Published in final edited form as:

Stat Med. 2008 July 30; 27(17): 3247–3268. doi:10.1002/sim.3265.

Bayesian Semiparametric Regression for Longitudinal Binary Processes with Missing Data

Li Su* and

Medical Research Council, Biostatistics Unit, Robinson Way, Cambridge CB2 0SR, U.K

Joseph W. Hogan

Center for Statistical Sciences, Department of Community Health, Box G-S121-7, Brown University, Providence, Rhode Island 02912, U.S.A

Summary

Longitudinal studies with binary repeated measures are widespread in biomedical research. Marginal regression approaches for balanced binary data are well developed, while for binary process data, where measurement times are irregular and may differ by individuals, likelihood-based methods for marginal regression analysis are less well developed. In this article, we develop a Bayesian regression model for analyzing longitudinal binary process data, with emphasis on dealing with missingness. We focus on the settings where data are missing at random, which require a correctly specified joint distribution for the repeated measures in order to draw valid likelihood-based inference about the marginal mean. To provide maximum flexibility, the proposed model specifies both the marginal mean and serial dependence structures using nonparametric smooth functions. Serial dependence is allowed to depend on the time lag between adjacent outcomes as well as other relevant covariates. Inference is fully Bayesian. Using simulations, we show that adequate modeling of the serial dependence structure is necessary for valid inference of the marginal mean when the binary process data are missing at random. Longitudinal viral load data from the HIV Epidemiology Research Study (HERS) are analyzed for illustration.

Keywords

Repeated measures; Marginal model; Nonparametric regression; Penalized splines; HIV/AIDS; Antiviral treatment

1 Introduction

1.1 Marginal modeling for longitudinal binary data

Longitudinal binary data arise in many fields of biomedical research, including cohort studies, intervention studies, and clinical trials. For example, in studies of antiviral therapy, detectability of plasma HIV viral load is an important endpoint for monitoring HIV progression and assessing treatment efficacy. For analyzing this type of data, multilevel modeling, such as generalized linear mixed models ([1]), is an effective approach that characterizes both mean and serial dependence structures flexibly. However, these models traditionally focus on conditional mean regression parameters, while marginal mean parameters might be of primary interest in some studies. In this article, we introduce a likelihood-based semiparametric regression model for the settings where marginal mean

*Email: li.su@mrc-bsu.cam.ac.uk; phone: 44-1223-330386; fax: 44-1223-330388.

structure is the target of inference and measurement times may be unbalanced or irregularly spaced due to missingness.

In terms of marginal models for longitudinal binary data, parametric regression using generalized estimating equations (GEE) ([2]) is a popular choice (the terms ‘parametric’ and ‘semi-parametric’ in this article refer to the regression function rather than the error distribution). With complete and balanced data that have a common set of equally spaced measurement times, it is well understood that GEE estimators of regression coefficients are consistent when the marginal mean model is correctly specified even if the dependence structure is misspecified ([3]). Further, to allow for more flexible dependence of the outcome variable on covariates, semiparametric moment-based regression has attracted substantial interests recently ([4,5,6,7,8,9,10]); estimating equations have been developed using conventional or modified profile-kernel methods as the extensions of GEE ([8,10]).

However, for longitudinal binary process data where measurement times are random and misaligned across individuals, estimation based solely on the first moment assumption (e.g., GEE) can break down. This is because the robust variance estimator in GEE uses replications across individuals to estimate the dependence structure nonparametrically, and this becomes problematical when there are many missing values and/or when measurement times are not common to all individuals ([3, page 80]). This problem can be overcome with likelihood methods, but unlike linear mixed models for continuous longitudinal data, models for the joint distribution of binary outcomes are less obvious. The canonical log-linear models and their extensions are computationally restricted to small and balanced binary series ([11]). Recently, Heagerty (2002) ([12]) introduced marginalized transition models (MTM) as an extension of the first-order MTM by Azzalini (1994) ([13]) to permit computationally feasible likelihood-based marginal regression analysis of balanced binary data in long series. In an MTM, the marginal mean and serial dependence within individuals are modeled separately and the serial dependence is captured using a Markov structure. In fact, the first-order MTM, or MTM(1), is a reparameterization of the canonical log-linear model when both the marginal mean and serial dependence are specified using logistic link functions. Nevertheless, for binary process data where measurements are unequally spaced and vary across individuals, there are relatively few methods available for marginal regression analysis in either moment-based or likelihood-based frameworks.

1.2 Missing data and likelihood methods

The motivation of our model originates from the consideration of missing data problem that can result in unequally spaced measurement times for longitudinal binary series. That is, our focus is on the binary processes with missing data.

