\alpha) + (1 - \rho)\mathbf{I}_m$. The matrix $\mathbf{W}^*(\alpha)$ has diagonal elements $w_{ii}^*(\alpha) = \sum_{j=1}^m w_{ij}(\alpha)$ and off-diagonal elements $w_{ij}^*(\alpha) = -w_{ij}(\alpha)$. Finally, this form of the adjacencies yields the joint distribution $f(\phi|\delta, \Sigma, \alpha, \rho) \sim \text{MVN}(\mathbf{1}_m \otimes \delta, \mathbf{Q}(\alpha, \rho)^{-1} \otimes \Sigma)$. We denote a random variable ϕ with this distribution as $\phi \sim \text{MCAR}(\alpha, \rho, \delta, \Sigma)$.

4.3. Hyperprior distributions

To complete the model specification, we select prior distributions for the hyperparameters. The prior for the mean δ , covariance Σ , and boundary detection parameter α are as follows;

$$\delta \sim \mathbf{N}(\mathbf{0}, \kappa^2 \mathbf{I}_p), \quad \Sigma \sim \text{Inverse-Wishart}(\xi, \Psi), \quad \text{and } \alpha \sim \text{Uniform}(a_\alpha, b_\alpha).$$

The mean δ is zero centered with variance, κ^2 , which is chosen to be large in order to yield a weakly informative prior. For the prior on Σ , we specify an inverse-Wishart distribution with degrees of freedom $\xi = p + 1$ and scale matrix, $\Psi = \mathbf{I}_p$, where p is the dimension of the spatial parameter, ϕ_i . This prior induces marginally uniform priors on the correlations of Σ and allows for the diagonals to be weakly informative (Gelman et al., 2014). The dissimilarity metric parameter, α , is forced to be non-negative with a uniform distribution. The bounds are defined as $a_\alpha = 0$ and $b_\alpha = -\log\{0.5\}/\min_{i,j}\{z_{ij}\}$. As in Lee and Mitchell (2011), the upper limit is set so that the smallest dissimilarity metric between two locations can still obtain a spatial adjacency of 0.50.

We fix ρ at 0.99 to force spatial dependency, in which there is precedent (Lee and Mitchell, 2011; Berchuck et al., 2018). In particular, Berchuck et al. (2018) demonstrated that ρ fixed at 0.99 has little impact on estimation of α . Multivariate spatial structure is introduced through $\phi_i = \{\beta_0(\mathbf{s}_i), \beta_1(\mathbf{s}_i), \beta_2(\mathbf{s}_i), \eta_1(\mathbf{s}_i), \lambda_0(\mathbf{s}_i), \lambda_1(\mathbf{s}_i), \lambda_2(\mathbf{s}_i), \eta_2(\mathbf{s}_i)\}^T$, which are assigned an $\text{MCAR}(\alpha, 0.99, \delta, \Sigma)$ process.

5. Monitoring visual field progression

5.1. Change points model for visual field data

We modify the general form of the model described in Section 4 for the purpose of monitoring a patient's VF data across space and time with the goal of improving disease management. Using the notation developed in Section 4, we define $Y_t^*(\mathbf{s})$ as the observed VF sensitivity at VF location \mathbf{s} and x_t days after the baseline visit ($x_1 = 0$), for $\{\mathbf{s}_i : i = 1, \dots, 52\}$, $t = 1, \dots, v$, where the number of visits, v , is patient specific. Note that since the timing of visits is the same for all VF locations for a single patient, $x_t(\mathbf{s}) \equiv x_t$. Furthermore, the number of VF locations, 52, represents the total number of locations on the VF after removing the two corresponding to the blind spot.

The data generating mechanism for VF sensitivities is such that observations are censored at zero due to limitations of the machine that collects the data. Therefore, we use a Tobit model (Tobin, 1958), in which there is precedent in the glaucoma progression literature (Betz-Stablein et al., 2013; Bryan et al., 2015; Berchuck et al., 2018). To induce the Tobit model, we define a latent process through the standard Tobit link, $Y_t^*(\mathbf{s}) = \max\{0, Y_t(\mathbf{s})\}$, where $Y_t(\mathbf{s}) = \mu_t(\mathbf{s}) + \epsilon_t(\mathbf{s})$, $\epsilon_t(\mathbf{s}) \stackrel{\text{ind}}{\sim} N(0, \sigma_t^2(\mathbf{s}))$. This is a special case of Eq. (3) where $\zeta_t(\mathbf{s}) = \sigma_t^2(\mathbf{s})$, $g_1(\cdot)$ is the identity link, and $g_2(\cdot) = \log\{\sqrt{\cdot}\}$.

To appropriately model the course of glaucoma progression, we modify the mean and variance structures. Patients with glaucoma are typically stable until the point of progression when their VF time series begin to deteriorate (see Fig. 1 for example). As such, for VF data we define the mean process so that at times before the CP there is a constant intercept at each VF location. To accomplish this, we set $\beta_1(\mathbf{s}) = 0$ in Eq. (3). The same assumption is used for the variance process, so that the variance can increase with the onset of deterioration (i.e., $\lambda_1(\mathbf{s}) = 0$). Furthermore, motivated by the disease course, we force the CPs in the mean and variance to occur simultaneously, $\theta(\mathbf{s}) \equiv \theta_1(\mathbf{s}) = \theta_2(\mathbf{s})$. To finalize the model, we specify hyperpriors based on details in Section 4.3, where $\kappa^2 = 1000$ and $p = 5$. We define the dissimilarity metric, z_i , as the Garway-Heath angle at VF location \mathbf{s}_i . A sensitivity analysis to determine the impact of the values of κ^2 and b_α on the final results is given in Appendix D.1.

We draw samples from the joint posterior by utilizing MCMC techniques (Metropolis et al., 1953; Geman and Geman, 1984; Gelfand and Smith, 1990). Many of the parameters have conjugate full conditionals. In particular, all of the parameters have closed form solutions, with the exception of α , $\lambda_0(\mathbf{s})$, $\lambda_2(\mathbf{s})$, and $\eta(\mathbf{s})$. A Gibbs sampler is used with Metropolis steps for parameters lacking conjugacy. For each eye, the MCMC sampler is run for 250,000 iterations after a 10,000 burn-in and thinned to a final sample size of 10,000 in order to reduce autocorrelation in the posterior samples and to reduce the data storage size of the samples collected after fitting the model to each patient eye. There were no obvious signs of non-convergence observed for any of the 194 VF series based on traceplots and the Geweke statistic (Geweke, 1992). The MCMC sampler is implemented in the R package spCP. More details about the MCMC sampler can be found in Appendix A and a sensitivity analysis to investigate the impact of thinning the full set of posterior samples can be found in Appendix D.2.

5.2. Model comparison

In order to establish the capability of the introduced model to predict VF sensitivities, predict the location and timing of future vision loss, and diagnose glaucoma progression, we compare it to a number of competing models with varying levels of complexity:

1. *Point-wise linear regression (PLR)*: Independent Tobit linear regression at each location.
2. *Non-spatial CP (D)*: Tobit linear regression with **discrete** CPs, based on Eq. (1).
3. *Non-spatial CP (C)*: Tobit linear regression with **continuous** CPs, where the CPs at each VF location are independent and uniformly distributed a priori.
4. *Non-spatial CP (L)*: Tobit linear regression with continuous CPs and an underlying normally distributed **latent** process, as specified in Section 4.1.

Table 1

Model fit and prediction diagnostics for model comparison. Model fit is determined using the deviance information criterion (DIC) and effective number of parameters (p_D). Prediction capability is determined using the mean squared prediction error (MSPE). The ability of each model to diagnose progression is assessed using a metric dependent on the latent change points (CPs). For each model with a latent CP process, the metric $\max_i \{P[\theta(\mathbf{s}_i) < t_v | \mathbf{Y}]\}$ is calculated and regressed against the clinically determined progression status. Diagnostic capability is measured using Akaike information criterion (AIC), area under the receiver operating characteristic curve (AUC), and the p -value for each measure.

Model	Model Fit		Prediction	Diagnostic		
	DIC	p_D	MSPE	AIC	AUC	p -value
PLR	31.11	120.30	0.20	–	–	–
Non-spatial CP (D)	–14.27	119.30	0.45	–	–	–
Non-spatial CP (C)	–18.76	112.13	0.32	–	–	–
Non-spatial CP (L)	–29.05	67.69	0.23	211.36	0.64	0.001
Spatial CP	–53.74	68.21	0.13	208.67	0.69	<0.001

5. Spatial CP: Spatial CPs model detailed in Section 5.1.

Each of the competing CP models specify a constant process before the CP, as specified in Section 5.1. In addition, each of the models utilize the Tobit likelihood in order to gain an understanding of key differences in their performance due to other factors. The motivation of fitting the simplified models, such as PLR and basic CP models, is to have a benchmark for comparison and to emphasize the importance of the latent CP specification and spatial structure. Note that each of the models are subsets of Spatial CP.

The models are compared using model fit diagnostics and hold-out prediction criterion. To assess model fit, the deviance information criterion (DIC) is used (Spiegelhalter et al., 2002). For prediction, the final visit for each of the VF series is held out, and the remaining visits are used to predict the hold-out visit. Prediction accuracy is measured using the mean squared prediction error (MSPE). Smaller values of DIC and MSPE are preferred.

The model fit and prediction diagnostics can be found in Table 1 averaged over all patients. The DIC for Spatial CP is superior to the rest of the models, with a minimal value of –53.74. Generally, there is a monotone decreasing trend in DIC as the models become more complex. The number of effective parameters also trends smaller with added model complexity, indicating that the added model structure provides a more concise framework for modeling VF data. The prediction measure also favors Spatial CP, with optimal MSPE of 0.13. MSPE for the three non-spatial CP models is inferior to PLR, indicating that the inclusion of a CP does not guarantee improved prediction. Only when accounting for the spatial nature of VF data does a CP model offer improved prediction over PLR, again illuminating the importance of modeling spatial variability. Overall, the results in Table 1 indicate that Spatial CP has superior model fit and prediction ability for VF data.

5.3. Diagnosing progression using change points

To illustrate the superiority of the introduced methodology, in addition to model fit and prediction, we demonstrate the diagnostic capability of the CPs. We hypothesize that the location of the CPs can serve as a proxy for disease progression and propose a diagnostic metric that is a function of the posterior probability that a CP has occurred up to the last VF visit, $p_i = P[\theta(\mathbf{s}_i) < t_v | \mathbf{Y}]$. We define the final metric for a patient eye as the maximum across all VF probabilities, $\max\{p_i : i = 1, \dots, 52\}$. This is intuitive because a longitudinal series of VFs with earlier occurring CPs is more likely to be progressing.

