



Diagnosing Glaucoma Progression With Visual Field Data Using a Spatiotemporal Boundary Detection Method

Samuel I. Berchuck, Jean-Claude Mwanza & Joshua L. Warren

To cite this article: Samuel I. Berchuck, Jean-Claude Mwanza & Joshua L. Warren (2019) Diagnosing Glaucoma Progression With Visual Field Data Using a Spatiotemporal Boundary Detection Method, Journal of the American Statistical Association, 114:527, 1063-1074, DOI: [10.1080/01621459.2018.1537911](https://doi.org/10.1080/01621459.2018.1537911)

To link to this article: <https://doi.org/10.1080/01621459.2018.1537911>

 View supplementary material [↗](#)

 Published online: 01 Apr 2019.

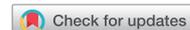
 Submit your article to this journal [↗](#)

 Article views: 543

 View related articles [↗](#)

 View Crossmark data [↗](#)

 Citing articles: 2 View citing articles [↗](#)



Diagnosing Glaucoma Progression With Visual Field Data Using a Spatiotemporal Boundary Detection Method

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

^aDepartment of Statistical Science and Forge, Duke University, Durham, NC; ^bDepartment of Ophthalmology, University of North Carolina-Chapel Hill, Chapel Hill, NC; ^cDepartment of Biostatistics, Yale University, New Haven, CT

ABSTRACT

Diagnosing glaucoma progression is critical for limiting irreversible vision loss. A common method for assessing glaucoma progression uses a longitudinal series of visual fields (VFs) acquired at regular intervals. VF data are characterized by a complex spatiotemporal structure due to the data generating process and ocular anatomy. Thus, advanced statistical methods are needed to make clinical determinations regarding progression status. We introduce a spatiotemporal boundary detection model that allows the underlying anatomy of the optic disc to dictate the spatial structure of the VF data across time. We show that our new method provides novel insight into vision loss that improves diagnosis of glaucoma progression using data from the Vein Pulsation Study Trial in Glaucoma and the Lions Eye Institute trial registry. Simulations are presented, showing the proposed methodology is preferred over existing spatial methods for VF data. Supplementary materials for this article are available online and the method is implemented in the R package `womb1R`.

ARTICLE HISTORY

Received July 2017
Revised October 2018

KEYWORDS

Areal wombling; Bayesian methods; Conditional autoregressive models; Dissimilarity metric.

1. Introduction

Glaucoma is a leading cause of blindness worldwide, with a prevalence of 4% in the population aged 40–80 (Tham et al. 2014). The most common form of glaucoma is primary open-angle glaucoma (POAG). The biological basis of the disease is not fully defined, however, the most significant risk factor for POAG is elevated intraocular pressure (IOP). Elevated IOP can be treated with eye drops, surgery, or laser. If this condition goes untreated the optic nerve may be damaged, resulting in permanent vision loss. Since visual impairment caused by glaucoma is irreversible and efficient treatments exist, early detection of the disease is essential. As such, patients diagnosed with glaucoma are monitored for disease progression even if they are receiving treatment, because the role of the treatment is to slow the progression. Determining if the disease is progressing remains one of the most challenging aspects of glaucoma management, since it is difficult to distinguish true progression from variability due to natural degradation or noise (Vianna and Chauhan 2015). Numerous techniques have been developed to monitor progression, but there is currently no consensus as to which method is best. In this study, we focus on visual field (VF) testing.

A VF test is a psychophysical procedure that assesses a patient's field of vision. The test results in a two-dimensional map that represents the level of eyesight uniformly across the retina. Glaucoma patients normally receive biannual VF tests and have follow-up lasting numerous years (Chauhan et al. 2008). The collection of VF data results in a longitudinal series of spatially referenced measurements that exhibit a complex spa-

tiotemporal structure. The spatial surface of the VF is observed on a lattice (i.e., uniform areal data), however, it is a complex projection of the underlying optic disc and exhibits anatomically induced spatial dependencies. In particular, localized damage to the optic disc can result in clinically deterministic deterioration across the VF (Quigley et al. 1992). Incorporating this nonstandard spatial dependence structure into our methodology is a priority for properly analyzing these data.

There are comparable methodological complexities that arise in Alzheimer's disease, attention deficit hyperactivity disorder, and multiple sclerosis, all related to white-matter connectivity of the brain (He, Chen, and Evans 2008; Konrad and Eickhoff 2010). However, complex brain imaging data are generally analyzed using point-referenced statistical models as opposed to areal data models (Bowman et al. 2008). The point-referenced framework has a rich theory that accounts for nonstandard spatial dependencies, mainly through the assumptions of nonstationarity (Sampson 2010) and anisotropy (Ecker and Gelfand 2003). Castruccio, Ombao, and Genton (2016) model anatomical regions of interest of the brain in fMRI data using these assumptions in a study of stroke rehabilitation, while another study accounts for brain connectivity in Alzheimer's patients (Thompson et al. 2004).

The literature surrounding complex spatial dependencies is less developed in the areal data setting with spatiotemporal methods even less common. One reason for this is that stationarity and isotropy (i.e., correlation as a function of distance alone) cannot be defined for areal data due to the contrasting definition of spatial proximity between the two frameworks. When

CONTACT Joshua L. Warren  joshua.warren@yale.edu  Department of Biostatistics, Yale University, New Haven, CT 06520.

Color versions of one or more of the figures in the article can be found online at www.tandfonline.com/r/JASA.

 Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JASA.

 These materials were reviewed for reproducibility.

© 2019 American Statistical Association

analyzing areal data, spatial similarity is typically dictated by a local neighborhood structure (Banerjee, Gelfand, and Carlin 2003). Over time, modifications to the neighborhood structure have been proposed, and in this manuscript we work within the boundary detection framework. An extension of directional gradients from point-referenced theory (Banerjee and Gelfand 2006), boundary detection was originally developed to identify boundaries on geographical maps in the context of disease mapping (Ma and Carlin 2007). The inherited method of modifying the local neighborhood structure provides an intricate framework for introducing complex spatial structure on the VF. We introduce a novel spatiotemporal boundary detection model that allows the underlying anatomy of the optic disc to dictate the spatial structure of the VF across time and show that it offers new and valuable information for improved progression detection through analysis of data from the Vein Pulsation Study Trial in Glaucoma (VPSG) and the Lions Eye Institute trial registry.

This article is outlined as follows. Section 2 details the data generating mechanism for VF data and the setting of glaucoma progression diagnostics. We briefly review spatial boundary detection methods in Section 3. In Section 4, our newly developed statistical methodology is described. We apply our method to a dataset of VF tests from glaucoma patients in Section 5 and compare its performance to an existing boundary detection method via simulation study in Section 6. We conclude in Section 7 with a discussion.

2. Visual Field Data

The VF is the spatial array of visual sensations that the brain perceives as vision (Smythies 1996). The most common technique for testing the VF is standard automated perimetry (SAP, Chauhan et al. 2008). In this study, we analyze data acquired with the Humphrey Field Analyzer-II (HFA-II) (Carl Zeiss Meditec Inc., Dublin, CA). The VF data generating process is displayed in Figure 1. We follow a single observation throughout the figure, presented as a diamond. In Figure 1(a), a patient (the first author) is tested on a HFA-II. SAP constructs a VF map by assessing a patient's response to varying levels of light. Patients are instructed to focus on a central fixation point as light is introduced randomly in a preceding manner over a grid on the VF. As light is observed, the patient presses a button and the current light intensity is recorded (Chauhan et al. 2008). The process is repeated until the entire VF is tested. In the figure, the first author stares at the background of the machine, waiting until he observes the stimulus to press the buzzer.

The HFA-II measures 54 test points, however, two of these correspond to a natural blind spot corresponding to the optic disc, resulting in 52 informative points. The points that correspond to the blind spot are highlighted yellow (Figure 1c). Figure 1(b) presents the anatomy of the eye. The retina is a light-sensitive layer at the back of the eye that absorbs stimulus from the HFA-II and transmits information to the brain through the optic disc along the optic nerve. The optic nerve is a bundle of more than one million fibers that carry visual information in the form of electrical signal from the retina to the brain. The retinal nerve fiber layer (RNFL), Figure 1(c), is made of retinal ganglion

cell (RGC) axons. All the axons converge at the optic disc, which is the departure point of the optic nerve.