Missing data problem is ubiquitous in longitudinal studies with long term follow-up of cohorts. With incomplete data, likelihood-based methods are attractive because a correctly specified likelihood model can be used to provide valid inference under ignorable missingness ([14]). In the case of longitudinal data, this requires that the joint distribution of outcomes must be properly modeled, including the dependence structure. For complete data or data missing completely at random (MCAR), parameters in the marginal mean and serial dependence structures of the original MTM(1) are orthogonal ([12,15]). However, this property does not hold under missing at random (MAR) condition. It follows that under MAR, correctly modeling dependence becomes necessary for valid inference about the marginal mean. Kurland and Heagerty (2004) ([16]) explored the misspecification bias for MTM under different missing data mechanisms, and showed that even under MAR, biased regression coefficient estimates can result from a model with correctly specified marginal mean structure if the serial dependence model is incorrect. Improving the fit of serial

dependence model reduces bias noticeably, which motivates our semiparametric regression approach for binary processes with missing data.

In this article we use the principle of the MTM for balanced binary data to develop likelihood-based semiparametric regression models for unbalanced data generated by longitudinal binary processes, focusing on dealing with missing data. We build on the MTM structure to incorporate nonparametric smooth functions in both the marginal mean and serial dependence models, allowing more flexible specification of both marginal covariate effects and dependence structure using semiparametric formulations. After including other covariates, serial dependence is modeled as an unspecified smooth function of the time lag between adjacent outcomes, providing robustness to possible misspecification with parametric functions. Our serial dependence model is similar to the *lorelogram* ([17]), an exploratory graphical display for dependence in longitudinal categorical data. To obtain a lorelogram, moment-based nonparametric regression is used to estimate marginal pairwise log-odds ratios, while data-driven methods for selecting smoothing parameters are not considered.

1.3 Motivating example

The HIV Epidemiology Research Study (HERS) is a multi-site longitudinal cohort study evaluating the natural history of HIV infection in women between 1993 and 2001 ([18]). At baseline the study enrolled 871 HIV-seropositive women and 439 HIV-seronegative women at risk for HIV infection. Participants were scheduled for 6-year follow-up, where a variety of clinical, behavioral and sociologic outcomes were recorded every 6 months.

The progression of HIV can be significantly altered by the suppression of plasma HIV RNA with highly active antiretroviral therapy (HAART), decreasing HIV-associated morbidity and mortality. Coinfection with Hepatitis C virus (HCV) is associated with HIV disease progression and HAART treatment efficacy ([19]). Using the HERS data, we aim to characterize the prevalence profile of detectable HIV viral load before and after HAART initiation and investigate whether there is an HCV coinfection effect associated with this time course, after accounting for the baseline (at enrollment) CD4 count and CD4 count changes from baseline.

The analysis focuses on whether viral load is detectable, i.e., whether plasma HIV RNA > 50 copies/mL. In clinical practice, the virus is considered suppressed when viral load is below this lower limit of detection. HAART is defined as a 3-drug regimen recommended by IAS-USA or PHS panels ([20]). We are able to use the HERS cohort to capture population-level effect of HAART because all initiations of HAART in this study were incident. Our analysis uses data from 378 participants with HAART use and baseline HCV status information available. HAART use, viral load and CD4 count were scheduled to be recorded every 6 months during the study. At each visit, participants reported only whether they received HAART in the past six months; exact dates for HAART initiation are not available. Therefore, in our analysis the time since HAART initiation is defined as the time since the earliest visit with HAART information reported. Because visit dates for the first reported HAART use and for viral load and CD4 measurements varied across individuals due to missed visits and/or assessment at non-scheduled time points, the measurement times correspond to dates and the unit of days is used. Figure 1 shows nine individual viral load profiles in the HERS. Exploratory analysis by moment-based smoothing splines in Figure 2 indicates substantial non-linearity in the population viral prevalence profile (without adjusting HCV status, baseline and longitudinal CD4 counts), which motivates a semiparametric regression approach with a nonparametric time effect for the marginal mean.

In analyzing these data, we pay attention to the missingness as showed by the differential follow-up times across individuals. Figure 3 shows the proportion of observations with viral load > 50 copies/mL (within individuals) as a function of follow-up times. It is apparent that individuals with longer follow-up times tended to have more observations with viral suppression. Thus MAR is probably more defensible than MCAR for missing data in our example, and carefully modeling the serial dependence structure is important in this case. In Section 6, we will apply our methods to the HERS data for allowing more flexible modeling of the binary processes, aiming to reduce any possible bias in the marginal mean estimation.