In order to assess the diagnostic ability of the CPs, we construct logistic regression models, regressing the defined metric on the clinical assessment of progression and present Akaike information criterion (AIC), area under the receiver operating characteristic curve (AUC), and a p -value for the slope parameter. These measures are presented in Table 1 for models that have a latent

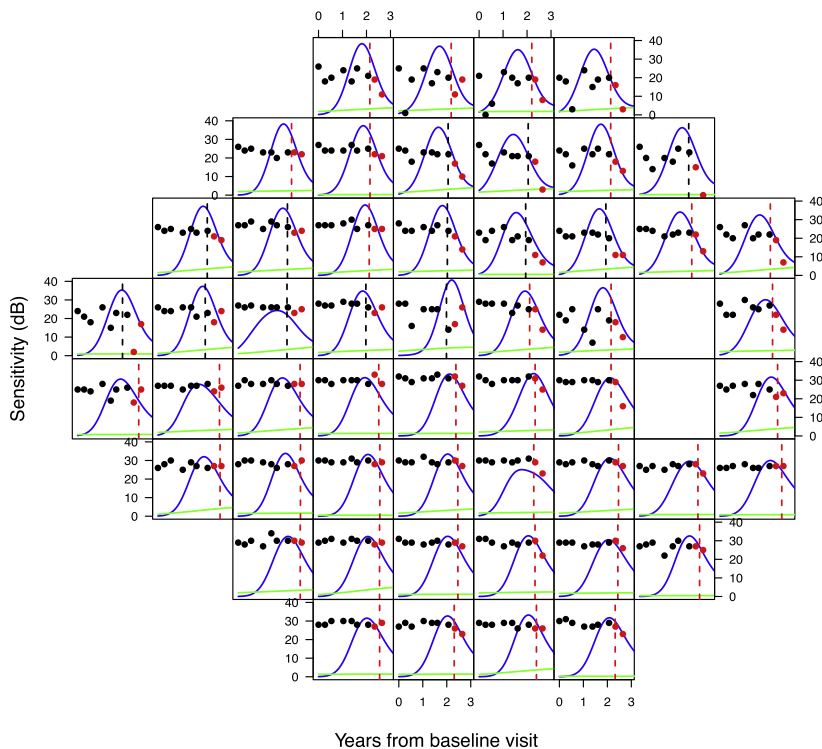


Fig. 2. Demonstrating the ability of Spatial CP to identify change points (CPs) that have yet to occur. At each cell, the observed DLS are presented for an example patient with nine visits. Based on only the first seven visits, the posterior density estimates of the CPs for Spatial CP and Non-spatial CP (L) are presented as density estimates in blue and green lines, respectively. The estimated CP based on Spatial CP using all of the nine VF visits are presented as dashed vertical lines. The DLS and CPs are shown in red if they occur after the seventh visit and thus are not observable at the time the posterior densities were estimated. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

CP process (the only models capable of estimating the CP probability). From Table 1, it is clear that the CPs arising from Spatial CP are highly associated with progression (p -value: < 0.001). The non-spatial model also yields a significant slope estimate, however the p -value is of a different magnitude (p -value: 0.001). Furthermore, the values of AIC and AUC are optimized under Spatial CP with minimal AIC, 208.67, and maximal AUC, 0.69. These results help to validate the claim that the CPs are representative of disease progression and that their estimates are clinically useful.

To further illustrate the diagnostic ability granted by the latent CP framework, we present an example patient as a proof of concept. This patient's information is appropriate for investigation because they have late occurring CPs that are truncated when removing visits near the end of follow-up. The patient has nine VF visits, however we only use the initial seven to fit Spatial CP and spatially interpolate the latent CPs. We then compare these predictions with the actual CPs estimated using the full set of data (i.e., all nine visits). The patient's VF data are presented in Fig. 2, where at each location the first seven VF sensitivities are shown using black dots and the last two are shown in red. Furthermore, the mean CPs estimated by Spatial CP for the full data are presented as dashed vertical lines. Any CP that was estimated to have occurred after the seventh visit is colored red and is technically truncated for purposes of clinical use at visit seven. Finally, the posterior latent CP density estimates (based on the reduced data) for Spatial CP and Non-spatial CP (L) are presented as density estimates in blue and green, respectively.

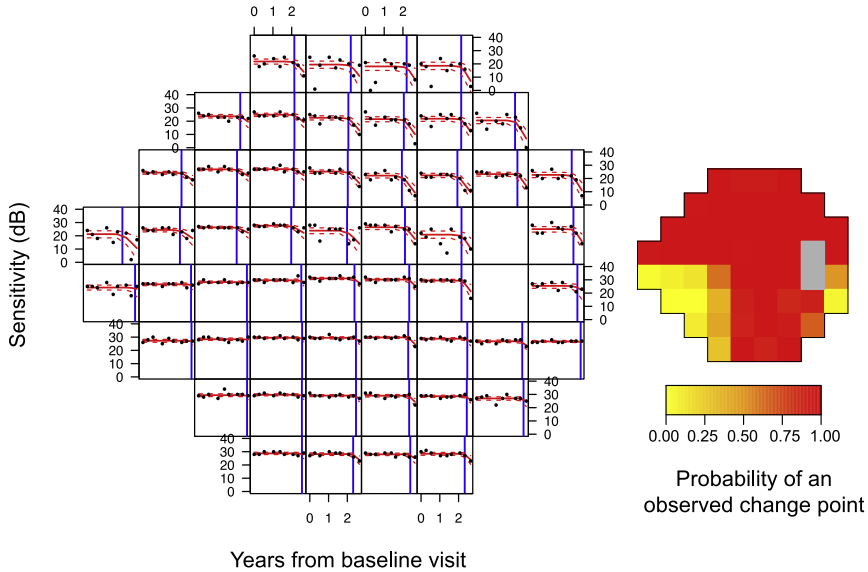


Fig. 3. Output generated from Spatial CP. The plot on the left presents the estimated mean (with 95% credible intervals) and CPs for the same example patient in Fig. 1. The mean process is presented in red and the change points are vertical blue lines. The panel on the right presents the posterior probability of an observed CP at each location on the VF.

From Fig. 2 it is clear that Spatial CP is capable of predicting future CPs that have not been observed. This is remarkable, since at the seventh visit there was no clear sign of a CP at many of these locations and furthermore, Non-spatial CP (L) performed poorly, indicating the importance of spatial structure. To quantify the conclusions drawn from Fig. 2, for each model we present diagnostics for prediction of the estimated CPs. Each model is compared to the “true” CPs estimated from the same model with full data. Spatial CP and Non-spatial (L) have MSPE of 4.70 and 8525.82, respectively, indicating Spatial CP provides superior prediction of the unobserved CPs. These powerful results indicate that the introduced methodology has the capability of predicting future unobserved CPs, and therefore progression. This is possible due to both the latent CP specification and multivariate spatial boundary detection structure. As shown in Table 1 and Fig. 2, without either of these specifications the model suffers in terms of model fit, prediction, and diagnostic capability.

5.4. Clinical implementation

Having established the potential for using Spatial CP in the clinical setting for prediction of VF sensitivities and CPs, we now present example output for use clinically. In Fig. 3, on the left, the posterior fitted trend is presented on the VF with the estimated posterior CPs. The plot presents the estimated regression mean (with 95% credible intervals) and CPs for the same example patient in Fig. 1. The mean process is presented in red and the CPs are vertical blue lines. This figure demonstrates the flexibility of Spatial CP to produce non-linear trends, not simply two lines connected at a point. This type of trajectory is a result of the CPs being unknown quantities. In addition to the predicted fit, it is useful to have a heatmap of the posterior probability that each of the CPs have occurred in the observed data (right panel of Fig. 3). This presentation corresponds to the estimated CPs in the left frame, but further provides a useful map for clinicians to understand the location and severity of progression.

Both of these plots in conjunction with the diagnostic potential from Spatial CP provide a useful framework for clinicians to assess progression. This presentation is clinically useful, and can be further improved when presented in video format. In Video B.1 of Appendix B, the heatmap of

Fig. 3 is presented over the full course of follow-up and into the future to provide clinicians an understanding of the disease pattern and growth over time and into the future. The combination of the latent CP process and novel spatial structure has produced a technique that goes beyond the simple CP models and PLR in terms of model fit, prediction, and progression diagnoses.

6. Simulation study

The purpose of the simulation study is to investigate the performance of Spatial CP and the comparison models detailed in Section 5.2 under different data generating scenarios. The scenarios are all subsets of the joint spatial structure, $MVN(\mathbf{1}_m \otimes \delta, \mathbf{Q}(\alpha, 0.99)^{-1} \otimes \Sigma)$, with parameters chosen to reflect those from an actual analyzed glaucoma patient:

$$\delta = \begin{bmatrix} 25 \\ -15 \\ -0.5 \\ 0.1 \\ 0.5 \end{bmatrix}, \quad \Sigma = \begin{bmatrix} 0.025 & -0.500 & -0.500 & -0.500 & 0.500 \\ -0.500 & 0.025 & 0.500 & 0.500 & -0.500 \\ -0.500 & 0.500 & 0.025 & 0.250 & -0.500 \\ -0.500 & 0.500 & 0.250 & 0.025 & -0.500 \\ 0.500 & -0.500 & -0.500 & -0.500 & 0.025 \end{bmatrix}, \quad (4)$$

and $\alpha = 0.10$. Based on these true parameters, we create five data generating settings: (1) Progressing: CPs have all occurred before the beginning of follow-up; (2) Stable: CPs have yet to occur; (3) Non-spatial CPs: CPs are observed throughout follow-up in the mean and variance process with no spatial dependency; (4) Spatial CPs with no covariance: spatially dependent CPs are observed throughout follow-up in the mean and variance process with no cross-covariance in the spatial effects; (5) Full spatial CP model: spatially dependent CPs are observed throughout follow-up in the mean and variance process with spatial cross-covariance.

Settings 1–4 are obtained by changing components of δ , Σ , and α . For Settings 1 and 2, α is fixed at an arbitrarily large value, 1000, to remove spatial dependency, and Σ is restricted to being diagonal, with the fourth entry equal to zero to remove a CP in the variance. This specification yields the PLR model as long as the CPs are not observed. As such, Setting 1 is obtained by fixing the fifth entry in δ to an arbitrarily large negative value, -10, while a large positive value, 10, is used for Setting 2. To guarantee no CP in the variance, the fourth entry in δ is fixed at zero. Setting 3 introduces a CP in the variance process by using the specification of δ from Eq. (4) and the fully diagonal Σ . The value of α is still fixed at 1000, so that there is no spatial dependence. Setting 4 is obtained by specifying Σ to be diagonal and represents a spatial CP model with no cross-covariance between spatially referenced effects. Setting 5 uses the full parameter specification.

A new dataset is generated for each simulated vector of ϕ to ensure that the results are not affected by a particular realization. For each setting, we simulate data based on a set of known hyperparameters and then fit each model to the same dataset to assess model fit, prediction, and estimation of CPs (both latent and observed). As in Section 5, for model fit we utilize DIC and p_D for exploring explanatory ability of each model, and MSPE for prediction. The framework for assessing prediction is innately linked to how the data were simulated. Each dataset was simulated assuming there were 21 VF visits (the maximum number of visits in our study data), however, the models were only fit to the first 14 visits. This allows us to use the last 7 visits for prediction assessment. The visit times were specified as the sequence of numbers between 0 and 1 by a 0.05 increment. Finally, estimation of both the observed CPs and latent CPs process are assessed using bias, MSE, and empirical coverage (EC) of the Bayesian credible intervals. In this context, the EC is the proportion of the simulated datasets where the calculated 95% credible interval contains the true parameter value.