These RGC axons are responsible for encoding visual information. The ganglion cells disperse across the retina, but are mostly concentrated in the center of the retina, and use photoreceptors to transmit information to the brain along their axons (Davson 2012). Both the RGCs and their axons may die progressively as a result of elevated IOP. In particular, damage to specific regions of the optic disc corresponds to loss of RGCs whose axons enter the damaged region. Thus, vision loss across the VF and the corresponding damage to the optic disc are the result of the death of RGCs and their axons. Furthermore, correlation between two test points on the VF is dependent on the spatial proximity that their underlying nerve fibers enter the optic disc. This indicates that variability in the neighborhood structure of the VF is possibly indicative of progression. The nerve fibers (not to be confused with the thicker and more sparse blood vessels) are shown extending throughout the retina with the VF test points represented by black dots.

Each test location has underlying fibers that track across the RNFL and enter the optic disc at a particular angle. Garway-Heath et al. (2000) studied the relationship between the VF test points and the underlying RNFL. They quantified the relationship between the VF and optic disc by estimating the angle that each test location's underlying fiber enters the optic disc, measured in degrees ($^{\circ}$). The measure ranges from 0° to 360° , where 0° is designated at the 9-o'clock position (right eye) and angles are counted counter clockwise. The Garway-Heath angle for the example location is 77° (Figure 1(d)).

In the course of a VF test all 52 informative locations are assessed, resulting in a grid of test values (Figure 1(e)). The two locations corresponding to the blind spot are gray. SAP measures the differential light sensitivity (DLS) across the VF. The measurement represents a contrast between the background of the machine, normally white, and the light stimulus. The intensity of the stimulus is initially similar to the background, but as the intensity increases the contrast grows and the probability of detecting the stimulus increases. The intensity of the stimulus is measured in Apostilbs (asb), where larger values represent a greater intensity. The HFA-II is capable of outputting intensities ranging from one (similar to the background) to 10,000 asb. All stimuli that are not detected by 10,000 asb are censored, due to the constraints of the machine. In practice, these intensities are converted to decibels (dB), where $\text{dB} = 40 - 10 \log_{10}(\text{asb})$. This inverts the scale, such that DLS values near 40 indicate good vision, while values near zero represent possible blindness. In the course of monitoring a patient with glaucoma, VF testing is performed on a regular basis and a longitudinal series of VFs is obtained (Figure 1(f)).

3. Spatial Boundary Detection

VF data exhibit a complex spatiotemporal structure that is characterized by localized spatial dependencies dictated by the underlying anatomy of the optic disc. These data are generated over a lattice, a subset of areal data in spatial literature. In spatial statistics, the foundational assumption states that dependence between observations weakens as the distance between locations

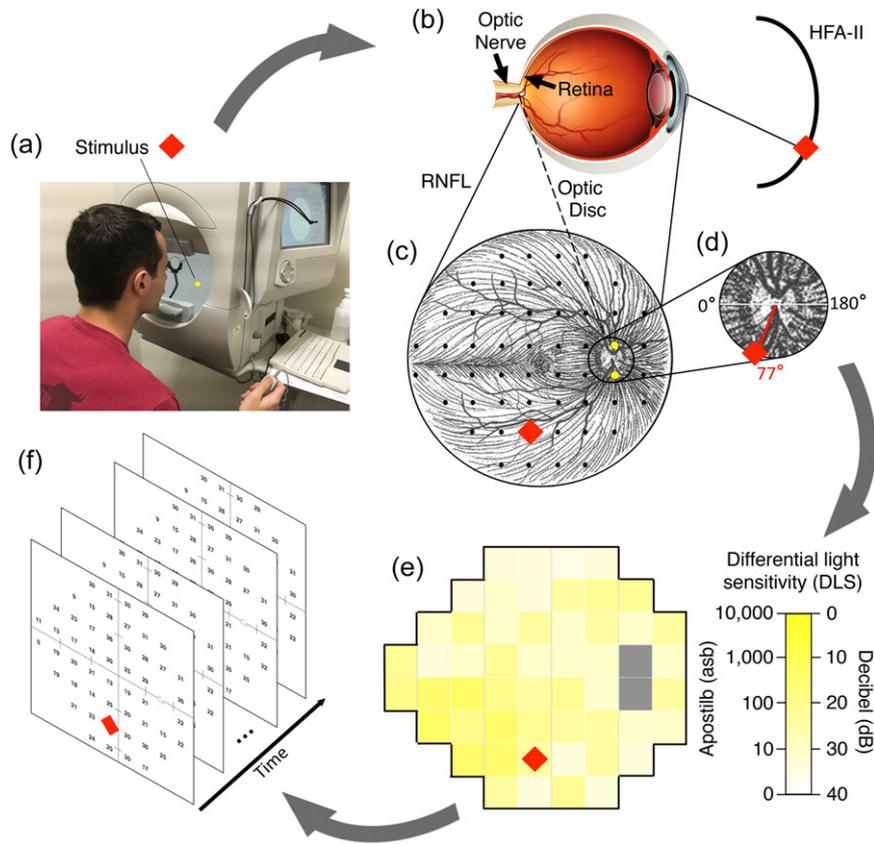


Figure 1. Demonstrating the VF data generating mechanism, (a) A patient (the first author) participates in a VF examination. (b) The HFA-II stimulus is absorbed by the retina and transmitted through the optic disc, and along the optic nerve to the brain. (c) Each VF test point corresponds to a location on the underlying RNFL. (d) The corresponding nerve fiber enters the optic disc at 77°. (e) Each VF test produces a map that shows the intensity a stimulus is detected at each test location. The gray locations represent the blind spot. (f) Over time the patient accrues VF tests.

increases (Gelfand et al. 2010). In areal data models, this assumption often manifests through Gaussian Markov random fields (GMRF) that induce neighborhood dependence across the region (Geman and Geman 1984). A common GMRF is the conditional autoregressive (CAR) process (Besag 1974).

The CAR model achieves spatial smoothing through random effects, φ_i at location i for $i = 1, \dots, n$, with spatial structure defined through the set of neighborhood adjacencies, $\{w_{ij}\}$. These adjacencies are fixed, such that $w_{ij} = 1(i \sim j)$, where $1(\cdot)$ is the indicator function and $i \sim j$ is the event that locations i and j share a border ($w_{ii} = 0$ for all i). This specification induces an elegant conditional distribution for each random effect. We present a general form of the CAR process originally introduced by Leroux, Lei, and Breslow (2000),

$$\varphi_i | \varphi_{-i}, \mu, \tau^2 \sim N \left(\frac{\rho \sum_{j=1}^n w_{ij} \varphi_j + (1 - \rho) \mu}{\rho \sum_{j=1}^n w_{ij} + 1 - \rho}, \frac{\tau^2}{\rho \sum_{j=1}^n w_{ij} + 1 - \rho} \right), \tag{1}$$

where $\varphi_{-i} = (\varphi_1, \dots, \varphi_{i-1}, \varphi_{i+1}, \dots, \varphi_n)^T$. Note that setting $\rho = 1$ reduces Equation (1) to the standard intrinsic CAR process. The mean is a weighted average of the neighbors with variance decreasing inversely with the number of neighbors. The standard CAR model provides an attractive representation and is flexible in its ability to model smooth spatial processes. However, it can be limited in settings where spatial structure

is fragmented into localized regions due to $\{w_{ij}\}$ being fixed. In the areal data setting, a flexible class of models called boundary detection can be used to remedy this issue (Banerjee, Gelfand, and Carlin 2003).

Boundary detection was originally explored by Womble (1951), but has gained a niche in the context of disease mapping. The motivation for boundary detection is to identify regions on the spatial surface where there are sharp changes in the response value (Jacquez and Greiling 2003). In disease mapping, these boundaries can take many forms, for example, geographic obstacles such as mountain ranges or socioeconomic boundaries caused by pockets of increased poverty. Standard methods attempt to control for disjoint spatial regions by including covariates in the mean structure of the CAR process. This technique can be effective in producing highly variable spatial surfaces, but is limited in producing truly localized spatial smoothing (Banerjee, Gelfand, and Carlin 2003). Boundary detection improves on this naive approach by carefully considering the form of the adjacencies.