The remainder of this article is organized as follows. In the next section we provide some necessary background information on penalized splines for semiparametric regression. The proposed Bayesian regression model is introduced in Section 3. Prior specification and posterior sampling are detailed in Section 4. Section 5 presents a simulation study to evaluate the performance of the proposed model under MAR. Our methods are implemented for the HERS viral load data in Section 6. We conclude in Section 7.

2 Penalized splines for semiparametric regression

This section gives a brief overview of penalized splines for semiparametric regression. For more detailed coverage, see [21], [22] and [23].

Smoothing splines are a highly flexible method for semiparametric regression of longitudinal data ([8,9,10,24,25,26,27,28]). However, when there are a large number of unique values for a covariate, smoothing splines become less practical because the number of knots is equal to the number of unique values. A more general approach is penalized splines, or *P-splines* ([21]), which use relatively large number of knots, but many fewer than smoothing splines. A roughness penalty is also chosen to avoid overfitting. Ruppert (2002) ([29]) discussed the choice for the number of knots and found that the exact value of the number of knots is not important if it exceeds a certain minimum value; five to twenty knots in general are sufficient for the types of regression functions found in practice.

Convenient classes of P-splines include penalized B-splines of Eilers and Marx (1996) ([21]) and truncated polynomial splines of Ruppert and Carroll (2000) ([22]). The latter are the q th-degree polynomial splines with K fixed knots, $v_1, \dots, v_K$. For example, in a normal-error semiparametric regression model for cross-sectional outcomes Y_i ($i = 1, \dots, N$), we assume

$$Y_i = \mathbf{X}_i \boldsymbol{\beta} + m(V_i) + \varepsilon_i, \quad (1)$$

where $\varepsilon_i \sim N(0, \sigma_\varepsilon^2)$, $\boldsymbol{\beta}$ is the vector characterizing the parametric part of the mean, and $m(\cdot)$ is an unknown smooth function of a single covariate V . Following Ruppert and Carroll (2000) ([22]), $m(V)$ can be modeled by

$$\mathbf{B}(V)^T \boldsymbol{\theta} = \theta_0 + \theta_1 V + \dots + \theta_q V^q + \sum_{k=1}^K \theta_{q+k} (V - v_k)_+^q, \quad (2)$$

where $q \geq 1$ is an integer, $(a)_+^q = a^q \mathbf{I}(a \geq 0)$, $\boldsymbol{\theta} = (\theta_0, \dots, \theta_{q+K})^T$ is the vector of regression coefficients, and $\mathbf{B}(V) = (1, V, \dots, V^q, (V - v_1)_+^q, \dots, (V - v_K)_+^q)^T$ is the bases. Then $\theta_{q+1}, \dots,$

θ_{q+K} are jump sizes in the q th derivative of $\mathbf{B}(V)^T \boldsymbol{\theta}$ at the knots $v_1, \dots, v_K$. In practice the knots could be placed at the observed quantiles of V .

In terms of estimation and inference, one advantage of P-splines is that they can be incorporated into a likelihood and fitted using penalized likelihood methods. For example, assuming that the dispersion parameter σ_ε^2 in model (1) is known, its corresponding penalized log-likelihood takes the form

$$pl(\boldsymbol{\beta}, \boldsymbol{\theta}, \lambda_1) = l(\boldsymbol{\beta}, \boldsymbol{\theta}) - \lambda_1 \boldsymbol{\theta}^T \mathbf{D} \boldsymbol{\theta}, \quad (3)$$

where $l(\boldsymbol{\beta}, \boldsymbol{\theta})$ is the original normal log-likelihood, λ_1 is the smoothing parameter, and $\mathbf{D}$ is a known positive semi-definite penalty matrix. A typical choice of $\mathbf{D}$ is

$$\mathbf{D} = \begin{bmatrix} \mathbf{0}_{(q+1) \times (q+1)} & \mathbf{0}_{(q+1) \times K} \\ \mathbf{0}_{K \times (q+1)} & \Sigma_{K \times K}^{-1} \end{bmatrix}, \quad (4)$$

where Σ is a positive definite matrix. This choice of $\mathbf{D}$ penalizes $\theta_{q+1}, \dots, \theta_{q+K}$ only. A

useful special case sets $\Sigma = \mathbf{I}_K$, so that $\boldsymbol{\theta}^T \mathbf{D} \boldsymbol{\theta} = \sum_{k=1}^K \theta_{q+k}^2$ is the sum of the squared jump sizes in (2).