The results from the study are displayed in Tables 2, 3, and 4 and relate to 250 simulated data sets for each setting. In Table 2, the model fit diagnostics, DIC, and p_D , are presented for the five models across all data generating settings. As expected, the models all have generally similar DIC values for the first two settings, which are just progressing and stable representations of PLR. As the data generating settings increase in model complexity, the DIC value corresponding to Spatial CP separates from the non-spatial CP models and PLR. The effective number of parameters, p_D , is similar for all models in Setting 1, but is generally smaller for Spatial CP and Non-spatial CP (L)

Table 2
Model fit diagnostics for the simulation study. The diagnostics included are deviance information criterion (DIC) and effective number of parameters (p_D) and are presented for each of the models introduced in Section 5.2. The data generating settings for the simulation increase in model complexity and are as follows, (1) progressing, (2) stable, (3) non-spatial CP, (4) spatial CP with no covariance, and (5) full spatial CP model. Each reported estimate is based on 250 simulated datasets.

Metric	Model	Simulation Setting				
		1	2	3	4	5
DIC	PLR	−466.33	−409.93	−360.16	−401.16	−276.76
	Non-spatial CP (D)	−457.89	−434.85	−426.27	−455.50	−293.29
	Non-spatial CP (C)	−412.21	−442.09	−414.03	−448.02	−293.54
	Non-spatial CP (L)	−473.37	−476.06	−446.44	−446.77	−318.36
	Spatial CP	−479.61	−455.21	−457.57	−468.10	−351.22
p_D	PLR	118.23	118.39	118.59	121.08	118.72
	Non-spatial CP (D)	111.19	118.51	117.44	134.85	125.36
	Non-spatial CP (C)	133.70	105.29	120.37	135.56	115.00
	Non-spatial CP (L)	109.34	69.65	87.82	78.27	77.70
	Spatial CP	118.80	92.41	110.22	94.71	90.56

Table 3
Prediction diagnostics for estimation of the visual field (VF) visits at future times 0.75 and 1. Mean squared prediction error is presented for each of the models introduced in Section 5.2. The data generating settings for the simulation increase in model complexity and are as follows, (1) progressing, (2) stable, (3) non-spatial CP, (4) spatial CP with no covariance, and (5) full spatial CP model. Each reported estimate is based on 250 simulated datasets.

Time	Model	Simulation Setting				
		1	2	3	4	5
0.75	PLR	215.14	577.70	450.40	358.81	531.48
	Non-spatial CP (D)	213.56	572.25	447.12	364.53	532.32
	Non-spatial CP (C)	215.35	577.26	451.93	362.37	531.95
	Non-spatial CP (L)	213.62	577.54	451.17	359.04	530.65
	Spatial CP	187.63	577.25	439.73	349.58	524.76
1	PLR	131.81	573.39	451.89	248.02	541.26
	Non-spatial CP (D)	129.24	558.73	443.05	258.77	539.04
	Non-spatial CP (C)	132.91	573.20	455.12	254.08	542.85
	Non-spatial CP (L)	131.47	574.85	452.77	247.87	541.05
	Spatial CP	109.62	575.05	443.39	239.76	527.46

for the remaining settings. It is encouraging that Spatial CP has a smaller p_D in the more complex settings, indicating that the increased model structure maintains parsimony even in the presence of complex data.

Next, we studied the capability of each model to predict future VF sensitivities. Prediction ability was assessed for each model by predicting the simulated data for the final 7 visits excluded from the modeling. In Table 3, we present results for two of these visits. For Setting 1, MSPE shows similar values for most models within each time point with the exception of Spatial CP which is superior to the competing methods. This is surprising, since this data generating setting is a subset of each of the competing models, but possibly indicates that Spatial CP is preferred even in the absence of clear spatial structure. In Setting 2, all of the models have comparable performance, while in Settings 3 and 4, Spatial CP generally shows small improvements. Finally, in the most complex setting, Spatial CP has superior prediction ability at both future visits.

Finally, we investigate the ability of each model to estimate both the observed, $\theta(\mathbf{s})$, and latent, $\eta(\mathbf{s})$, CPs. In Table 4, we present bias, MSE, and EC for models that are capable of estimating either of these CPs. Note that Non-spatial CP (D) was not included in the observed CP section since it is estimating a different, discrete, CP. The results for estimating the observed CP demonstrate the

Table 4

Summary diagnostics for estimation of the observed, $\theta(\mathbf{s})$, and latent, $\eta(\mathbf{s})$, change points (CPs). The diagnostics included are bias, mean squared error (MSE)), and empirical coverage (EC) (95% nominal) and are presented for each of the models introduced in Section 5.2 that are capable of estimating either the observed or latent CP. The data generating settings for the simulation increase in model complexity and are as follows, (1) progressing, (2) stable, (3) non-spatial CP, (4) spatial CP with no covariance, and (5) full spatial CP model. Each reported estimate is based on 250 simulated datasets.

Estimand	Metric	Model	Simulation setting				
			1	2	3	4	5
$\theta(\mathbf{s})$	Bias	Non-spatial CP (C)	−0.11	0.19	0.07	0.07	0.09
		Non-spatial CP (L)	−0.00	0.06	0.01	0.04	0.01
		Spatial CP	−0.00	0.00	0.00	0.01	0.00
	MSE	Non-spatial CP (C)	0.02	0.07	0.05	0.03	0.05
		Non-spatial CP (L)	0.00	0.04	0.03	0.04	0.04
		Spatial CP	0.00	0.00	0.00	0.01	0.00
	EC	Non-spatial CP (C)	0.00	0.00	0.16	0.58	0.16
		Non-spatial CP (L)	1.00	1.00	0.98	0.90	0.98
		Spatial CP	1.00	1.00	0.98	0.97	0.99
$\eta(\mathbf{s})$	Bias	Non-spatial CP (L)	16.00	−17.69	−3.82	−11.18	−18.99
		Spatial CP	−1.59	1.01	−0.00	−0.01	0.01
	MSE	Non-spatial CP (L)	4088.01	5399.53	4615.01	4453.89	5531.15
		Spatial CP	33.33	26.85	1.90	0.04	0.56
	EC	Non-spatial CP (L)	0.99	0.99	0.90	0.84	0.96
		Spatial CP	0.85	0.84	0.90	0.91	0.95

importance of the latent CP process for properly modeling the true CP. The most important finding in these results is that Non-spatial CP (C) is incapable of estimating the observed CP, with inferior bias, MSE, and EC. In particular, the EC is estimated to be zero for Settings 1–2. This is clearly due to the lack of a latent CP specification, as the two models with this feature have improved EC. Finally, it is important to note that Spatial CP has superior bias and MSE in estimating the observed CPs, indicating that in addition to the latent CPs, spatial structure is also crucial.

In the simulation study for estimation of the latent CPs, only Spatial CP and Non-spatial CP (L) are compared (i.e., the only models with latent CPs). It is immediately clear that the inclusion of spatial structure is critical for estimating the latent CPs. In all simulation settings, the bias and MSE are improved for Spatial CP. The bias and MSE are generally larger in the Settings 1–2 because there are no observed CPs in the generated data. From Table 4 it is clear that Spatial CP is needed for accurately estimating both the latent and observed CPs.

Finally, an additional simulation study is presented in C to demonstrate the performance of Spatial CP in estimating the observed CP across various points of follow-up. In Fig. C.1 and Table C.1, the results illustrate the importance of Spatial CP, especially in models with CPs near the bounds of follow-up.

7. Discussion

In this paper, we presented a unified framework that combines multiple vital aspects of glaucoma management; (i) prediction of future VF sensitivities, (ii) predicting the timing and spatial location of future vision loss, and (iii) making clinical decisions regarding progression. Furthermore, this framework incorporates important anatomical information. Previous statistical work has focused on one or some of these aspects typically, but to our knowledge, no methods have considered all simultaneously; in particular the latent CP process. Although motivated by VF data, the methodology was introduced in a general manner that permits the model to be applied in broad areal data settings where a CP is appropriate. The methodology is built upon theory in the CP literature, utilizing a Bayesian hierarchical modeling framework for making inference.

We extended previous attempts at modeling CPs across spatial units, that either ignore spatial dependencies (Wagner and Midway, 2014) or do not include a CP in the variance process (Warren

et al., 2017). The inclusion of the variance CP process can add power to detect CPs and was motivated by the known inverse relationship between decreases in VF sensitivity and variability (Russell et al., 2012). Finally, we extended the univariate spatial boundary detection framework for allowing a dissimilarity metric to dictate the local neighborhood structure on the VF. The novel multivariate extension allows the anatomy of the optic disc to dictate the neighborhood adjacencies for the mean, variance, and CP processes. This specification functions jointly and accounts for known dependencies between the separate processes through a cross-covariance. The combination of these model features leads to a superior modeling technique for longitudinal series of VFs.

The results from applying our method to VF data (Section 5) demonstrated the benefit of our methodology in both prediction of future VF sensitivities and model fit. Spatial CP was compared to non-spatial CP models and PLR and was shown to have superior DIC and MSPE (Table 1); this was later confirmed in simulation (Tables 2 and 3). These results indicate the promise of Spatial CP clinically, as predicting future VF sensitivities is incredibly important for managing progression (Crabb et al., 1997). These improvements can be attributed to the inclusion of a spatial dependency structure and the flexible trajectories produced by the CPs across time. The need for spatiotemporal treatment of VF data is demonstrated here by a reduction in prediction errors and observed clinically meaningful results.

In addition to an increase in prediction precision, the CP framework allows for prediction of the timing and spatial location of future vision loss. This information is accessible due to both the modeling of spatial dependencies and the latent specification of the CP. The benefits of the latent CP are that it provides a biological interpretation of true disease progression and informs about VF sensitivity outside of the observed follow-up period. The latent CP is modeled using the novel multivariate spatial process, which allows the anatomy of the optic disc to dictate the local neighborhood structure. As such, the patterns of future vision loss based on the latent CP are anatomically driven. This can be seen in the estimated probabilities of a CP for an example patient (right panel of Fig. 3 and Video B.1 of Appendix B), which display a clear differentiation in the inferior and superior sectors as defined by Garway-Heath et al. (2000). The practical utility of this latent CP specification is rooted in the biological motivation, but also in that it provides an alternate method for predicting severity across the VF.

The latent CP clearly represents an improvement over standard CP specifications that limit the range of the CP to the observed follow-up period. Furthermore, it allows us to predict unobserved CPs. This was exemplified with an example patient with many observed CPs near the end of follow-up. For this patient, we showed that when removing the last two visits (and thus truncating these CPs), Spatial CP was still capable of detecting the presence of an imminent CP. This served as a proof of practice for predicting unobserved CPs and was facilitated by the spatial correlation between VF locations, as Non-spatial CP (L) was incapable of detecting these truncated CPs (Fig. 2). This was further confirmed in simulation, as Spatial CP estimated the latent CP with lower bias and MSE (Table 4).