Initially, boundary detection methods were parameterized for use in disease mapping, defining boundaries as a function of the difference in standardized incidence ratios (Lu and Carlin 2005). This method is limited, since it can be difficult to have knowledge of boundaries a priori and does not generalize outside of disease mapping. Numerous methods in the boundary detection literature treat the adjacencies as random variables and construct hierarchical models to estimate the adjacency

matrix (Lu et al. 2007; Ma and Carlin 2007), and even provide extensions to the spatiotemporal setting (Rushworth, Lee, and Sarran 2017). However, inference from these models can be highly sensitive to prior specifications on certain parameters (Li et al. 2015). Furthermore, these methods introduce $(n)(n-1)/2$ additional random variables leading to potential identifiability issues. Li, Banerjee, and McBean (2011) proposed another class of methods that enumerate all possible permutations of the adjacencies in parallel models, using the Bayesian information criterion to choose between them. This class of methods was formalized when Lee, Rushworth, and Sahu (2014) introduced a novel joint prior distribution for the spatial random effects and adjacency matrix. This prior has been extended to spatiotemporal models (Lee and Mitchell 2014).

A final class of methods models the adjacencies using dissimilarity metrics. The method introduced in Lee and Mitchell (2011) generalizes the form of Equation (1) to allow for the adjacency weights to be modeled as a function of a small number of regression parameters, $\alpha = (\alpha_1, \dots, \alpha_q)^T$. According to Lee and Mitchell (2011), “boundaries in the risk surface are likely to occur between populations that are very different because homogeneous populations should have similar risk profiles.” They define q nonnegative dissimilarity metrics $z_{ij} = (z_{ij1}, \dots, z_{ijq})^T$, where $z_{ijk} = |z_{ik} - z_{jk}|$ for $k = 1, \dots, q$. The q covariates, z_{ik} at location i , drive detection of boundaries and are characterized by their importance in defining the neighborhood structure. The choice of q is problem specific and based on the availability of useful explanatory information for describing the boundaries. The adjacencies are modeled as follows,

$$w_{ij}(\alpha) = 1(i \sim j)1\left(\exp\left\{-z_{ij}^T \alpha\right\} \geq 0.5\right), \quad (2)$$

where each α_k is constrained to be nonnegative so that a larger dissimilarity metric indicates a higher likelihood of a boundary (or zero weight). In this model, ρ is fixed at 0.99 to force the spatial correlation structure to be determined locally by $\{w_{ij}(\alpha)\}$, rather than globally by μ . This form has many appealing properties that make it amenable to boundary detection. In particular, if there are no adjacencies (i.e., $w_{ij}(\alpha) = 0$ for all $i \neq j$) the conditional mean and variance are still defined. The method proposed by Lee and Mitchell (2011) provides a parsimonious framework for introducing localized smoothing.

4. Methods

Following the approach of Lee and Mitchell (2011), we propose modeling localized spatial correlation through a set of weights $\{w_{ij}(\alpha_t)\}$ as a parsimonious function of dissimilarity metrics and their regression parameters. However, we propose extending the framework to account for spatiotemporal localized smoothing and therefore define $\alpha_t = (\alpha_{t1}, \dots, \alpha_{tq})^T$, for $t = 1, \dots, v$. With appropriate temporal dependency structures in place, this specification allows for localized smoothing in instances of true temporal correlation. Inference for this model is based on Markov chain Monte Carlo (MCMC) simulation, and a description of the algorithm is given in the online supplementary materials. Spatiotemporal models are computationally intensive, so the MCMC algorithm is implemented using Rcpp

(Eddelbuettel et al. 2011) and is available from the R package `womb1R` (R Core Team 2016).

4.1. Observational Model

We begin by describing our new methodology generally before applying it to VF data in Section 5. Let Y_{it} denote an observation from spatial location i at time t , $i = 1, \dots, n_t$, for $t = 1, \dots, v$. The number of locations can vary over time. Define $\varphi_t = (\varphi_{1t}, \dots, \varphi_{n_t t})^T$, with φ_{-it} missing the i th entry. The observational model is given by,

$$Y_{it} | \vartheta_{it}, \zeta \stackrel{\text{ind}}{\sim} f(Y_{it} | \vartheta_{it}, \zeta) \quad \text{for } i = 1, \dots, n_t, t = 1, \dots, v, \quad (3)$$

$$g(\vartheta_{it}) = \varphi_{it},$$

$$\varphi_{it} | \varphi_{-it}, \mu_t, \tau_t^2, \alpha_t \stackrel{\text{ind}}{\sim} \text{N} \left(\frac{\rho \sum_{j=1}^{n_t} w_{ij}(\alpha_t) \varphi_{jt} + (1 - \rho) \mu_t}{\rho \sum_{j=1}^{n_t} w_{ij}(\alpha_t) + 1 - \rho}, \frac{\tau_t^2}{\rho \sum_{j=1}^{n_t} w_{ij}(\alpha_t) + 1 - \rho} \right).$$

The parameter ϑ_{it} describes the distribution of Y_{it} and our novel spatiotemporal random effect, φ_{it} , is introduced as a linear predictor of $g(\vartheta_{it})$, with $g(\cdot)$ a link function. Finally, ζ is a vector of variance (or nuisance) parameters, for example, the overdispersion parameter in the negative binomial distribution. This modeling framework is general and accommodates generalized linear mixed models (GLMM). The GLMM setting can be obtained by setting $\vartheta_{it} = \mathbb{E}[Y_{it} | \vartheta_{it}]$. Due to the general specification, our methodology can be used to induce spatiotemporal localized smoothing in a general areal data setting, such as disease mapping.

The random effect for φ_{it} represents an extension of the Lee and Mitchell (2011) specification with temporally referenced parameters, μ_t , τ_t^2 , and α_t (referred to as observational level parameters). The ρ parameter acts as a gauge of present spatial variation, where $\rho = 0$ corresponds to global independence and $\rho \rightarrow 1$ implies strong spatial correlation. In the same vein as Lee and Mitchell (2011), we fix $\rho = 0.99$ to guarantee that spatial correlation can be determined locally by the set of weights $\{w_{ij}(\alpha_t)\}$. We fully explore the impact of this decision through simulation as described in Section 5.6 and in the online supplementary materials. Overall, we find that the results are robust to this choice.

The conditional distributions of the random effect can be written jointly using Brook’s lemma (Banerjee, Gelfand, and Carlin 2003), $\varphi_t | \mu_t, \tau_t^2, \alpha_t \stackrel{\text{ind}}{\sim} \text{MVN}(\mu_t \mathbf{1}_{n_t}, \tau_t^2 \mathbf{Q}(\alpha_t)^{-1})$, $t = 1, \dots, v$, where $\mathbf{Q}(\alpha_t) = [\rho \mathbf{W}^*(\alpha_t) + (1 - \rho) \mathbf{I}_{n_t}]$, and $\mathbf{1}_n$ and \mathbf{I}_n are an n dimensional column of ones and identity matrix, respectively. The matrix $\mathbf{W}^*(\alpha_t)$ has diagonal elements $w_{ii}^*(\alpha_t) = \sum_{j=1}^{n_t} w_{ij}(\alpha_t)$ and off-diagonal elements $w_{ij}^*(\alpha_t) = -w_{ij}(\alpha_t)$.

4.2. Neighborhood Model

We use a similar framework as Lee and Mitchell (2011) to model the adjacency weights by writing them as a function of

dissimilarity metrics, z_{ij} . The weights are defined as follows,

$$w_{ij}(\alpha_t) = 1(i \sim j) \exp \left\{ -z_{ij}^T \alpha_t \right\}. \tag{4}$$

Unlike the original Lee and Mitchell (2011) specification (Equation (2)), the weights are not forced to be binary. We do specify the components of α_t to be nonnegative, forcing the adjacencies in the open unit interval. These constraints on α_{tk} yield intuitive interpretations at extreme values. As $\alpha_{tk} \rightarrow \infty$ the adjacencies become zero, resulting in an independence model, while $\alpha_{tk} \rightarrow 0$ reduces the adjacencies to a standard CAR process (Equation (1)). This new specification changes how neighbors share information. It is best understood through the conditional mean from Equation (1), which becomes a weighted average of neighbors under our new definition of an adjacency, versus a simple average.