Penalized likelihood models such as in (3) can also be viewed as the particular cases of generalized linear mixed models (GLMM) ([21, 23, 25, 30]). If we divide $\boldsymbol{\theta}$ in (2) into two parts, $\boldsymbol{\theta}_1 = (\theta_0, \dots, \theta_q)^T$ and $\boldsymbol{\theta}_2 = (\theta_{q+1}, \dots, \theta_{q+K})^T$, and assign the elements of $\boldsymbol{\theta}_2$ a distribution, the penalty $\lambda_1 \boldsymbol{\theta}^T \mathbf{D} \boldsymbol{\theta}$ in (3) with the choice of $\mathbf{D}$ as in (4) is equivalent to the negative log-likelihood of $\boldsymbol{\theta}_2$ assuming $\boldsymbol{\theta}_2 \sim N(\mathbf{0}, \Sigma/(2\lambda_1))$. P-spline regression formulated as GLMM can be limited by the need to approximate high dimensional integrals and replacing true parameters with estimated parameters in confidence intervals. On the other hand, Bayesian P-splines enjoy both the flexibility of semiparametric models and the exact inference provided by the Bayesian inferential machinery, thanks to the recent advances in Markov chain Monte Carlo (MCMC) simulation techniques. MCMC techniques provide rich output for inference, prediction and model comparison; no approximations based on asymptotic arguments have to be made, and data-driven smoothing parameters are incorporated into the model, for example, by assigning a non-informative hyper-prior to λ_1 in (3) within a hierarchical model setting. Crainiceanu et al. (2005) ([31]) showed how to implement Bayesian analysis for a variety of P-splines regression models using WinBUGS package ([32]). In this article we focus on truncated polynomial bases in Ruppert and Carroll (2000) ([22]) and adopt a fully Bayesian approach for P-splines.

3 The semiparametric regression model for binary process data

3.1 Model formulation

Suppose that the data come from N individuals, and for the i th ($i = 1, \dots, N$) individual there is a binary outcome process $\{Y_i(t): t \geq 0\}$ with mean $\pi_i(t)$. At time points $t_{i1}, \dots, t_{in_i}$ the binary outcome measurements $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{in_i})^T$ are observed. We assume that associated exogenous but possibly time-varying covariates $(\mathbf{X}_{ij}, V_{ij})$ ($j = 1, \dots, n_i$) are recorded for each individual at each occasion, where $\mathbf{X}_{ij}$ is a $1 \times p$ vector and V_{ij} is a scalar. Our primary objective is to infer the marginal mean $\pi_{ij}^M = E(Y_{ij} | \mathbf{X}_{ij}, V_{ij})$. We assume that

$$g(\pi_{ij}^M) = \mathbf{X}_{ij}\beta + m(V_{ij}), \quad (5)$$

where $g(\cdot)$ is a link function, β is a $p \times 1$ vector, and $m(\cdot)$ is an unknown smooth function. Note that additive structure of smooth functions similar to generalized additive models ([33]) can also be introduced; without loss of generality we focus on the single nonparametric term in this article. It is assumed that the semiparametric regression model in (5) satisfies the full covariate conditional mean condition ([34]) such that $E(Y_{ij} | \mathbf{X}_{i1}, \dots, \mathbf{X}_{in_i}, V_{i1}, \dots, V_{in_i}) = E(Y_{ij} | \mathbf{X}_{ij}, V_{ij})$. Generally, this condition assumes that stochastic time-varying covariates are properly modeled through $(\mathbf{X}_{ij}, V_{ij})$, possibly by including appropriate lagged values; it also assumes that Y_{ij} , the outcome at time t_{ij} , is not a predictor of future covariates $(\mathbf{X}_{ij+l}, V_{ij+l})$ for $l \geq 1$ ([3]).

The marginal mean structure in (5) allows flexible dependence of the average response on covariates by introducing a semiparametric formulation. However, it only defines one facet of the multivariate probability structure of $\mathbf{Y}_i$; additional assumptions regarding the serial dependence structure are required to identify a joint probability model for likelihood-based analysis. Following Heagerty (2002) ([12]), we use a first-order Markov model to describe the dependence structure. Specifically, we assume that Y_{ij} is dependent on its history only through the immediate previous observed outcome $Y_{i,j-1}$, such that $E(Y_{ij} | \mathbf{X}_{ij}, V_{ij}, Y_{i1}, \dots, Y_{i,j-1}) = E(Y_{ij} | \mathbf{X}_{ij}, V_{ij}, Y_{i,j-1})$. Define $\pi_{ij}^C = E(Y_{ij} | \mathbf{X}_{ij}, V_{ij}, Y_{i,j-1})$. The dependence structure follows a first-order transition model

$$\text{logit}(\pi_{ij}^C) = \Delta_{ij}(\mathbf{X}_{ij}, V_{ij}) + \gamma_{ij}(\mathbf{X}_{ij}, V_{ij}) \cdot Y_{i,j-1}. \quad (6)$$