We defined a diagnostic metric as the maximum probability across the VF of a CP occurring in the observed follow-up period and showed that it is a significant predictor of clinically determined glaucoma progression (Table 1). This contribution is important clinically, since a detectable change in a longitudinal series of VFs can be highly suggestive of progression, even if standard VF metrics may appear well within normal limits (Artes and Chauhan, 2005). This finding supports the hypothesis that a CP can intrinsically parameterize the timing of progression. Furthermore, it confirms the importance of Spatial CP for characterizing structural change on the VF that can be used for improved clinical decision making.

This paper introduced a framework for assessing glaucoma progression from numerous angles and demonstrates the clinical utility of CPs acting as a proxy for disease progression. To validate these findings it will be important to compare the diagnostic efficiency of the CPs to other established metrics in the glaucoma progression literature, in addition to the ones presented here. Furthermore, the estimation of progression can be improved when combining functional (e.g., VF) and structural (e.g., optic disc) testing techniques (Nicolela et al., 2003). As such, a synthesis of these two data sources could provide improvement in estimating the CPs. This would provide stronger clinical evidence of progression as damage to the optic disc could be corroborated with the VF and

vice-versa (Artes and Chauhan, 2005). Finally, this work opens up numerous avenues for future statistical research including an extension to consider multiple CPs within a single spatial location and the generalization of the separable multivariate spatial process to a more flexible structure with varying dissimilarity metric parameters, α_p (e.g., linear model of coregionalization). When extending the model to different data settings, we expect longer time series to have only minimal impact on the computational efficiency of the algorithm while including more spatial locations would have a larger impact due to the introduction of additional model parameters. We do not expect the introduction of multiple CPs to impact computational performance greatly, though identifiability issues could result if all CPs are not well separated in time or if the variance in the data is too large.

Acknowledgments

This work was partially supported by the National Center for Advancing Translational Science, USA (JLW; UL1 TR001863, KL2 TR001862) and the Research to Prevent Blindness, USA (Department of Ophthalmology, University of North Carolina-Chapel Hill). The authors thank Brigid D. Betz-Stablein (School of Medical Sciences, University of New South Wales), William H. Morgan (Lions Eye Institute, University of Western Australia), Philip H. House (Lions Eye Institute, University of Western Australia), and Martin L. Hazelton (Institute of Fundamental Sciences, Massey University) for providing the dataset from their original study for use in this analysis.

Appendix A. Implementation of the model

For a full introduction of the model, see Sections 4 and 5. Recall that to induce the Tobit model, we define a latent process through the standard Tobit link, $Y_t^*(\mathbf{s}) = \max\{0, Y_t(\mathbf{s})\}$. As mentioned in Section 5.1, to appropriately model the course of glaucoma progression, we will modify the mean and variance structures to reflect the glaucoma follow-up. The mean and variance processes are defined as follows,

$$\mu_t(\mathbf{s}) = \begin{cases} \beta_0(\mathbf{s}) & \text{if } 0 \leq x_t \leq \theta(\mathbf{s}), \\ \beta_0(\mathbf{s}) + \beta_1(\mathbf{s})\{x_t - \theta(\mathbf{s})\} & \text{if } \theta(\mathbf{s}) \leq x_t \leq x_v, \end{cases}$$

and

$$\log\{\sigma_t(\mathbf{s})\} = \begin{cases} \lambda_0(\mathbf{s}) & \text{if } 0 \leq x_t \leq \theta(\mathbf{s}), \\ \lambda_0(\mathbf{s}) + \lambda_1(\mathbf{s})\{x_t - \theta(\mathbf{s})\} & \text{if } \theta(\mathbf{s}) \leq x_t \leq x_v. \end{cases}$$

To obtain this form from the specification detailed in Section 4.1, set $\beta_1(\mathbf{s}) = 0$ and then recast $\beta_2(\mathbf{s})$ as $\beta_1(\mathbf{s})$. The same is done with the variance parameters. Finally, we force the CPs from the mean and variance to be equal.

As mentioned in Section 5, we draw samples from the joint posterior using MCMC techniques (Metropolis et al., 1953; Geman and Geman, 1984; Gelfand and Smith, 1990). The Tobit model specification is particularly amenable to MCMC, since the latent process has a closed form full conditional distribution. To guarantee that Metropolis acceptance rates are within an acceptable range, pilot adaptation is used (Banerjee et al., 2003). To improve MCMC mixing, the VF data were scaled by 10, time was parameterized in years, and the dissimilarity metric was scaled by 100. The MCMC sampler is implemented in the R package spCP.

A.1. Change point likelihood

The change point likelihood can be written for the latent Tobit process,

$$\begin{aligned} & f\left\{Y_t(\mathbf{s}_i) \mid \beta_0(\mathbf{s}_i), \beta_1(\mathbf{s}_i), \lambda_0(\mathbf{s}_i), \lambda_1(\mathbf{s}_i), \theta(\mathbf{s}_i)\right\} \\ &= \prod_{t=1}^v \prod_{i=1}^m \left[N\left(\beta_0(\mathbf{s}_i), e^{2\lambda_0(\mathbf{s}_i)}\right) \right]^{1(0 \leq x_t \leq \theta(\mathbf{s}_i))} \\ & \quad \times \left[N\left(\beta_0(\mathbf{s}_i) + \beta_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}, e^{2[\lambda_0(\mathbf{s}_i) + \lambda_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}]} \right) \right]^{1(\theta(\mathbf{s}_i) \leq x_t \leq x_v)}. \end{aligned}$$

This likelihood is informative, however for computational purposes we will write the likelihood in matrix form. The latent responses can be written more concisely using an indicator variable. Define the event $\gamma_t(\mathbf{s}_i) = 1\{\theta(\mathbf{s}_i) \leq x_t \leq x_v\}$. Then,

$$\begin{aligned} Y_t(\mathbf{s}_i) &= \beta_0(\mathbf{s}_i) + \beta_1(\mathbf{s}_i) \{x_t - \theta(\mathbf{s}_i)\} \gamma_t(\mathbf{s}_i) + \epsilon_t(\mathbf{s}_i) \\ &= x_t(\mathbf{s}_i; \theta) \boldsymbol{\beta}(\mathbf{s}_i) + \epsilon_t(\mathbf{s}_i), \end{aligned}$$

where $\boldsymbol{\beta}(\mathbf{s}_i) = \{\beta_0(\mathbf{s}_i), \beta_1(\mathbf{s}_i)\}^T$ and $x_t(\mathbf{s}_i; \theta) = [1, \{x_t - \theta(\mathbf{s}_i)\} \gamma_t(\mathbf{s}_i)]$. Now, we determine the vector of latent observations at each time point, defining $\mathbf{Y}_t = \{Y_t(\mathbf{s}_1), Y_t(\mathbf{s}_2), \dots, Y_t(\mathbf{s}_m)\}^T$,

$$\mathbf{Y}_t = X_t(\theta) \boldsymbol{\beta} + \boldsymbol{\epsilon}_t,$$

where $\boldsymbol{\epsilon}_t = \{\epsilon_t(\mathbf{s}_1), \epsilon_t(\mathbf{s}_2), \dots, \epsilon_t(\mathbf{s}_m)\}^T$, the regression coefficient vector $\boldsymbol{\beta} = \{\boldsymbol{\beta}(\mathbf{s}_1)^T, \boldsymbol{\beta}(\mathbf{s}_2)^T, \dots, \boldsymbol{\beta}(\mathbf{s}_m)^T\}^T$ and $X_t(\theta)$ is an $m \times 2m$ block diagonal matrix with blocks entries $x_t(\mathbf{s}_i; \theta)$. Finally, we write the full latent process, $\mathbf{Y} = \{\mathbf{Y}_1^T, \mathbf{Y}_2^T, \dots, \mathbf{Y}_v^T\}^T$,

$$\mathbf{Y} = \mathbf{X}(\theta) \boldsymbol{\beta} + \boldsymbol{\epsilon},$$

where $\boldsymbol{\epsilon} = \{\boldsymbol{\epsilon}_1^T, \boldsymbol{\epsilon}_2^T, \dots, \boldsymbol{\epsilon}_v^T\}^T$ and $\mathbf{X}(\theta)$ is an $mv \times 2m$ matrix stacking the $X_t(\theta)$ matrices.

Having written the mean process in matrix form, we turn our attention to the variance process. Define, $\sigma_t^2 = \{\sigma_t^2(\mathbf{s}_1), \sigma_t^2(\mathbf{s}_2), \dots, \sigma_t^2(\mathbf{s}_m)\}^T$. Then $\sigma^2 = \{(\sigma_1^2)^T, (\sigma_2^2)^T, \dots, (\sigma_v^2)^T\}^T$. The variance process can be written using the same form as the mean process,

$$\log\{\sigma\} = \mathbf{X}(\theta)\boldsymbol{\lambda} \implies \sigma^2 = \exp\{2\mathbf{X}(\theta)\boldsymbol{\lambda}\},$$

where $\boldsymbol{\lambda} = \{\boldsymbol{\lambda}(\mathbf{s}_1)^T, \boldsymbol{\lambda}(\mathbf{s}_2)^T, \dots, \boldsymbol{\lambda}(\mathbf{s}_m)^T\}^T$, with $\boldsymbol{\lambda}(\mathbf{s}_i) = \{\lambda_0(\mathbf{s}_i), \lambda_1(\mathbf{s}_i)\}^T$. Define the full likelihood covariance as $\boldsymbol{\Omega} = \text{Diag}(\sigma^2)$. Then, we can write the matrix likelihood as follows,

$$\begin{aligned} f(\mathbf{Y}|\boldsymbol{\beta}, \boldsymbol{\lambda}, \theta) &\sim \text{MVN}(\mathbf{X}(\theta) \boldsymbol{\beta}, \boldsymbol{\Omega}) \\ \log\{\sigma\} &= \mathbf{X}(\theta) \boldsymbol{\lambda}, \end{aligned}$$

where $\theta = \{\theta(\mathbf{s}_1), \theta(\mathbf{s}_2), \dots, \theta(\mathbf{s}_m)\}^T$.

Having written the likelihood in a convenient matrix form, we now write the full data likelihood which will be used to find the full conditionals. Note that the likelihood does not include θ , but rather is parameterized to include the generating parameters,

$$\begin{aligned} f(\mathbf{Y}, \boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) &\propto f(\mathbf{Y}|\boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}) \\ &\times f(\{\boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}\}|\boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) f(\boldsymbol{\delta}) f(\boldsymbol{\Sigma}) f(\alpha), \end{aligned}$$

where $\boldsymbol{\eta} = \{\eta(\mathbf{s}_1), \eta(\mathbf{s}_2), \dots, \eta(\mathbf{s}_m)\}^T$.