4.3. Temporal Model

From the joint specification of the random effects, we see that spatial structure is introduced through the covariance of φ_t at each time point and there is conditional independence ($\varphi_{t_1} \perp \varphi_{t_2} | \mu_t, \tau_t^2, \alpha_t : t = t_1, t_2$). To induce temporal dependence between the φ_t , we specify a separable temporal structure on the observational level parameters. Define,

$$\theta = [\theta_{\cdot 1} \cdots \theta_{\cdot v}] = \begin{bmatrix} \theta_{1\cdot} \\ \theta_{2\cdot} \\ \vdots \\ \theta_{2+q\cdot} \end{bmatrix} = \begin{bmatrix} \mu_1 & \cdots & \mu_v \\ \log(\tau_1) & \cdots & \log(\tau_v) \\ \log(\alpha_{11}) & \cdots & \log(\alpha_{v1}) \\ \vdots & & \vdots \\ \log(\alpha_{1q}) & \cdots & \log(\alpha_{vq}) \end{bmatrix}.$$

Using properties of the vectorization function, $\text{vec}(\cdot)$, and the Kronecker product, \otimes , a separable process is specified such that

$$\text{vec}(\theta) | \delta, \mathbf{T}, \phi \sim \text{MVN}(\mathbf{1}_v \otimes \delta, \Sigma(\phi) \otimes \mathbf{T}). \tag{5}$$

This process yields elegant interpretations for the row and column moments of θ . From the moments, we see that δ is a constant that corresponds to the mean of the observational level parameters at time t . The matrix \mathbf{T} can be interpreted as the local covariance of the observational level parameters at each time t .

Finally, the correlation matrix, $\Sigma(\phi)$, represents the temporal correlation of each observational level parameter over time. Due to the properties of the separable covariance, each of the observational level parameters has the same temporal structure dictated by the form of $\Sigma(\phi)$. The form of $\Sigma(\phi)$ is general such that any standard temporal correlation function may be specified (e.g., exponential or first-order autoregressive). The parameter ϕ acts as a temporal tuning parameter describing the strength of correlation across time and can be interpreted within the context of each specific temporal structure.

4.4. Specifying Hyperprior Distributions

To complete the model specification, we define hyperprior distributions for the introduced parameters such that

$$\begin{aligned} \delta &\sim \text{MVN}(\mu_\delta, \Omega_\delta), & \mathbf{T} &\sim \text{Inverse-Wishart}(\xi, \Psi), \\ \phi &\sim \text{Uniform}(a_\phi, b_\phi). \end{aligned}$$

The choice of the entries in μ_δ can be informative or noninformative depending on the context and user. It is important to judiciously consider the entries of Ω_δ . For simplification, we detail a situation where Ω_δ is diagonal, $\Omega_\delta = \text{diag}(1000, 1000, \nu_1, \dots, \nu_q)$. The entries in Ω_δ that are of importance are those that correspond to $\log(\alpha_{tk})$, since the large variances for μ_t and $\log(\tau_t)$ induce approximately flat priors. More care is needed in specifying ν_1, \dots, ν_v . These hyperprior variances are chosen for purposes of regularization, to encourage $\log(\alpha_{tk})$ to be in a realistic range. Regularization is a common use of Bayesian priors (Gelman and Shalizi 2013). In particular, these variances are chosen so that α_{tk} do not become larger than α_k^* , such that $[\alpha_k^* : \exp\{-\alpha_k^* z_k\} = 0.5]$ with $z_k = \min_{i,j} \{z_{ijk}\}$. This condition comes from Equation (4), where we isolate each dissimilarity metric individually.

For our prior on \mathbf{T} , we use an inverse-Wishart distribution with degrees of freedom $\xi = (q + 2) + 1$ and scale matrix, $\Psi = \mathbf{I}_{q+2}$. This prior is appealing since it induces marginally uniform priors on the correlations of \mathbf{T} and allows for the diagonals to be weakly informative (Gelman et al. 2014). Finally, we specify the hyperprior distribution for the temporal tuning parameter ϕ for correlation structures with one parameter. The bounds for ϕ cannot be specified arbitrarily since it is important to account for the magnitude of time elapsed. We specify the following conditions for finding the bounds, $[a_\phi : [\Sigma(a_\phi)]_{t,t'} = 0.95, |x_t - x_{t'}| = x_{\max}]$ and $[b_\phi : [\Sigma(b_\phi)]_{t,t'} = 0.01, |x_t - x_{t'}| = x_{\min}]$, where x_{\min} and x_{\max} are the minimum and maximum temporal differences between visits. These conditions state that the lower bound of ϕ is small enough so that the greatest length of time between time points can yield a correlation of 0.95 and the upper bound is set so that the shortest length of time between time points can reach 0.01. These conditions were specified so that ϕ can dictate a temporal correlation close to independence ($\phi \rightarrow b_\phi$) or strong correlation ($\phi \rightarrow a_\phi$), resulting in a weakly informative prior distribution.

4.5. Prediction

Once posterior samples have been obtained, prediction is often a priority. In particular, obtaining samples from the posterior predictive distribution (PPD) is of interest, $f(\mathbf{Y}_{v+1} | \mathbf{Y})$, where $\mathbf{Y}_t = (Y_{1t}, \dots, Y_{nt})^T$ and $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_v)^T$. We express the PPD as an integral $\int_{\Omega} f(\mathbf{Y}_{v+1} | \Omega, \mathbf{Y}) f(\Omega | \mathbf{Y}) d\Omega$ and then further partition the integral,

$$\int_{\Omega} \underbrace{f(\mathbf{Y}_{v+1} | g^{-1}(\varphi_{v+1}), \xi)}_1 \underbrace{f(\varphi_{v+1} | \theta_{\cdot v+1})}_2 \underbrace{f(\theta_{\cdot v+1} | \theta, \delta, \mathbf{T}, \phi)}_3 \underbrace{f(\xi, \theta, \delta, \mathbf{T}, \phi | \mathbf{Y})}_4 d\Omega, \tag{6}$$

where $\Omega = (\varphi_{v+1}, \theta_{\cdot v+1}, \xi, \theta, \delta, \mathbf{T}, \phi)$. The convenient form of Equation (6) is a function of four known densities that are defined as a consequence of the methodology introduced in Section 4. As such, the PPD can be obtained through composition sampling (Tanner 1996). This theory is presented with one future time point, but is easily generalized to multiple. Full prediction theory details are presented in the online supplementary materials.

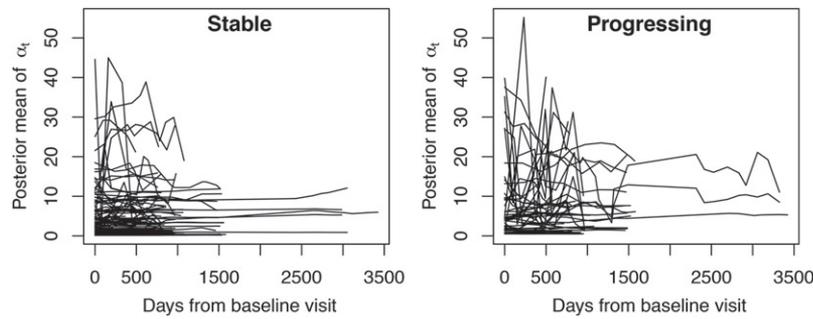


Figure 3. Posterior mean estimates of α_{tGH} plotted across time for all patients, subset by progression status defined at the end of the study.

of variation (CV) of α_{tGH} for $t = 1, \dots, \nu$. The CV is the ratio of the SD and mean and is an ideal summary metric, because it accounts for the variability in a parameter while standardizing by its mean, allowing it to be comparable over populations. This metric is novel clinically as α_{tGH} does not describe mean trend in DLS over time, but rather is representative of optic disc damage and changes in the spatial covariance structure across time. We refer to our metric based on the spatiotemporal (ST) model as “ST CV.” In Figure 3, we display posterior mean estimates of α_{tGH} over time for all patients, subset by progression status. The figure suggests that there is more variability in the estimates across time for the progressing patients generally, and that there is not a clear trend in the estimates across time for either group. This further motivates a metric like CV to quantify general variability across time for diagnostic purposes.