Note that other link functions can also be used; we use logit link in this article for illustration purpose. The log odds ratio γ_{ij} measures the association between $Y_{i,j-1}$ and Y_{ij} ; the intercept Δ_{ij} describes the conditional mean of Y_{ij} when $Y_{i,j-1} = 0$, and its value is determined such that the marginal mean structure in (5) and the Markov dependence structure in (6) are simultaneously satisfied. In fact, given β , $m(\cdot)$ and γ_{ij} , the value of Δ_{ij} can be derived in closed form ([13]) through the relationship between marginal and conditional means:

$$\begin{aligned} \pi_{ij}^M &= \sum_{y=0}^1 \text{pr}(Y_{ij}=1 | \mathbf{X}_{ij}, V_{ij}, Y_{i,j-1}=y) \cdot \text{pr}(Y_{i,j-1}=y | \mathbf{X}_{i,j-1}, V_{i,j-1}) \\ &= \sum_{y=0}^1 \pi_{ij}^C(y) \cdot \pi_{i,j-1}^M(y). \end{aligned}$$

Further details are given in Heagerty (2002) ([12]).

To accommodate the situations where individuals might not be observed at a common set of equally spaced time points and measurement times can be continuous and misaligned across individuals, we allow the first-order serial dependence in (6) to vary nonparametrically as a function of the time lag between adjacent outcomes. This is accomplished by modeling the serial dependence γ_{ij} using a semiparametric regression model,

$$\begin{aligned} \gamma_{ij} &= \mathbf{Z}_{ij}\alpha + h(t_{ij} - t_{i,j-1}) \\ &= \mathbf{Z}_{ij}\alpha + h(u_{ij}), \end{aligned} \quad (7)$$

where $\mathbf{Z}_{ij}$ in general can be a subset of the covariates $(\mathbf{X}_{ij}, V_{ij})$, $u_{ij} = t_{ij} - t_{i,j-1}$ ($j = 2, \dots, n_i$) is the time elapsed between adjacent outcome measurements, and $h(\cdot)$ is an unknown smooth function. The parameter vector α and smooth function $h(\cdot)$ determine how the dependence of Y_{ij} on $Y_{i,j-1}$ varies as a function of the covariates $\mathbf{Z}_{ij}$ and the time lag variable u_{ij} . For example, if $\gamma_{ij} = \alpha_0 + \alpha_1 \cdot Z_i$ (e.g., Z_i is a treatment group indicator), individuals for whom $Z_i = 1$ are allowed to have different serial dependence compared with individuals for whom $Z_i = 0$; and here serial dependence is assumed to be independent of the measurement time and the time lag between adjacent outcomes. If $\gamma_{ij} = h(u_{ij})$, we allow the first-order serial dependence to depend on the time lag nonparametrically, and it follows that the specification and estimation of the time lag effect can be made more flexible than in a parametric serial dependence model. For the setting with common and equally spaced measurement times, $h(\cdot)$ reduces to a constant. The nonparametric formulation here is not restricted to the time lag effect; it can be applied to other important covariates, such as the time effect when the dependence structure is not assumed to be stationary over time.

P-splines introduced in Section 2 are used to model the two smooth functions $m(\cdot)$ and $h(\cdot)$, with smoothing parameters λ_1 and λ_2 respectively. The covariate effect $m(V_{ij})$ in the marginal mean model (5) is specified using truncated polynomial bases as in (2). Similarly, for time lag effect $h(u_{ij})$ in the serial dependence model (7), we have

$$h(u) = \varphi_0 + \varphi_1 u + \dots + \varphi_s u^s + \sum_{k=1}^{K^*} \varphi_{s+k} (u - \xi_k)_+^s, \quad (8)$$

where $s \geq 1$ is an integer, K^* is the total number of knots for u_{ij} , $\boldsymbol{\varphi} = (\varphi_0, \dots, \varphi_{s+K^*})^T$ is a vector of regression coefficients, and $\xi_1 < \dots < \xi_{K^*}$ are fixed knots such that ξ_k ($k = 1, \dots, K^*$) is the sample quantile corresponding to probability $k/(K^* + 1)$ for pooled u_{ij} s of all individuals across time.