A.2. Full conditional distributions

1. Gibbs Step for $\boldsymbol{\delta}$:

Note that $(\mathbf{1}_m \otimes \boldsymbol{\delta}) = (\mathbf{1}_m \otimes \mathbf{I}_5) \boldsymbol{\delta} = \mathbf{Z}_\delta \boldsymbol{\delta}$. Then,

$$\begin{aligned} f(\boldsymbol{\delta}|\cdot) &\propto f(\boldsymbol{\phi}|\boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \times f(\boldsymbol{\delta}) \\ &\propto \exp\left\{-\frac{1}{2}\left[\{\boldsymbol{\phi} - \mathbf{Z}_\delta \boldsymbol{\delta}\}^T (\mathbf{Q}(\alpha) \otimes \boldsymbol{\Sigma}^{-1}) \{\boldsymbol{\phi} - \mathbf{Z}_\delta \boldsymbol{\delta}\} + \frac{\boldsymbol{\delta}^T \boldsymbol{\delta}}{\kappa^2}\right]\right\} \\ &\propto \exp\left\{-\frac{1}{2}\left[\boldsymbol{\delta}^T \left\{\mathbf{Z}_\delta^T (\mathbf{Q}(\alpha) \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{Z}_\delta + \frac{\mathbf{I}_5}{\kappa^2}\right\} \boldsymbol{\delta} \right. \right. \\ &\quad \left. \left. - 2\boldsymbol{\delta}^T \{\mathbf{Z}_\delta^T \boldsymbol{\Omega}^{-1} \boldsymbol{\phi}\}\right]\right\} \\ &\sim \text{MVN}(\mathbb{E}_\delta, \mathbb{C}_\delta), \end{aligned}$$

where $\mathbb{C}_\delta = \left\{\mathbf{Z}_\delta^T (\mathbf{Q}(\alpha) \otimes \boldsymbol{\Sigma}^{-1}) \mathbf{Z}_\delta + \frac{\mathbf{I}_5}{\kappa^2}\right\}^{-1}$, $\mathbb{E}_\delta = \mathbb{C}_\delta \{\mathbf{Z}_\delta^T \boldsymbol{\Omega}^{-1} \boldsymbol{\phi}\}$.

2. Gibbs Step for β :

It is not immediately obvious, but with some manipulation we can find a closed form full conditional distribution for β . The key is that we can find an analytical form of the distribution $f(\beta|\lambda, \eta, \delta, \Sigma, \alpha)$. Recall that, $\mathbb{C}(\phi) = \mathbf{Q}(\alpha)^{-1} \otimes \Sigma$. Define, $k = \{1, 2\}$ and $j = \{3, 4, 5\}$ as the indices of the components of β in the vector of random effects ϕ_i . The mean can be partitioned as,

$$\mathbf{1}_m \otimes \delta = \begin{bmatrix} \mathbf{1}_m \otimes \delta_k \\ \mathbf{1}_m \otimes \delta_j \end{bmatrix}.$$

The covariance matrix can then be partitioned in block form as follows,

$$\Sigma = \begin{bmatrix} \Sigma_{kk} & \Sigma_{kj} \\ \Sigma_{jk} & \Sigma_{jj} \end{bmatrix},$$

so that,

$$\mathbf{Q}(\alpha)^{-1} \otimes \Sigma = \begin{bmatrix} \mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{kk} & \mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{kj} \\ \mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{jk} & \mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{jj} \end{bmatrix}.$$

Define ϕ_{ij} as the original ϕ_i vector, only keeping the indices included in j and $\phi^j = \{\phi_{1j}^T, \phi_{2j}^T, \dots, \phi_{mj}^T\}^T$. Using the properties of the multivariate normal distribution we can compute the moments of the distribution $f(\beta|\lambda, \eta, \delta, \Sigma, \alpha) \sim \text{MVN}(\mathbb{E}_{k|j}, \mathbb{C}_{k|j})$,

$$\begin{aligned} \mathbb{E}_{k|j} &= \mathbb{E}[\beta|\lambda, \eta, \delta, \Sigma, \alpha] = (\mathbf{1}_m \otimes \delta_k) + [\mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{kj}] \\ &\quad \times [\mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{jj}]^{-1} (\phi^j - (\mathbf{1}_m \otimes \delta_j)) \\ &= (\mathbf{1}_m \otimes \delta_k) + [\mathbf{I}_n \otimes \Sigma_{kj} \Sigma_{jj}^{-1}] \\ &\quad \times (\phi^j - (\mathbf{1}_m \otimes \delta_j)), \end{aligned}$$

and

$$\begin{aligned} \mathbb{C}_{k|j} &= \mathbb{C}(\beta|\lambda, \eta, \delta, \Sigma, \alpha) = [\mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{kk}] \\ &\quad - [\mathbf{I}_n \otimes \Sigma_{kj} \Sigma_{jj}^{-1}] [\mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{jk}] \\ &= [\mathbf{Q}(\alpha)^{-1} \otimes \Sigma_{k|j}], \end{aligned}$$

where $\Sigma_{k|j} = (\Sigma_{kk} - \Sigma_{kj} \Sigma_{jj}^{-1} \Sigma_{jk})$. Now, we can derive the full conditional for β .

$$\begin{aligned} f(\beta|\cdot) &\propto f(\mathbf{Y}|\beta, \lambda, \eta) \times f(\beta|\lambda, \eta, \delta, \Sigma, \alpha) \\ &\propto \exp \left\{ -\frac{1}{2} \left[\{\mathbf{Y} - \mathbf{X}(\theta) \beta\}^T \Omega^{-1} \{\mathbf{Y} - \mathbf{X}(\theta) \beta\} \right. \right. \\ &\quad \left. \left. + (\beta - \mathbb{E}_{k|j})^T \mathbb{C}_{k|j}^{-1} (\beta - \mathbb{E}_{k|j}) \right] \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left[\beta^T \{\mathbf{X}(\theta)^T \Omega^{-1} \mathbf{X}(\theta) + \mathbb{C}_{k|j}^{-1}\} \beta \right. \right. \\ &\quad \left. \left. - 2\beta^T \{\mathbf{X}(\theta)^T \Omega^{-1} \mathbf{Y} + \mathbb{C}_{k|j}^{-1} \mathbb{E}_{k|j}\} \right] \right\} \\ &\sim \text{MVN}(\mathbb{E}_\beta, \mathbb{C}_\beta), \end{aligned}$$

where the covariance is $\mathbb{C}_\beta = \{\mathbf{X}(\theta)^T \Omega^{-1} \mathbf{X}(\theta) + \mathbb{C}_{k|j}^{-1}\}^{-1}$, and the mean is $\mathbb{E}_\beta = \mathbb{C}_\beta \{\mathbf{X}(\theta)^T \Omega^{-1} \mathbf{Y} + \mathbb{C}_{k|j}^{-1} \mathbb{E}_{k|j}\}$.

3. Metropolis Step for $\lambda_0(\mathbf{s}_i)$:

Each $\lambda_0(\mathbf{s}_i)$, $i = 1, \dots, m$, is sampled using a Metropolis step given the following quantity that is proportional to the full conditional density,

$$\begin{aligned} f\{\lambda_0(\mathbf{s}_i) | \cdot\} &\propto f(\mathbf{Y} | \boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}) \times f(\boldsymbol{\phi} | \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \\ &\propto \prod_{t=1}^v [N(\beta_0(\mathbf{s}_i), e^{2\lambda_0(\mathbf{s}_i)})]^{1-\tau_t(\mathbf{s}_i)} \\ &\quad \times [N(\beta_0(\mathbf{s}_i) + \beta_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}, e^{2[\lambda_0(\mathbf{s}_i) + \lambda_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}])}]^{\tau_t(\mathbf{s}_i)} \\ &\quad \times f(\lambda_0 | \{\boldsymbol{\beta}, \boldsymbol{\lambda}_1, \boldsymbol{\eta}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha), \end{aligned}$$

where $f(\lambda_0 | \{\boldsymbol{\beta}, \boldsymbol{\lambda}_1, \boldsymbol{\eta}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \sim \text{MVN}(\mathbb{E}_{k|j}, \mathbb{C}_{k|j})$, with $k = \{3\}$ and $j = \{1, 2, 4, 5\}$. Here $\boldsymbol{\lambda}_j = \{\lambda_j(\mathbf{s}_1), \lambda_j(\mathbf{s}_2), \dots, \lambda_j(\mathbf{s}_m)\}^T$, $j = 0, 1$.

4. Metropolis Step for $\lambda_1(\mathbf{s}_i)$:

Each $\lambda_1(\mathbf{s}_i)$, $i = 1, \dots, m$, is sampled using a Metropolis step given the following quantity that is proportional to the full conditional density,

$$\begin{aligned} f\{\lambda_1(\mathbf{s}_i) | \cdot\} &\propto f(\mathbf{Y} | \boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}) \times f(\boldsymbol{\phi} | \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \\ &\propto \prod_{t=1}^v [N(\beta_0(\mathbf{s}_i), e^{2\lambda_0(\mathbf{s}_i)})]^{1-\tau_t(\mathbf{s}_i)} \\ &\quad \times [N(\beta_0(\mathbf{s}_i) + \beta_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}, e^{2[\lambda_0(\mathbf{s}_i) + \lambda_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}])}]^{\tau_t(\mathbf{s}_i)} \\ &\quad \times f(\lambda_1 | \{\boldsymbol{\beta}, \boldsymbol{\lambda}_0, \boldsymbol{\eta}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha), \end{aligned}$$

where $f(\lambda_1 | \{\boldsymbol{\beta}, \boldsymbol{\lambda}_0, \boldsymbol{\eta}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \sim \text{MVN}(\mathbb{E}_{k|j}, \mathbb{C}_{k|j})$, with $k = \{4\}$ and $j = \{1, 2, 3, 5\}$.

5. Metropolis Step for $\eta(\mathbf{s}_i)$:

Each $\eta(\mathbf{s}_i)$, $i = 1, \dots, m$, is sampled using a Metropolis step given the following quantity that is proportional to the full conditional density,

$$\begin{aligned} f\{\eta(\mathbf{s}_i) | \cdot\} &\propto f(\mathbf{Y} | \boldsymbol{\beta}, \boldsymbol{\lambda}, \boldsymbol{\eta}) \times f(\boldsymbol{\phi} | \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \\ &\propto \prod_{t=1}^v [N(\beta_0(\mathbf{s}_i), e^{2\lambda_0(\mathbf{s}_i)})]^{1-\tau_t(\mathbf{s}_i)} \\ &\quad \times [N(\beta_0(\mathbf{s}_i) + \beta_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}, e^{2[\lambda_0(\mathbf{s}_i) + \lambda_1(\mathbf{s}_i)\{x_t - \theta(\mathbf{s}_i)\}])}]^{\tau_t(\mathbf{s}_i)} \\ &\quad \times f(\eta | \{\boldsymbol{\beta}, \boldsymbol{\lambda}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha), \end{aligned}$$

where $f(\eta | \{\boldsymbol{\beta}, \boldsymbol{\lambda}\}, \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \sim \text{MVN}(\mathbb{E}_{k|j}, \mathbb{C}_{k|j})$, with $k = \{5\}$ and $j = \{1, 2, 3, 4\}$.