To assess the novelty of our metric, we compare it to two metrics that aim to be representative of a class of standard VF trend-based methods. Trend-based VF diagnostic techniques can be grouped in two categories, global and point-wise, and study associations between VF changes and progression (Vianna and Chauhan 2015). We define the global metric “Mean CV” as the CV of the VF wide mean of DLS at each visit. If the patient is stable, the means should be similar at each visit and the CV of the metric should be near zero. On the other hand, a progressing patient will have varying means at each visit and the CV should increase as visits accrue. We also compare our model to a commonly used point-wise linear regression (PLR) method. The progression metric is defined as the minimum p -value for the slope parameter from separately run simple linear regression analyses across all VF locations, where observed DLS is the dependent variable and time from first visit is the independent variable (Smith, Katz, and Quigley 1996). Using both “Mean CV” and PLR as comparators is far-reaching, as global methods are generally more robust, while point-wise methods identify local changes in visual ability.

To assess the diagnostic capability of our metric, we construct logistic regression models regressing various metrics on the clinical assessment of progression. We compare our method to “Mean CV,” PLR, and also the posterior mean CV of α_{tGH} from the Lee and Mitchell (2011) model, where α_{tGH} is estimated independently at each visit. We refer to this metric as “Space CV.” To obtain “Space CV,” we apply the Lee and Mitchell (2011) methodology with a Tobit likelihood. Thus, the appreciable differences between the two models that produce “Space CV” and “ST CV” are the definition of the weights, cross-covariance,

Table 1. Assessing the diagnostic capability of VF metrics.

Metrics	Estimate	SE	z-Value	Pr(> z)
Mean CV	0.40	0.16	2.55	0.011*
PLR	-0.59	0.25	-2.37	0.018*
Space CV	-0.07	0.17	-0.41	0.680
ST CV	0.39	0.16	2.44	0.015*

NOTE: Each metric is regressed against the clinical assessment of progression using a logistic regression model with the slope estimates being displayed (*significance level of 0.05). “Mean CV” is the CV of the VF wide mean of DLS at each visit, PLR is the minimum p -value of the location specific regression slopes, and “Space CV” and “ST CV” represent the mean posterior CV of α_{tGH} from the Lee and Mitchell (2011) and our spatiotemporal models, respectively. Each of the metrics are standardized.

and temporal correlation structure for the observational level parameters.

In Table 1, we present the results from the logistic regression analyses using a significance level of 0.05. Each predictor is standardized to facilitate comparisons of the different metrics. We can see that as expected “Mean CV” and PLR are significantly associated with glaucoma progression, with p -values of 0.011 and 0.018, respectively. For “Mean CV,” the estimated slope coefficient of 0.40 indicates that as a patient’s “Mean CV” increases, their risk of progression increases. For PLR, a smaller minimum p -value suggests an increased risk of progression (estimated slope coefficient of -0.59). Based on the glaucoma literature, we would expect trend-based methods such as “Mean CV” and PLR to have good discriminatory capability between stable and progressing eyes. It was less clear for our new metrics dependent on α_{tGH} . We see that “Space CV” is not significantly associated with glaucoma progression with a p -value of 0.680, while “ST CV” is significantly associated with a p -value of 0.015. This result is encouraging, yet surprising, since the models are similar and illuminates their differences in the VF data setting. However, for this newly defined metric to be impactful, we must verify that it is explaining a novel pathway in glaucoma progression, independent of existing metrics.

To assess whether “ST CV” provides novel diagnostic capabilities, we explore the correlation between the metrics. In Figure 1 of the online supplementary materials, we show pairwise correlation plots. We also present Pearson correlation estimates (ρ) and p -values from the hypothesis test: $H_0 : \rho = 0, H_1 : \rho \neq 0$. “ST CV” is uncorrelated with both “Mean CV” and PLR with estimated correlations of 0.06 and 0.11, respectively, and large p -values. This result has important implications; indicating that in addition to being highly predictive of progression, “ST CV” is uncorrelated with the standard trend-based metrics.

Table 2. Operating characteristics for diagnostic metrics.

	AIC	AUC	pAUC	p -Value ₁	p -Value ₂	p -Value ₃
Mean CV & PLR	209.81	0.68	0.21	—	—	—
Mean CV & PLR + Space CV	213.46	0.70	0.22	0.503	0.196	0.410
Mean CV & PLR + ST CV	204.46	0.74	0.29	0.010	0.042	0.059

NOTE: p -Values assess statistical significance of improvement over the combined trend-based metric (Mean CV & PLR) for, (1) additionally included regression parameters (using a nested likelihood ratio test), (2) AUC, and (3) pAUC.

This suggests that “ST CV” and the trend-based metrics can be used in conjunction to diagnose progression. Finally, the estimated correlation between “Space CV” and “ST CV” is 0.48 and the p -value is highly significant (<0.001), which is not surprising since “Space CV” and “ST CV” are estimating the same quantity. It is, however, interesting that these two metrics have such different associations with glaucoma progression (see Table 1). When “Space CV” was calculated with continuous weights this association did not change (p -value: 0.326, not included in Table 1), indicating the importance of temporal correlation and cross-covariance for properly modeling VF data. Presumably due to our enhanced methodology, “ST CV” is more precisely estimating the CV of α_{IGH} by smoothing the α_{IGH} and eliminating temporal noise and cross-covariance dependencies. To formalize this hypothesis a simulation study is designed in Section 6.

5.5.2. Extension to the Clinical Setting

Having demonstrated the novelty of using variability in the boundary detection parameter as a progression diagnostic, we now establish its clinical utility. We combine the two trend-based metrics into a composite model that represents an ideal synthesis of global and local methods. This metric includes “Mean CV” and PLR and their interaction. We then compare the changes in operating characteristics after separately adding “Space CV” and “ST CV” to this composite model, including the main effect and pairwise interactions. These results are presented in Table 2, where summary statistics include the Akaike information criterion (AIC), area under the receiver operating characteristic (ROC) curve (AUC), and the partial AUC (pAUC). Here, pAUC is limited to the clinically significant region of specificity of 85–100% (Zhu et al. 2014).

The results indicate the clinical importance of “ST CV,” producing optimal values of AIC, AUC, and pAUC. Furthermore, adding “ST CV” produces statistically significant improvements over the composite trend-based model in terms of the additional predictors, and both AUC and pAUC (although at the $\alpha = 0.10$ level for pAUC). Meanwhile, there are no improvements in the operating characteristics when adding “Space CV,” which is also confirmed when investigating the ROC curves (Figure 4a). In Figure 4a, it is clear that the inclusion of “ST CV” increases the discriminatory ability of the composite trend-based model, in particular within the clinically significant region, left of the dashed line.

In addition to operating characteristics calculated at the end of the study, we are also interested in the diagnostic performance of each metric in earlier stages of the disease, where accurate discrimination is more important for preserving a patient’s visual ability. In Figure 4(b), a smoothed LOESS trajectory of pAUC is

presented for each of the models in Table 2 for the first 4.5 years of follow-up. This analysis reflects a true clinical setting where data are analyzed as they are collected for a patient over time. We calculate “Mean CV,” PLR, “Space CV,” and “ST CV” for each patient during each visit. Then, using the fitted logistic regression models from the end of the study, we calculate predicted probabilities of VF progression. Finally, these probabilities are converted into binary disease progression diagnoses based on probability cutoffs obtained from the full study data. We forgo the commonly used Youden’s index (which maximizes sensitivity and specificity), for a criteria that forces specificity to be in the clinically significant range, while maximizing sensitivity.

These results indicate that the improvement in operating characteristics attributed to “ST CV” begins early during the follow-up period, in particular for pAUC. This result establishes that “ST CV” is clinically useful, and further enforces the importance of the introduced methodology as “Space CV” is incapable of improving operating characteristics during the same time period. In addition, Figure 4(c) presents boxplots of the predicted probabilities at half-year intervals, by progression status as defined at the end of the study. In Figure 4(c), there is a clear separation between the progressing and non progressing patients’ predicted probabilities of progression beginning around 1.5 years after baseline visit. This separation is more pronounced with the addition of “ST CV.”