3.2 Implications for situations with missing data

From a Bayesian perspective, likelihood-based methods can provide valid inference when missingness is MAR and the priors of parameters for describing the measurement process and the missingness process are independent ([14]). However, this requires correct specification of the entire joint distribution of the repeated measures, and therefore correct specification of the dependence structure plays a critical role in valid inference about the marginal mean as discussed in Section 1.2.

To see this in more details, consider a simple case of correlated bivariate binary outcomes $\{(Y_{i1}, Y_{i2}), i = 1, \dots, N\}$. Suppose we have complete observations for individuals $i = 1, \dots, r$ (completers), but the values of Y_{i2} are missing for individuals $i = (r + 1), \dots, N$ (dropouts). A missingness indicator R_i is defined as $R_i = 1$ if Y_{i2} is observed and $R_i = 0$ if Y_{i2} is missing. The observed data can be tabulated in 2×2 contingency tables as in Table 1.

We define the marginal means such that $E(Y_{ij}) = \pi_j^M$, where $j = 1, 2$. The conditional means measuring the association between Y_{i1} and Y_{i2} are

$$E(Y_{i2}|Y_{i1}=0) = \pi^C(0), \quad E(Y_{i2}|Y_{i1}=1) = \pi^C(1);$$

hence $\pi_2^M = \pi_1^M \cdot \pi^C(1) + (1 - \pi_1^M) \cdot \pi^C(0)$. If we incorrectly assume no association between these bivariate binary variables, the maximum likelihood estimators for the marginal means are

$$\tilde{\pi}_1^M = \sum_{i=1}^N Y_{i1} / N, \quad \tilde{\pi}_2^M = \sum_{i=1}^N Y_{i2} R_i / \sum_{i=1}^N R_i.$$

After some algebra, it can be shown that the bias of $\tilde{\pi}_2^M$ is

$$E(\tilde{\pi}_2^M) - \pi_2^M = \{\pi^C(1) - \pi^C(0)\} \{E(\sum_{i=1}^N Y_{i1} R_i / \sum_{i=1}^N R_i) - \pi_1^M\}. \quad (9)$$

From (9) it appears that the bias for $\tilde{\pi}_2^M$ depends on two components: the true dependence for (Y_{i1}, Y_{i2}) and the marginal distributions of Y_{i1} for completers and dropouts. When the marginal distributions for Y_{i1} are the same for completers and dropouts and

$E(\sum_{i=1}^N Y_{i1} R_i / \sum_{i=1}^N R_i) = \pi_1^M$, $\tilde{\pi}_2^M$ remains unbiased. In general this happens when the data are MCAR, that is, missingness is independent of (Y_{i1}, Y_{i2}) . If the data are MAR and in particular the probability that Y_{i2} is missing depends on Y_{i1} but not Y_{i2} , then

$E(\sum_{i=1}^N Y_{i1} R_i / \sum_{i=1}^N R_i) \neq \pi_1^M$ and the bias in $\tilde{\pi}_2^M$ depends on both the dependence structure parameters and the marginal mean of Y_{i1} for completers. On the other hand, if we model the joint distribution of (Y_{i1}, Y_{i2}) , the maximum likelihood estimator for π_2^M is

$$\hat{\pi}_2^M = \hat{\pi}_1^M \cdot \hat{\pi}^C(1) + (1 - \hat{\pi}_1^M) \cdot \hat{\pi}^C(0),$$

where $\hat{\pi}_1^M = \hat{\pi}_1^M$, $\hat{\pi}^C(1) = n_{11}^c / n_{1+}^c$ and $\hat{\pi}^C(0) = n_{01}^c / n_{0+}^c$.

When the data are either MCAR or MAR, it can be shown that $E(\hat{\pi}_2^M) = \pi_2^M$. Thus in this example using information on the dependence structure successfully removes the bias from naive maximum likelihood estimates. The situation is analogous for longer sequences of binary data. In summary, when there are missing data and the missingness is not MCAR, it is important to carefully model the dependence structure as well as the marginal mean structure.