6. Metropolis Step for α :

We transform α to the real line to facilitate sampling. Define a new parameter $\Delta = h(\alpha) = \log\left(\frac{\alpha - a_\alpha}{b_\alpha - \alpha}\right)$, such that $h^{-1}(\Delta) = (b_\alpha \exp\{\Delta\} + a_\alpha) / (1 + \exp\{\Delta\})$ and $\left|\frac{\partial}{\partial \Delta} h^{-1}(\Delta)\right| \propto \exp\{\Delta\} / (1 + \exp\{\Delta\})^2$. Now we can sample from the transformed proposal distribution, $\Delta^* \sim N(\Delta, \delta)$, where δ is a tuning parameter. Then we can obtain a proposal of α , $\alpha^* = h^{-1}(\Delta^*)$.

$$f(\alpha | \cdot) \propto f(\boldsymbol{\phi} | \boldsymbol{\delta}, \boldsymbol{\Sigma}, \alpha) \times f(\alpha) \times \left| \frac{\partial}{\partial \Delta} h^{-1}(\Delta) \right|.$$

7. Gibbs Step for $\boldsymbol{\Sigma}$:

Define $\boldsymbol{\Phi}_{5 \times n} = \{\boldsymbol{\beta}_0, \boldsymbol{\beta}_1, \boldsymbol{\lambda}_0, \boldsymbol{\lambda}_1, \boldsymbol{\eta}\}^T$. Note that $\boldsymbol{\phi} = \text{vec}(\boldsymbol{\Phi})$. Therefore, using the following properties of the Kronecker product and trace, $(B^T \otimes A)\text{vec}(X) = \text{vec}(AXB)$ and $\text{tr}(AB) =$

$\text{vec}(A)^T \text{vec}(B)$ and defining $\mathbf{M} = \delta \mathbf{1}_m^T$, we may rewrite the MCAR prior, $f(\boldsymbol{\phi} | \delta, \Sigma, \alpha)$, as follows,

$$\begin{aligned}
 & \propto \exp \left\{ -\frac{1}{2} (\boldsymbol{\phi} - (\mathbf{1}_m \otimes \delta))^T (\mathbf{Q}(\alpha) \otimes \Sigma^{-1}) (\boldsymbol{\phi} - (\mathbf{1}_m \otimes \delta)) \right\} \\
 & = \exp \left\{ -\frac{1}{2} (\text{vec}(\boldsymbol{\Phi}) - \text{vec}(\mathbf{M}))^T (\mathbf{Q}(\alpha) \otimes \Sigma^{-1}) (\text{vec}(\boldsymbol{\Phi}) - \text{vec}(\mathbf{M})) \right\} \\
 & = \exp \left\{ -\frac{1}{2} \text{vec}(\boldsymbol{\Phi} - \mathbf{M})^T (\mathbf{Q}(\alpha) \otimes \Sigma^{-1}) \text{vec}(\boldsymbol{\Phi} - \mathbf{M}) \right\} \\
 & = \exp \left\{ -\frac{1}{2} \text{vec}(\boldsymbol{\Phi} - \mathbf{M})^T \text{vec}(\Sigma^{-1}(\boldsymbol{\Phi} - \mathbf{M})\mathbf{Q}(\alpha)) \right\} \\
 & = \exp \left\{ -\frac{1}{2} \text{tr}((\boldsymbol{\Phi} - \mathbf{M})^T \Sigma^{-1}(\boldsymbol{\Phi} - \mathbf{M})\mathbf{Q}(\alpha)) \right\} \\
 & = \exp \left\{ -\frac{1}{2} \text{tr}((\boldsymbol{\Phi} - \mathbf{M})\mathbf{Q}(\alpha)(\boldsymbol{\Phi} - \mathbf{M})^T \Sigma^{-1}) \right\} \\
 & = \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}_{\boldsymbol{\Phi}, \alpha} \Sigma^{-1}) \right\},
 \end{aligned}$$

where $\mathbf{S}_{\boldsymbol{\Phi}, \alpha} = (\boldsymbol{\Phi} - \mathbf{M})\mathbf{Q}(\alpha)(\boldsymbol{\Phi} - \mathbf{M})^T$. Now, the full conditional for Σ is straight forward,

$$\begin{aligned}
 f(\Sigma | \cdot) & \propto f(\boldsymbol{\phi} | \delta, \Sigma, \alpha) \times f(\Sigma) \\
 & \propto |\mathbf{Q}(\alpha) \otimes \Sigma^{-1}|^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}_{\boldsymbol{\Phi}, \alpha} \Sigma^{-1}) \right\} \\
 & \quad \times |\Sigma|^{\frac{-(\xi+p+1)}{2}} \exp \left\{ -\frac{1}{2} \text{tr}(\Psi \Sigma^{-1}) \right\} \\
 & \propto (|\Sigma^{-1}|^m |\mathbf{Q}(\alpha)|^p)^{\frac{1}{2}} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}_{\boldsymbol{\Phi}, \alpha} \Sigma^{-1}) \right\} \\
 & \quad \times |\Sigma|^{\frac{-(\xi+p+1)}{2}} \exp \left\{ -\frac{1}{2} \text{tr}(\Psi \Sigma^{-1}) \right\} \\
 & \propto |\Sigma^{-1}|^{\frac{m}{2}} \exp \left\{ -\frac{1}{2} \text{tr}([\mathbf{S}_{\boldsymbol{\Phi}, \alpha} + \Psi] \Sigma^{-1}) \right\} |\Sigma|^{\frac{-(\xi+p+1)}{2}} \\
 & \propto |\Sigma|^{-\frac{m}{2}} \exp \left\{ -\frac{1}{2} \text{tr}([\mathbf{S}_{\boldsymbol{\Phi}, \alpha} + \Psi] \Sigma^{-1}) \right\} |\Sigma|^{\frac{-(\xi+p+1)}{2}} \\
 & \propto |\Sigma|^{\frac{-(m+\xi+p+1)}{2}} \exp \left\{ -\frac{1}{2} \text{tr}([\mathbf{S}_{\boldsymbol{\Phi}, \alpha} + \Psi] \Sigma^{-1}) \right\} \\
 & \sim \text{Inverse-Wishart}(m + \xi, \mathbf{S}_{\boldsymbol{\Phi}, \alpha} + \Psi).
 \end{aligned}$$

8. Gibbs Step for $Y_t(\mathbf{s}_i)$:

The latent process $Y_t(\mathbf{s}_i)$ is sampled from their full conditional distribution of,

$$f\{Y_t(\mathbf{s}_i) | Y_t^*(\mathbf{s}_i), \Omega\} \propto f\{Y_t^*(\mathbf{s}_i) | Y_t(\mathbf{s}_i)\} \times f\{Y_t(\mathbf{s}_i) | \Omega\},$$

where Ω is a vector of the model parameters. Then, we know that there are two possibilities, $Y_t^*(\mathbf{s}_i) = 0$ or $Y_t^*(\mathbf{s}_i) = Y_t(\mathbf{s}_i)$. If $Y_t^*(\mathbf{s}_i) = 0$, then,

$$f\{Y_t^*(\mathbf{s}_i) | Y_t(\mathbf{s}_i)\} = P[Y_t^*(\mathbf{s}_i) = 0 | Y_t(\mathbf{s}_i)] = \begin{cases} 1 & Y_t(\mathbf{s}_i) \leq 0, \\ 0 & Y_t(\mathbf{s}_i) > 0. \end{cases}$$

Therefore, $f\{Y_t(\mathbf{s}_i) | Y_t^*(\mathbf{s}_i), \Omega\} \propto 1\{Y_t(\mathbf{s}_i) \leq 0\} f\{Y_t(\mathbf{s}_i) | \Omega\}$. Furthermore, note that when $Y_t^*(\mathbf{s}_i) = Y_t(\mathbf{s}_i)$,

$$f\{Y_t(\mathbf{s}_i) | Y_t^*(\mathbf{s}_i), \Omega\} = f\{Y_t(\mathbf{s}_i) | Y_t(\mathbf{s}_i), \Omega\} = Y_t(\mathbf{s}_i).$$

Table C.1
Summary diagnostics for estimation of the observed CPs in the simulation study. The diagnostics included are bias, MSE, and EC (95% nominal) and are presented for each of the models that are capable of estimating the continuous observed CP. Standard errors (SE) correspond to the variation across all VF locations and simulations. Each reported estimate is based on 1000 simulated datasets.

Metric	Model	Estimate (SE)
Bias	Non-spatial CP (C)	0.147 (0.171)
	Non-spatial CP (L)	0.058 (0.123)
	Spatial CP	−0.001 (0.008)
MSE	Non-spatial CP (C)	0.102 (0.065)
	Non-spatial CP (L)	0.083 (0.087)
	Spatial CP	0.002 (0.002)
EC	Non-spatial CP (C)	0.33 (0.461)
	Non-spatial CP (L)	0.94 (0.146)
	Spatial CP	0.99 (0.018)

Then, we see that the full conditional for the latent variable is,

$$f\left\{Y_t(\mathbf{s}_i) \mid Y_t^*(\mathbf{s}_i), \Omega\right\} \sim \begin{cases} N^-\left(\mu_t(\mathbf{s}_i), \sigma_t^2(\mathbf{s}_i)\right) & Y_t^*(\mathbf{s}_i) = 0, \\ Y_t^*(\mathbf{s}_i) & Y_t^*(\mathbf{s}_i) = Y_t(\mathbf{s}_i). \end{cases}$$

where, $N^-(\mu, \sigma^2)$ specifies a truncated normal distribution (truncated above by zero).

9. Repeat steps 1–8 until convergence has been achieved and an adequate number of posterior samples have been obtained post-convergence.

Appendix B. Prediction of future change points

See Video B.1 for a map of VF progression over the course of follow-up and into the future. The print version of this manuscript shows a still frame that matches the right panel of Fig. 3.

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.spasta.2019.02.001>.

Appendix C. Estimation of the observed change point

In Section 5, we established the importance of estimating the observed CPs for identifying the moment of progression. In Section 6, we established through simulation the importance of the latent CP specification and spatial structure for estimating both the observed and latent CPs. As a result, it is important to understand the behavior of each modeling feature in estimating CPs, especially at the censored spatial locations. Therefore, we develop a second simulation study that is based on a fixed ϕ , so that the CPs are fixed across all simulated datasets. The realization of ϕ was chosen to produce CPs that occur across all stages of follow-up. To generate the ϕ , we use the same data generating scheme as Section 6, with the same values of Σ and α , with $\delta = (25, -30, 1, 0.5, 0.5)$. Based on the value of ϕ , 1000 datasets are simulated. In this simulation, the VF visit times are defined as all 21 visits from Section 6.