5.6. Sensitivity Analyses

In the online supplementary materials, we present a number of sensitivity analyses and an additional simulation study to explore various modeling assumptions, including hyperparameter choices, the correlation structure, the bounds of ϕ , and misspecification of ρ and the dissimilarity metric. Overall, we find that the results are robust to these assumptions. For full details on these analyses, please see the online supplementary materials.

6. Model Performance Simulation

A simulation study is designed to assess the performance of the proposed model in the presence of temporal dependence and cross-covariance. We focus on the estimation of CV of α_{IGH} , since we propose using this posterior distribution in diagnosing progression.

To understand the gains of using our methodology in the presence of temporal correlation and cross-covariance dependence, we design a simulation study comparing the spatial model of Lee and Mitchell (2011), referred to as Space, and our spatiotemporal method (ST). We simulate data based on a set of known hyperparameters and then estimate the posterior mean of the CV of α_{IGH} using both models, comparing the estimates to the known truth using bias, mean squared error (MSE), and empirical coverage (EC). The simulation is designed to explore model performance in a typical patient from our study data. As such, we fix the true hyperparameters in the simulation study to the posterior means obtained from Section 5 for an average patient. Here an average patient is one whose posterior mean estimates are average among all patients. The

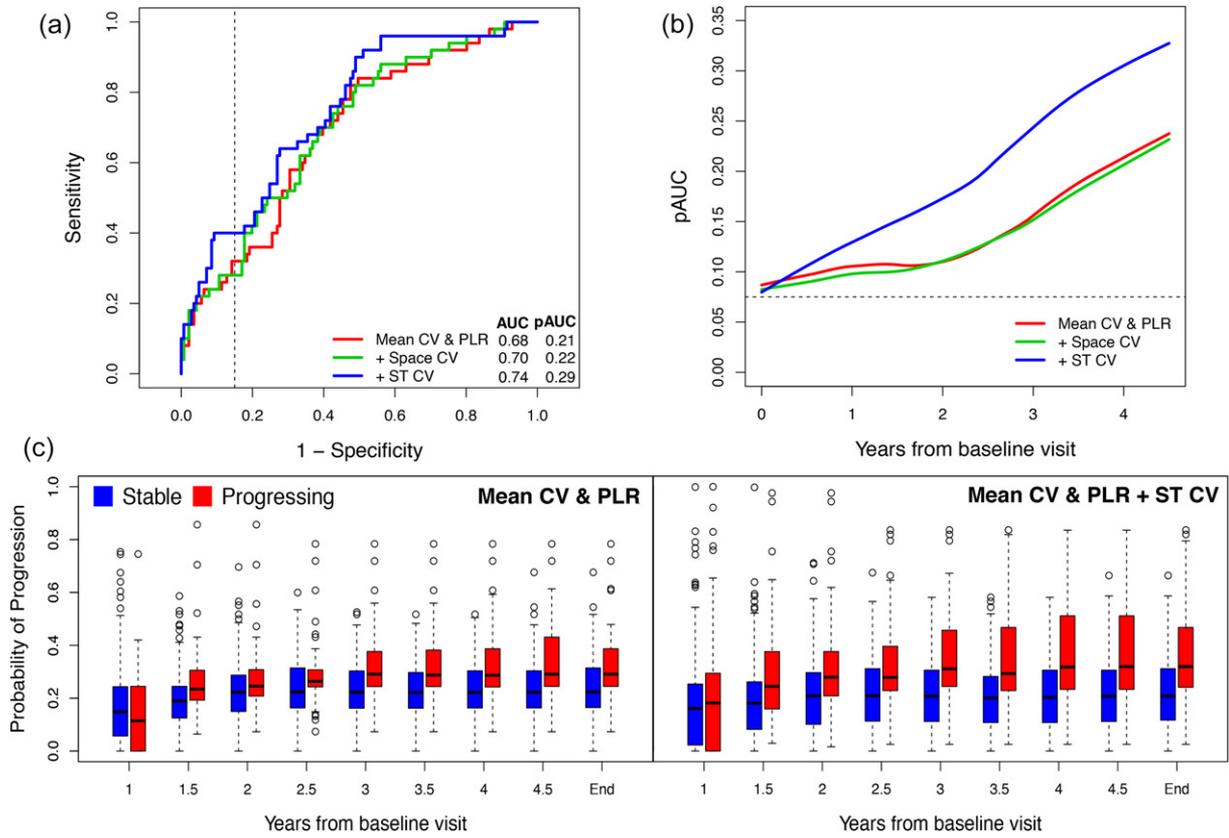


Figure 4. Demonstrating the performance of diagnostic metrics: (a) ROC curves for the combined trend-based metric (Mean CV and PLR), including the addition of “Space CV” and “ST CV” (clinically meaningful region to the left of dashed line), (b) pAUC in the initial years from baseline visit (estimates are presented as smooth LOESS curves and the dashed line indicates no discrimination), (c) boxplots of the predicted probabilities of progression presented by disease status over time as defined at the end of the study period.

true hyperparameters used in the simulation study are as follows,

$$\delta = \begin{bmatrix} 2.446 \\ 0.070 \\ 0.974 \end{bmatrix}, \quad \mathbf{T} = \begin{bmatrix} 0.820 & 0.004 & -0.028 \\ 0.004 & 0.380 & -0.191 \\ -0.028 & -0.191 & 0.840 \end{bmatrix},$$

$\phi = 0.163.$

Simulation settings are developed to incrementally understand the impact of the cross-covariance, \mathbf{T} , and temporal correlation, $\Sigma(\phi)$. To facilitate this analysis, Equation (5) is used. Define $\mathbf{T}_{\text{diag}} = \text{diag}(\mathbf{T})$, such that \mathbf{T}_{diag} has zeros on the off-diagonal and $\phi_I = 100$ (note that a large ϕ for the exponential correlation implies temporal independence). The simulation settings are as follows with the data generating covariance given in parentheses, A: no temporal correlation, no cross-covariance ($\Sigma(\phi_I) \otimes \mathbf{T}_{\text{diag}}$), B: no temporal correlation, cross-covariance ($\Sigma(\phi_I) \otimes \mathbf{T}$), C: temporal correlation, no cross-covariance ($\Sigma(\phi) \otimes \mathbf{T}_{\text{diag}}$), D: temporal correlation, cross-covariance ($\Sigma(\phi) \otimes \mathbf{T}$). Finally, we must specify the days and number of VF visits. To obtain the visit days, we sampled from a Poisson distribution with rate parameter equal to the average difference in days between VF visits (rate = 117.25 days). We present the simulation at the median (7) and maximum (21) number of visits in our study data.

For each simulation setting, we generate 100 values of θ and then use each θ to simulate 10 datasets. This yields 1000 simulated datasets for each of the simulation settings. In Table 3,

Table 3. Results from simulation study estimating the posterior mean of the CV of α_{tGH} in setting A: no temporal correlation, no cross-covariance, B: no temporal correlation, cross-covariance, C: temporal correlation, no cross-covariance, and D: temporal correlation, cross-covariance.

Setting	Model	# Visits					
		7 (Median)			21 (Maximum)		
		Bias	MSE	EC	Bias	MSE	EC
A	ST	0.032	0.107	0.97	0.023	0.084	0.95
	Space	0.102	0.111	0.87	0.174	0.098	0.81
B	ST	0.047	0.125	0.96	0.034	0.101	0.95
	Space	0.119	0.120	0.80	0.250	0.176	0.68
C	ST	-0.113	0.088	0.97	0.002	0.037	0.98
	Space	-0.306	0.172	0.51	-0.182	0.088	0.60
D	ST	-0.103	0.085	0.98	0.015	0.060	0.98
	Space	-0.299	0.172	0.51	-0.169	0.087	0.59

NOTE: Each setting is also implemented for the median (7) and maximum (21) number of VF visits. Each reported estimate is based on 1000 simulated datasets.

the bias, MSE, and EC are presented for the two models across all settings and at the median and maximum number of visits. We begin by noting the average (across all settings) simulation standard errors of the bias (0.223, 0.063), MSE (0.039, 0.028), and EC (0.141, 0.013) for the space and spatiotemporal models, respectively. In general the spatiotemporal model has smaller standard errors than the space model, indicating less variability in estimation across the simulated datasets.