4 Prior specification and posterior inference

With marginal mean parameters (β, θ) and serial dependence parameters (α, ϕ) , the

probability density for $\mathbf{y}_i$ conditioning on covariates $\mathbf{X}_i = (\mathbf{X}_{i1}^T, \dots, \mathbf{X}_{in_i}^T)^T$, $\mathbf{V}_i = (V_{i1}, \dots, V_{in_i})^T$ can be formed as follows:

$$\begin{aligned}
& p(\mathbf{y}_i | \mathbf{X}_i, \mathbf{V}_i; \boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi}) \\
&= p(y_{i1} | \mathbf{X}_{i1}, V_{i1}; \boldsymbol{\beta}, \theta) p(y_{i2} | \mathbf{X}_{i2}, V_{i2}, y_{i1}; \boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi}) \\
&\times \cdots \times p(y_{in_i} | \mathbf{X}_{in_i}, V_{in_i}, y_{i1}, \dots, y_{i, n_i-1}; \boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi}) \\
&= \{\pi_{i1}^M(\boldsymbol{\beta}, \theta)\}^{y_{i1}} \{1 - \pi_{i1}^M(\boldsymbol{\beta}, \theta)\}^{(1-y_{i1})} \\
&\times \prod_{j=2}^{n_i} \{\pi_{ij}^C(\boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi})\}^{y_{ij}} \{1 - \pi_{ij}^C(\boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi})\}^{(1-y_{ij})}.
\end{aligned}$$

Note that $\boldsymbol{\theta}$ and $\boldsymbol{\varphi}$ are coefficients in the P-splines for $m(\cdot)$ and $h(\cdot)$. Assuming independence between individuals, the full data likelihood is

$L(\boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi} | \mathbf{y}, \mathbf{X}, \mathbf{V}) \propto \prod_{i=1}^N p(\mathbf{y}_i | \mathbf{X}_i, \mathbf{V}_i; \boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi})$. It follows that the joint posterior distribution for $(\boldsymbol{\beta}, \alpha, \boldsymbol{\theta}, \boldsymbol{\varphi})$ and associated smoothing parameters $\boldsymbol{\lambda} = (\lambda_1, \lambda_2)$ is proportional to the product of the prior and the likelihood,

$$p(\boldsymbol{\beta}, \alpha, \theta, \boldsymbol{\varphi}, \boldsymbol{\lambda} | \mathbf{y}, \mathbf{X}, \mathbf{V}) \propto p(\boldsymbol{\lambda}) p(\boldsymbol{\beta}, \theta | \boldsymbol{\lambda}) p(\alpha, \boldsymbol{\varphi} | \boldsymbol{\lambda}) L(\boldsymbol{\beta}, \theta, \alpha, \boldsymbol{\varphi} | \mathbf{y}, \mathbf{X}, \mathbf{V}),$$

where $p(\boldsymbol{\lambda})$ is a hyper-prior for $\boldsymbol{\lambda}$, $p(\boldsymbol{\beta}, \theta | \boldsymbol{\lambda})$ and $p(\alpha, \boldsymbol{\varphi} | \boldsymbol{\lambda})$ are the priors for $(\boldsymbol{\beta}, \theta)$ and $(\alpha, \boldsymbol{\varphi})$ conditioning on $\boldsymbol{\lambda}$.

We use standard prior distributions for the HERS example in Section 6 and the simulation study in Section 5. Conditioning on $\boldsymbol{\lambda}$, all members of $(\boldsymbol{\beta}, \theta)$ and $(\alpha, \boldsymbol{\varphi})$ are assumed *a priori* independent and normally distributed. Similar to $\boldsymbol{\theta} = (\theta_1, \theta_2)$ in (2), we define $\boldsymbol{\varphi}_1 = (\varphi_0, \dots, \varphi_s)^T$ and $\boldsymbol{\varphi}_2 = (\varphi_{s+1}, \dots, \varphi_{s+K^*})^T$ in (8). For $(\boldsymbol{\beta}, \boldsymbol{\theta}_1)$ and $(\alpha, \boldsymbol{\varphi}_1)$, we assign vague normal priors with mean zero and large variance; for $\boldsymbol{\theta}_2$ and $\boldsymbol{\varphi}_2$, we assign normal priors with mean zero and common variances σ_1^2 and σ_2^2 . Therefore in terms of the original setup of P-splines in Section 2, we have penalty matrices

$$\begin{aligned}
\mathbf{D}_1 &= \begin{bmatrix} \mathbf{0}_{(q+1) \times (q+1)} & \mathbf{0}_{(q+1) \times K} \\ \mathbf{0}_{K \times (q+1)} & \mathbf{I}_{K \times K} \end{bmatrix}, \\
\mathbf{D}_2 &= \begin{bmatrix} \mathbf{0}_{(s+1) \times (s+1)} & \mathbf{0}_{(s+1) \times K^*} \\ \mathbf{0}_{K^* \times (s+1)} & \mathbf{I}_{K^* \times K^*} \end{bmatrix},
\end{aligned}$$