Each of the models were fit to the simulated datasets yielding an estimate of the observed CPs at each VF location. In Fig. C.1, the density estimates of the distribution of estimated observed CPs (posterior means) are presented for each model across the VF (excluding PLR). Furthermore, the true observed CP is presented as a vertical black dotted line and the threshold sensitivity values for an example simulated dataset are presented as gray dots. This figure provides important information regarding the introduced model. First, in VF locations where the true CP is near the middle of the monitoring period (i.e., easier to identify), Spatial CP has minimal bias and has the narrowest bounds (representative of the variability in the posterior mean estimates). In these settings, the other models generally have small bias, however the density estimates are much wider.

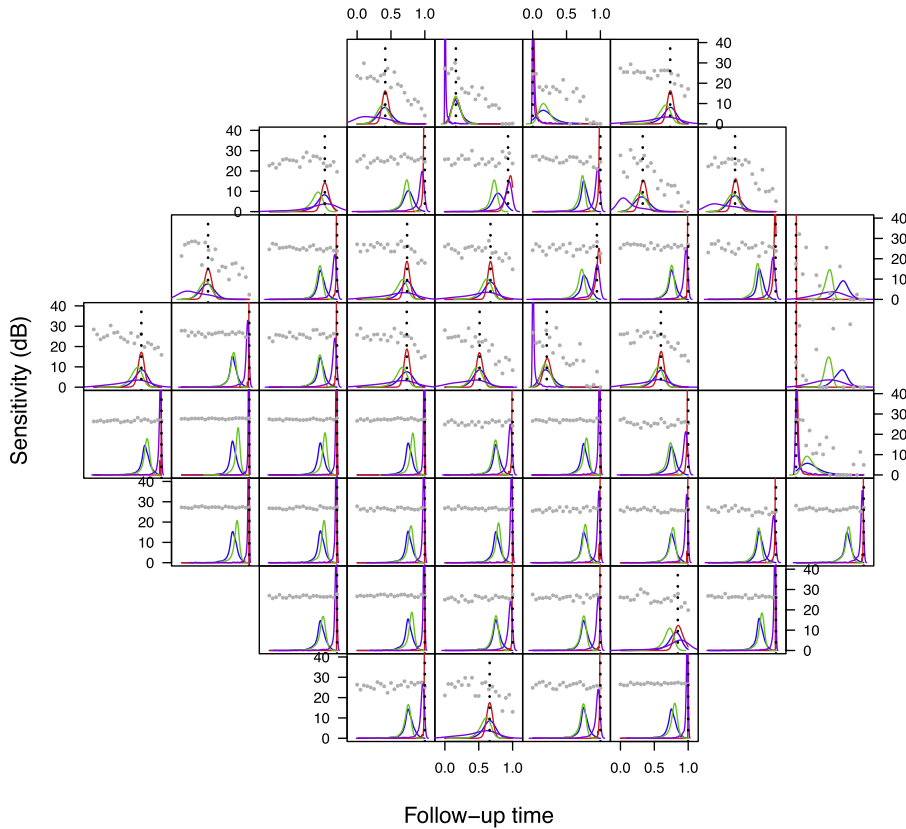


Fig. C.1. Simulation study comparing the performance of Spatial CP and non-spatial CP models in estimating the CPs across the VF. The true value of the CP is presented as a dotted black line and the threshold sensitivity values for an example simulated dataset are presented as gray dots. The posterior mean estimates of the CPs across all 1000 datasets for Spatial CP, Non-spatial CP (D), Non-spatial CP (C), and Non-spatial CP (L) are presented as density estimates in red, blue, green, and purple lines, respectively. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

In locations where the true CP is near or beyond the beginning or end of the monitoring period, Spatial CP is clearly superior. The two models that have no latent CP process are incapable of approaching the edges of the follow-up times, with apparent bias. Non-spatial CP (L) provides an improvement, allowing the density estimate to approach the bounds. However, it is important to note that this appears to be a trade-off with wider densities in locations with a true CP in the middle of follow-up. The only model that is capable of producing acceptable CP estimation in the middle and bounds of follow-up is Spatial CP. Interestingly, at certain VF locations at the boundaries, Spatial CP is capable of producing a point mass, meaning all posterior means are the same.

To formalize the results seen in Fig. C.1, summary diagnostics are presented in Table C.1. The results for Non-spatial CP (D) are excluded, since the results are not directly comparable. The results in Table C.1 confirm the trends observed in Fig. C.1 as the bias and MSE are improved with the introduction of a latent process and then spatial structure. Spatial CP has superior bias and MSE, -0.001 and 0.002 , respectively, and has the smallest standard errors, indicating less variability in estimation across the simulated datasets. The EC is only acceptable with the inclusion of the latent CP process. Non-spatial CP (C) has an EC of 0.33 , which is clearly a result of all of the true CPs near the bounds. Spatial CP slightly overestimates the nominal coverage, which corresponds to the

Table D.1

Sensitivity analysis to various prior distribution settings in the analysis of visual field data and glaucoma progression risk. For each different combination of prior values, the analysis of Table 1 is repeated for the Spatial CP model. Settings of interest include (i) the upper bound for α , b_α , and (ii) κ^2 .

b_α	κ^2	Model fit		Prediction	Diagnostic		
		DIC	p_D	MSPE	AIC	AUC	p -value
0.25	100	−54.03	66.58	0.13	208.83	0.68	<0.001
0.50	100	−55.73	67.41	0.13	215.59	0.67	0.005
0.75	100	−55.48	66.17	0.13	216.59	0.65	0.008
0.25	1,000	−50.98	67.03	0.13	211.32	0.68	<0.001
0.50	1,000	−50.03	68.21	0.13	209.43	0.70	<0.001
0.75	1,000	−58.04	66.82	0.13	212.18	0.65	0.001
0.25	10,000	−48.75	66.68	0.13	212.59	0.70	0.001
0.50	10,000	−60.51	67.60	0.13	212.30	0.68	0.001
0.75	10,000	−61.67	66.94	0.13	214.48	0.67	0.003

simulation results from Section 6 and makes sense since the observed CP is a censored parameter. Overall, this additional simulation provides a closer look into the estimation of the observed CP, demonstrating the behavior of each model at various locations over follow-up. These results again illuminate the importance of Spatial CP for estimating the observed CP, especially in models with CPs near the bounds of follow-up, which are commonly the interesting patients.

Appendix D. Sensitivity analyses

D.1. Choice of prior settings

In Section 5, we applied the introduced methodology to VF data. We assess sensitivity of those results to two modeling assumptions: (i) the upper bound of α , b_α , and (ii) the variance for the mean vector, δ , κ^2 . We refit the Spatial CP model for every patient eye under each sensitivity setting and display the results in Table D.1.

In the main analysis, the upper bound for α is set so that the smallest dissimilarity metric can yield an adjacency of 0.50 (Lee and Mitchell, 2011), while values of 0.25 and, 0.75 are used for comparison in the sensitivity analysis. To explore the sensitivity of the results to the weakly-informative variance of the mean vector, δ , we compare the original variance of 1000 to 100 and 10,000. Although there is slight variation in the results, the conclusions drawn from each setting do not differ in absolute, and more importantly, compared to the Non-spatial models from the original Table 1. Overall, from Table D.1, it is clear that the results are generally robust to the modeling assumptions and prior choices made in the VF data setting.

D.2. Thinning of posterior samples

In the main analyses, the MCMC sampler is run for 250,000 iterations after a burn-in period of 10,000 iterations and is thinned to a final sample size of 10,000. The thinning was performed in order to reduce autocorrelation in the posterior samples and to reduce the data storage size of the samples collected after fitting the model to each patient eye. To assess the sensitivity of model results to the thinning of posterior samples, we present trace plots and posterior means and standard deviations for an example patient (the patient included in the spCP R package). In the top row of Fig. D.1, we present the trace plot of α using the thinned set of posterior samples (a) and using the complete set of posterior samples (b). The trace plots do not show obvious differences, indicating that thinning is likely having only minimal impact on the resulting inference. Furthermore, in the bottom row of Fig. D.1, we present scatterplots of the posterior means and standard deviations for all model parameters from the example patient model fit, using the set of

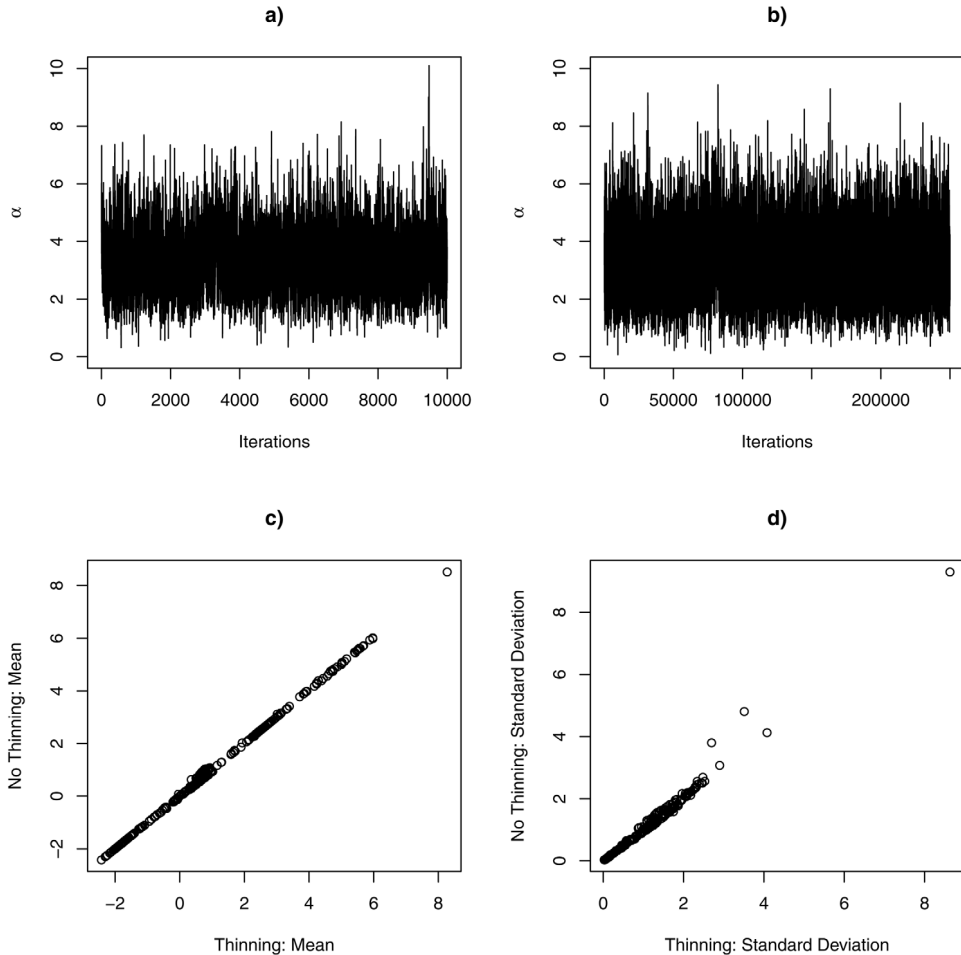


Fig. D.1. Demonstrating the impact of thinning the posterior samples on inference. The top row shows trace plots of α with, (a) the thinned set of posterior samples, and (b) the complete set of posterior samples, for the example patient included in the R package *spCP*. The bottom row shows the scatterplots of the posterior (c) means, and (d) standard deviations for all model parameters for the example patient using the different sets of posterior samples.

thinned posterior samples (c), and the complete set of posterior samples (d). The correlation for each of these scatterplots is nearly 1, indicating that thinning is having negligible impact on the final inference in this example.