The bias results suggest that our spatiotemporal model produces an estimator of CV with bias closer to zero than the spatial

model on average across all settings. In the settings with no temporal correlation (i.e., A and B) the MSE is quite similar between the models, except for Setting B where the maximum number of visits is considered. In this setting, the MSE is lower for the spatiotemporal model, revealing the importance of the cross-covariance for properly estimating the posterior mean of the CV of α_{tGH} . In all other settings, the estimated MSEs for the spatiotemporal model are smaller than for the spatial model, with the largest differences seen in Settings C and D for the median number of visits.

The EC is the proportion of the time that the estimated Bayesian credible intervals contain the true CV value. We define the nominal coverage as 95%. The EC results favor the spatiotemporal model across all settings. The spatial model has an EC of 0.87 and 0.81 in setting A for the median and maximum settings, respectively. This is the most ideal setting for the spatial model, because it was the setting in which the model was introduced (although within the context of disease mapping). The EC deteriorates in the spatial model as cross-covariance and temporal structure are introduced in the simulated data, falling as low as 0.51, while the spatiotemporal model performs consistently.

7. Discussion

In this article, we proposed a modeling framework for incorporating local neighborhood structure into complex spatiotemporal areal data. Based on this framework, we developed an innovative and highly predictive diagnostic of glaucoma progression that outperformed the spatial-only model. 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, including disease mapping and more generally GLMM. The methodology is built upon theory in boundary detection literature, using a Bayesian hierarchical modeling framework for inference. We extended the spatial-only method introduced by Lee and Mitchell (2011), that elegantly introduced a dissimilarity metric in a parsimonious framework. Our method allows for the local neighborhood structure to adapt over time as a function of changing dissimilarity metric parameters. Furthermore, the temporal correlation and cross-covariance are accounted for, eliminating known sources of excess variability. This parsimonious method induces nonstandard local spatial surfaces in areal data that are capable of mirroring complex processes, such as the surface of the VF.

We have shown (Figure 2) that the spatiotemporal covariance structure specified in our methodology successfully induces local neighborhood structure across the VF. The novelty of employing the Garway-Heath angles in the form of a dissimilarity metric provides a connection between the VF and the underlying optic disc, resulting in a neighborhood structure on the VF that is representative of the RNFL. Other methods incorporated the Garway-Heath angles into statistical models, either collapsed into regions of anatomical interest (Betz-Stablein et al. 2013; Warren et al. 2016) or by VF location (Zhu et al. 2014), but none allowed the effect to change dynamically. The dissimilarity metric parameter, α_{tGH} , dictates the VF spatial surface at each

visit, allowing the neighborhood structure to adapt alongside changes in DLS.

The results from applying our method to VF data (Section 5) demonstrated the added benefit of using our methodology, in both the clinical and statistical frameworks. We defined the diagnostic metric “ST CV” as the posterior mean of the CV of α_{tGH} , and showed that it is a significant predictor of clinically determined glaucoma progression while being uncorrelated with standard VF trend-based metrics (“Mean CV,” PLR). Since “ST CV” is independent of “Mean CV” and PLR, and each one is an effective predictor, there is clinical utility to combining their diagnostic capability (i.e., “ST CV” is not meant to compete with “Mean CV” and PLR but should be used together). The “Space CV” metric was not significantly associated with progression.

The significant association observed between “ST CV” and glaucoma progression was statistically note worthy, but taken in isolation had limited clinical implications. We showed the clinical utility of “ST CV” in Section 5.5.2, where the operating characteristics of the trend-based methods, both at the end of the study and in the early period, are improved with the addition of “ST CV” (Table 2, Figure 4). This fundamental finding shows that “ST CV” constitutes an alternate pathway for studying glaucoma progression using VF data. This pathway is facilitated by the dissimilarity metric framework introduced by the proposed methodology and consequently “ST CV” has a novel interpretation among glaucoma progression diagnostic metrics. In particular, “ST CV” quantifies damage to the optic disc over time as a function of stability on the spatial structure of the VF. It provides a method for discussing underlying damage to the optic disc using VF data.

Our simulation study indicated that the presence of temporal correlation and cross-covariance dependence impacted estimation of the posterior mean of the CV of α_{tGH} . The results (Table 3) illustrated superior bias, MSE, and EC for our proposed spatiotemporal model over the spatial model of Lee and Mitchell (2011). These results imply that our method provides a framework for more precisely estimating the CV of α_{tGH} by smoothing the α_{tGH} and eliminating temporal noise and cross-covariance dependencies. These results are consistent with studies that have shown when temporal correlation is ignored, estimators are often biased and variances are poorly estimated (West 1996).

In our implementation of the model for glaucoma data, we are not particularly interested in the specific location of the spatial boundaries across the VF, unlike traditional boundary detection applications. In isolation we do not believe those findings would be predictive of glaucoma progression, but would only inform about the current level of vision loss for a patient. Our results suggest that changes in these boundaries across time represent an informative, innovative, and unique metric for diagnosing glaucoma progression. Using the dissimilarity metric gives us the ability to directly quantify the boundary changes across time and therefore, to make clinical progression determinations for a patient.

The need for spatiotemporal boundary detection with a dissimilarity metric is also driven by the unique features of our VF dataset. Using a more basic CAR structure with fixed neighborhood adjacencies would result in ignoring the fact that these

neighborhood definitions are potentially changing across time due to disease progression. It would also require us to fix the adjacencies a priori. On the other hand, using a point-referenced geostatistical model with a spatiotemporal covariance function would result in ignoring the inherent discreteness of the spatial domain, a grid overlaying the VF. Spatiotemporal boundary detection with a dissimilarity metric represents an ideal blend between the two methods by allowing for discreteness in the spatial domain and flexibility in defining neighbors through use of a continuous covariate such as the Garway-Heath angles.

Finally, this work opens up numerous avenues for future statistical research. If data are available, covariates can be incorporated in the observational model naturally via the link function, $g(\vartheta_{it}) = x_{it}\beta + \varphi_{it}$. Currently, the specification of the Leroux, Lei, and Breslow (2000) likelihood includes a μ_t in the mean structure, thus to identify the μ_t and β_0 we must either apply a constraint, $\sum_{t=0}^v \mu_t = 0$, or reparameterize the likelihood to be a function of only τ_t and α_t . The incorporation of covariates is of particular importance in disease mapping, where standard incidence rates are often mapped over varying risk factors (Lee and Mitchell 2013). Another natural extension to this approach includes generalizing the separable temporal covariance to allow unique temporal decay parameters.

Supplementary Material

The supplementary materials contain details for implementing the MCMC sampler, including derivation of full conditionals, along with prediction theory, additional sensitivity analyses, and figures.

Acknowledgments

The authors thank the Lions Eye Institute for Figure 1(b) and Wallace L. M. Alward (Ophthalmology and Visual Sciences, University of Iowa Carver College of Medicine) for Figure 1(c); and also 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.

Funding

This work was partially supported by the National Institute of Environmental Health Sciences (SIB; T32ES007018), the Research to Prevent Blindness (DO, UNC-CH), and the National Center for Advancing Translational Science (JLW; UL1 TR001863, KL2 TR001862).