References

- Artes, P.H., Chauhan, B.C., 2005. Longitudinal changes in the visual field and optic disc in glaucoma. *Prog. Retinal Eye Res.* 24 (3), 333–354.
- Banerjee, S., Gelfand, A.E., Carlin, B.P., 2003. *Hierarchical Modeling and Analysis for Spatial Data*. CRC Press.
- Beckage, B., Joseph, L., Belisle, P., Wolfson, D.B., Platt, W.J., 2007. Bayesian change-point analyses in ecology. *New Phytologist* 174 (2), 456–467.
- Berchuck, S.I., Mwanza, J., Warren, J.L., 2018. Diagnosing glaucoma progression with visual field data using a spatiotemporal boundary detection method. *J. Amer. Statist. Assoc.* <http://dx.doi.org/10.1080/01621459.2018.1537911>.
- Besag, J., 1974. Spatial interaction and the statistical analysis of lattice systems. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 36 (2), 192–236.

- Betz-Stablein, B.D., Morgan, W.H., House, P.H., Hazelton, M.L., 2013. Spatial modeling of visual field data for assessing glaucoma progression.. *Investig. Ophthalmol. Vis. Sci.* 54 (2), 1544–1553.
- Brook, D., 1964. On the distinction between the conditional probability and the joint probability approaches in the specification of nearest-neighbour systems. *Biometrika* 51 (3–4), 481–483.
- Bryan, S.R., Eilers, P.H., Lesaffre, E.M., Lemij, H.G., Vermeer, K.A., 2015. Global visit effects in point-wise longitudinal modeling of glaucomatous visual fields. *Investig. Ophthalmol. Vis. Sci.* 56 (8), 4283–4289.
- Cai, B., Lawson, A.B., McDermott, S., Aelion, C.M., 2016. A Bayesian semiparametric approach with change points for spatial ordinal data. *Stat. Methods Med. Res.* 25 (2), 644–658.
- Carlin, B.P., Gelfand, A.E., Smith, A.F., 1992. Hierarchical Bayesian analysis of changepoint problems. *J. R. Stat. Soc. Ser. C. Appl. Stat.* 41 (2), 389–405.
- Crabb, D.P., Fitzke, F.W., McNaught, A.I., Edgar, D.F., Hitchings, R.A., 1997. Improving the prediction of visual field progression in glaucoma using spatial processing. *Ophthalmology* 104 (3), 517–524.
- Davson, H., 2012. *Physiology of the Eye*. Elsevier.
- Eddelbuettel, D., François, R., Allaire, J., Chambers, J., Bates, D., Ushey, K., 2011. Rcpp: Seamless R and C++ integration. *J. Stat. Software* 40 (8), 1–18.
- Fitzke, F.W., Hitchings, R.A., Poinoosawmy, D., McNaught, A.I., Crabb, D.P., 1996. Analysis of visual field progression in glaucoma. *Br. J. Ophthalmol.* 80 (1), 40–48.
- Garway-Heath, D.F., Poinoosawmy, D., Fitzke, F.W., Hitchings, R.A., 2000. Mapping the visual field to the optic disc in normal tension glaucoma eyes. *Ophthalmology* 107 (10), 1809–1815.
- Gelfand, A.E., Smith, A.F., 1990. Sampling-based approaches to calculating marginal densities. *J. Amer. Statist. Assoc.* 85 (410), 398–409.
- Gelfand, A.E., Vounatsou, P., 2003. Proper multivariate conditional autoregressive models for spatial data analysis. *Biostatistics* 4 (1), 11–15.
- Gelman, A., Carlin, J.B., Stern, H.S., Rubin, D.B., 2014. *Bayesian Data Analysis*, Vol. 2. Chapman & Hall/CRC Boca Raton, FL, USA.
- Geman, S., Geman, D., 1984. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. Pattern Anal. Mach. Intell. PAMI-6* (6), 721–741.
- Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In: Bernardo, J.M., Berger, J., Dawid, A.P., Smith, J.F.M. (Eds.), *Bayesian Statistics 4*. Oxford University Press, Oxford, pp. 169–193.
- Heijl, A., Leske, M.C., Bengtsson, B., Hyman, L., Bengtsson, B., Hussein, M., 2002. Reduction of intraocular pressure and glaucoma progression: results from the Early Manifest Glaucoma Trial. *Arch. Ophthalmol.* 120 (10), 1268–1279.
- Heijl, A., Lindgren, G., Olsson, J., 1987. A package for the statistical analysis of visual fields. In: *Seventh International Visual Field Symposium*, Amsterdam, September 1986. Springer, pp. 153–168.
- Hinkley, D.V., 1970. Inference about the change-point in a sequence of random variables. *Biometrika* 57 (1), 1–17.
- Jampel, H.D., Singh, K., Lin, S.C., Chen, T.C., Francis, B.A., Hodapp, E., Samples, J.R., Smith, S.D., 2011. Assessment of visual function in glaucoma: a report by the American Academy of Ophthalmology. *Ophthalmology* 118 (5), 986–1002.
- Jay, J., Murdoch, J., 1993. The rate of visual field loss in untreated primary open angle glaucoma. *Br. J. Ophthalmol.* 77 (3), 176–178.
- Jin, X., Banerjee, S., Carlin, B.P., 2007. Order-free co-regionalized areal data models with application to multiple-disease mapping. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 69 (5), 817–838.
- Jin, X., Carlin, B.P., Banerjee, S., 2005. Generalized hierarchical multivariate CAR models for areal data. *Biometrics* 61 (4), 950–961.
- Julious, S.A., 2001. Inference and estimation in a changepoint regression problem. *J. R. Stat. Soc. Ser. D* 50 (1), 51–61.
- Katz, J., Gilbert, D., Quigley, H.A., Sommer, A., 1997. Estimating progression of visual field loss in glaucoma. *Ophthalmology* 104 (6), 1017–1025.
- Lee, D., Mitchell, R., 2011. Boundary detection in disease mapping studies. *Biostatistics* 15 (3), 457–469.
- Leroux, B.G., Lei, X., Breslow, N., 2000. Estimation of disease rates in small areas: A new mixed model for spatial dependence. In: *Statistical Models in Epidemiology, the Environment, and Clinical Trials*. Springer, pp. 179–191.
- MacNab, Y.C., Gustafson, P., 2007. Regression b-spline smoothing in Bayesian disease mapping: with an application to patient safety surveillance. *Stat. Med.* 26 (24), 4455–4474.
- Majumdar, A., Gelfand, A.E., Banerjee, S., 2005. Spatio-temporal change-point modeling. *J. Statist. Plann. Inference* 130 (1–2), 149–166.
- Mardia, K., 1988. Multi-dimensional multivariate Gaussian Markov random fields with application to image processing. *J. Multivariate Anal.* 24 (2), 265–284.
- Menzeffricke, U., 1981. A Bayesian analysis of a change in the precision of a sequence of independent normal random variables at an unknown time point. *J. R. Stat. Soc. Ser. C. Appl. Stat.* 30 (2), 141–146.
- Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E., 1953. Equation of state calculations by fast computing machines. *J. Chem. Phys.* 21 (6), 1087–1092.
- Minin, V.N., Dorman, K.S., Fang, F., Suchard, M.A., 2007. Phylogenetic mapping of recombination hotspots in human immunodeficiency virus via spatially smoothed change-point processes. *Genetics* 175 (4), 1773–1785.
- Nicolela, M.T., McCormick, T.A., Drance, S.M., Ferrier, S.N., LeBlanc, R.P., Chauhan, B.C., 2003. Visual field and optic disc progression in patients with different types of optic disc damage: a longitudinal prospective study. *Ophthalmology* 110 (11), 2178–2184.
- Page, E.S., 1954. Continuous inspection schemes. *Biometrika* 41 (1), 100–115.

- Quandt, R.E., 1958. The estimation of the parameters of a linear regression system obeying two separate regimes. *J. Amer. Statist. Assoc.* 53 (284), 873–880.
- R Core Team, 2016. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL: <https://www.R-project.org/>.
- Russell, R.A., Crabb, D.P., Malik, R., Garway-Heath, D.F., 2012. The relationship between variability and sensitivity in large-scale longitudinal visual field data. *Investig. Ophthalmol. Vis. Sci.* 53 (10), 5985–5990.
- Smith, A.C., Hudson, M.-A.R., Downes, C.M., Francis, C.M., 2015. Change points in the population trends of aerial-insectivorous birds in North America: synchronized in time across species and regions. *PLoS One* 10 (7), e0130768.
- Spiegelhalter, D.J., Best, N.G., Carlin, B.P., Van Der Linde, A., 2002. Bayesian measures of model complexity and fit. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 64 (4), 583–639.
- Tham, Y.-C., Li, X., Wong, T.Y., Quigley, H.A., Aung, T., Cheng, C.-Y., 2014. Global prevalence of glaucoma and projections of glaucoma burden through 2040: a systematic review and meta-analysis. *Ophthalmology* 121 (11), 2081–2090.
- Tobin, J., 1958. Estimation of relationships for limited dependent variables. *Econometrica* 26 (1), 24–36.
- Vianna, J.R., Chauhan, B.C., 2015. How to detect progression in glaucoma. *Prog. Brain Res.* 221, 135–158.
- Wagner, T., Midway, S.R., 2014. Modeling spatially varying landscape change points in species occurrence thresholds. *Ecosphere* 5 (11), 1–16.
- Warren, J.L., Mwanza, J.-C., Tanna, A.P., Budenz, D.L., 2016. A statistical model to analyze clinician expert consensus on glaucoma progression using spatially correlated visual field data. *Transl. Vis. Sci. Technol.* 5 (4), 14–14.
- Warren, J.L., Pingali, S.C., Weinberger, D.M., 2017. Spatial variability in the persistence of pneumococcal conjugate vaccine-targeted pneumococcal serotypes among adults. *Epidemiology* 28 (1), 119.
- Yu, Q., Scribner, R., Carlin, B., Theall, K., Simonsen, N., Ghosh-Dastidar, B., Cohen, D., Mason, K., 2008. Multilevel spatio-temporal dual changepoint models for relating alcohol outlet destruction and changes in neighbourhood rates of assaultive violence. *Geospatial Health* 2 (2), 161.
- Zhu, H., Russell, R.A., Saunders, L.J., Ceccon, S., Garway-Heath, D.F., Crabb, D.P., 2014. Detecting changes in retinal function: analysis with non-stationary Weibull error regression and spatial enhancement (ANSWERS). *PLoS One* 9 (1), 1–11.