References

- Banerjee, S., and Gelfand, A. E. (2006), "Bayesian Wombling: Curvilinear Gradient Assessment Under Spatial Process Models," *Journal of the American Statistical Association*, 101, 1487–1501. [1064]
- Banerjee, S., Gelfand, A. E., and Carlin, B. P. (2003), *Hierarchical Modeling and Analysis for Spatial Data*, Boca Raton, FL: CRC Press. [1064,1065,1066]
- Besag, J. (1974), "Spatial Interaction and the Statistical Analysis of Lattice Systems," *Journal of the Royal Statistical Society, Series B*, 36, 192–236. [1065]
- Betz-Stablein, B. D., Morgan, W. H., House, P. H., and Hazelton, M. L. (2013), "Spatial Modeling of Visual Field Data for Assessing Glaucoma Progression," *Investigative Ophthalmology & Visual Science*, 54, 1544–1553. [1068,1072]
- Bowman, F. D., Caffo, B., Bassett, S. S., and Kilts, C. (2008), "A Bayesian Hierarchical Framework for Spatial Modeling of fMRI Data," *NeuroImage*, 39, 146–156. [1063]
- Bryan, S. R., Eilers, P. H., Lesaffre, E. M., Lemij, H. G., and Vermeer, K. A. (2015), "Global Visit Effects in Point-Wise Longitudinal Modeling of Glaucomatous Visual Fields," *Investigative Ophthalmology & Visual Science*, 56, 4283–4289. [1068]
- Castruccio, S., Ombao, H., and Genton, M. G. (2016), "A Multi-resolution Spatio-Temporal Model for Brain Activation and Connectivity in fMRI Data," arXiv no. 1602.02435. [1063]
- Chauhan, B. C., Garway-Heath, D. F., GoñI, F. J., Rossetti, L., Bengtsson, B., Viswanathan, A. C., and Heijl, A. (2008), "Practical Recommendations for Measuring Rates of Visual Field Change in Glaucoma," *British Journal of Ophthalmology*, 92, 569–573. [1063,1064]
- Davson, H. (2012), *Physiology of the Eye*, Amsterdam: Elsevier. [1064]
- Ecker, M. D., and Gelfand, A. E. (2003), "Spatial Modeling and Prediction Under Stationary Non-geometric Range Anisotropy," *Environmental and Ecological Statistics*, 10, 165–178. [1063]
- Eddelbuettel, D., François, R., Allaire, J., Chambers, J., Bates, D., and Ushey, K. (2011), "Rcpp: Seamless R and C++ Integration," *Journal of Statistical Software*, 40, 1–18. [1066]
- Garway-Heath, D. F., Poinoosawmy, D., Fitzke, F. W., and Hitchings, R. A. (2000), "Mapping the Visual Field to the Optic Disc in Normal Tension Glaucoma Eyes," *Ophthalmology*, 107, 1809–1815. [1064]
- Gelfand, A. E., Diggle, P., Guttorp, P., and Fuentes, M. (2010), *Handbook of Spatial Statistics*, Boca Raton, FL: CRC Press. [1065]
- Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B. (2014), *Bayesian Data Analysis* (Vol. 2), Boca Raton, FL: Chapman & Hall/CRC. [1067]
- Gelman, A., and Shalizi, C. R. (2013), "Philosophy and the Practice of Bayesian Statistics," *British Journal of Mathematical and Statistical Psychology*, 66, 8–38. [1067]
- Geman, S., and Geman, D. (1984), "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* 6, 721–741. [1065]
- He, Y., Chen, Z., and Evans, A. (2008), "Structural Insights Into Aberrant Topological Patterns of Large-Scale Cortical Networks in Alzheimer's Disease," *Journal of Neuroscience*, 28, 4756–4766. [1063]
- Jacquez, G. M., and Greiling, D. A. (2003), "Geographic Boundaries in Breast, Lung and Colorectal Cancers in Relation to Exposure to Air Toxics in Long Island, New York," *International Journal of Health Geographics*, 2, 4. [1065]
- Konrad, K., and Eickhoff, S. B. (2010), "Is the ADHD Brain Wired Differently? A Review on Structural and Functional Connectivity in Attention Deficit Hyperactivity Disorder," *Human Brain Mapping*, 31, 904–916. [1063]
- Lee, D., and Mitchell, R. (2011), "Boundary Detection in Disease Mapping Studies," *Biostatistics*, 15, 457–469. [1066,1067,1069,1070,1072]
- (2013), "Locally Adaptive Spatial Smoothing Using Conditional Auto-regressive Models," *Journal of the Royal Statistical Society, Series C*, 62, 593–608. [1073]
- (2014), "Controlling for Localised Spatio-Temporal Autocorrelation in Long-Term Air Pollution and Health Studies," *Statistical Methods in Medical Research*, 23, 488–506. [1066]
- Lee, D., Rushworth, A., and Sahu, S. K. (2014), "A Bayesian Localized Conditional Autoregressive Model for Estimating the Health Effects of Air Pollution," *Biometrics*, 70, 419–429. [1066]
- Leroux, B. G., Lei, X., and 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*, eds. M. E. Halloran and D. Berry, New York, NY: Springer, pp. 179–191. [1065,1073]
- Li, P., Banerjee, S., Hanson, T. A., and McBean, A. M. (2015), "Bayesian Models for Detecting Difference Boundaries in Areal Data," *Statistica Sinica*, 25, 385–402. [1066]
- Li, P., Banerjee, S., and McBean, A. M. (2011), "Mining Boundary Effects in Areal Referenced Spatial Data Using the Bayesian Information Criterion," *Geoinformatica*, 15, 435–454. [1066]

- Lu, H., and Carlin, B. P. (2005), "Bayesian Areal Wombling for Geographical Boundary Analysis," *Geographical Analysis*, 37, 265–285. [1065]
- Lu, H., Reilly, C. S., Banerjee, S., and Carlin, B. P. (2007), "Bayesian Areal Wombling via Adjacency Modeling," *Environmental and Ecological Statistics*, 14, 433–452. [1066]
- Ma, H., and Carlin, B. P. (2007), "Bayesian Multivariate Areal Wombling for Multiple Disease Boundary Analysis," *Bayesian Analysis*, 2, 281–302. [1064,1066]
- Quigley, H. A., Katz, J., Derick, R. J., Gilbert, D., and Sommer, A. (1992), "An Evaluation of Optic Disc and Nerve Fiber Layer Examinations in Monitoring Progression of Early Glaucoma Damage," *Ophthalmology*, 99, 19–28. [1063]
- R Core Team (2016), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing. [1066]
- Rushworth, A., Lee, D., and Sarran, C. (2017), "An Adaptive Spatiotemporal Smoothing Model for Estimating Trends and Step Changes in Disease Risk," *Journal of the Royal Statistical Society, Series C*, 66, 141–157. [1066]
- Sampson, P. D. (2010), "Constructions for Nonstationary Spatial Processes," in *The Handbook of Spatial Statistics*, eds. A. E. Gelfand, P. J. Diggle, M. Fuentes, and P. Guttorp, Boca Raton, FL: CRC Press, pp. 119–130, chapter 9. [1063]
- Smith, S. D., Katz, J., and Quigley, H. A. (1996), "Analysis of Progressive Change in Automated Visual Fields in Glaucoma," *Investigative Ophthalmology & Visual Science*, 37, 1419–1428. [1069]
- Smythies, J. (1996), "A Note on the Concept of the Visual Field in Neurology, Psychology, and Visual Neuroscience," *Perception*, 25, 369–372. [1064]
- Tanner, M. (1996), *Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions*, New York: Springer-Verlag. [1067]
- Tham, Y.-C., Li, X., Wong, T. Y., Quigley, H. A., Aung, T., and Cheng, C.-Y. (2014), "Global Prevalence of Glaucoma and Projections of Glaucoma Burden Through 2040: A Systematic Review and Meta-Analysis," *Ophthalmology*, 121, 2081–2090. [1063]
- Thompson, P. M., Hayashi, K. M., Sowell, E. R., Gogtay, N., Giedd, J. N., Rapoport, J. L., de Zubicaray, G. I., Janke, A. L., Rose, S. E., Semple, J., and Doodrell, D. M. (2004), "Mapping Cortical Change in Alzheimer's Disease, Brain Development, and Schizophrenia," *NeuroImage*, 23, S2–S18. [1063]
- Tobin, J. (1958), "Estimation of Relationships for Limited Dependent Variables," *Econometrica*, 26, 24–36. [1068]
- Vianna, J. R., and Chauhan, B. C. (2015), "How to Detect Progression in Glaucoma," *Progress in Brain Research*, 221, 135–158. [1063,1069]
- Warren, J. L., Mwanza, J.-C., Tanna, A. P., and Budenz, D. L. (2016), "A Statistical Model to Analyze Clinician Expert Consensus on Glaucoma Progression Using Spatially Correlated Visual Field Data," *Translational Vision Science & Technology*, 5, 14–14. [1068,1072]
- West, M. (1996), *Bayesian Forecasting*, Wiley Online Library. [1072]
- Womble, W. H. (1951), "Differential Systematics," *Science*, 114, 315–322. [1065]
- Zhu, H., Russell, R. A., Saunders, L. J., Ceccon, S., Garway-Heath, D. F., and Crabb, D. P. (2014), "Detecting Changes in Retinal Function: Analysis With Non-stationary Weibull Error Regression and Spatial Enhancement (ANSWERS)," *PLoS One*, 9, 1–11. [1070,1072]