and smoothing parameters $\lambda_1 - (2\sigma_1^2)^{-1}$ and $\lambda_2 - (2\sigma_2^2)^{-1}$ respectively. Note that we assume global smoothing parameters here in the sense that all elements of $\boldsymbol{\theta}_2$ and $\boldsymbol{\varphi}_2$ receive the same amount of penalty ([22]). The hyper-priors for σ_1^2 and σ_2^2 also need to be specified. In this article we use uniform priors with wide ranges for hierarchical standard deviations σ_1 and σ_2 , as recommended by Gelman (2005) ([35]).

We use MCMC to sample from the posterior distributions; summary statistics such as posterior means and 95% credible bands are then used for inference. The MCMC is implemented using WinBUGS package (version 1.4.1) and its development interface ([32]).

5 Simulation study

Our model gives the analyst flexibility in specifying the marginal mean and serial dependence structures for longitudinal binary processes with missing data. In this section,

we conduct a simulation study to evaluate the performance of the proposed model and other models with parametric serial dependence structures under MAR.

5.1 Simulation setup

For simplicity, we focus on the estimation of $m(\cdot)$ in the marginal mean. Longitudinal binary data with MAR dropout are simulated via a first-order marginalized transition model. The true marginal mean of the binary outcome is a nonlinear function of the continuous measurement time on the logit scale, while the first-order serial dependence remains high when the time lag is small and then decays rapidly to a positive constant as the time lag increases. The dropout probability at each measurement time depends only on the previous observed outcome and the first outcome in the sequence is always observed. The details of simulation setup are summarized in Table 2.

By comparing the following models with different serial dependence structures, we aim to investigate the role of serial dependence on valid inference about the marginal mean.

- a. Bayesian P-splines fit assuming independence between observations with $\gamma_{ij} = 0$.
- b. Semiparametric MTM(1) with constant serial dependence $\gamma_{ij} = \alpha$.
- c. Semiparametric MTM(1) with a linear parametric serial dependence structure $\gamma_{ij} = \alpha_1 + \alpha_2 \cdot u_{ij}^{1/2}$.
- d. Semiparametric MTM(1) with a semiparametric serial dependence structure $\gamma_{ij} = h(u_{ij}^{1/2})$ and $h(\cdot)$ fitted by P-splines.

The marginal mean structures are correctly specified for all models as in Table 2. In each simulation, Bayesian P-splines with truncated linear bases are used for estimating $m(V_{ij})$, with knots set at 5% quantiles of the V_{ij} s. To reduce the right skewness of the time lag distribution, the MTM(1) with linear parametric and semiparametric serial dependence structures are specified using the square root transformation of u_{ij} ; for $h(u_{ij})$, truncated linear bases on 10% quantiles are used. We also impose a restriction on the estimation of $h(u_{ij}^{1/2})$ by forcing the $h(\cdot)$ function to be constant after the 90th percentile.

Further, the following prior distributions are used: $\theta_1 \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$, $\theta_2 | \sigma_1 \sim N(\mathbf{0}, \sigma_1^2 \cdot \mathbf{I})$, $\sigma_1 \sim \text{Unif}(0, 5)$; $\alpha \sim N(0, 10^4)$, $\alpha_1 \sim N(0, 10^4)$, $\alpha_2 \sim N(0, 10^4)$, $\phi_1 \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$,

$\theta_2 | \sigma_2 \sim N(\mathbf{0}, \sigma_2^2 \cdot \mathbf{I})$ and $\sigma_2 \sim \text{Unif}(0, 5)$. It is impractical to assess the convergence of MCMC chains for all data points on $m(\cdot)$ in all simulated data sets. Instead, we take a random sample of $m(V_{ij})$ s for each data set and use history plots and Gelman and Rubin convergence statistics ([36]) provided by WinBUGS to check the convergence. For each simulated data set, we run three chains with relatively diverse initial values; and after convergence the pooled 3000 posterior samples from 1000 iterations are used for inference.

5.2 Misspecification bias

We focus on misspecification bias to evaluate the model performance. First define the mean squared bias as

$$\text{Mean squared bias} = \left(\sum_{i=1}^N n_i \right)^{-1} \sum_{i=1}^N \sum_{j=1}^{n_i} \{ \bar{m}(V_{ij}) - m(V_{ij}) \}^2,$$

where $N = 500$, $n_i = 5$ and $\bar{m}(V_{ij})$ are the sample averages of the estimated posterior means for $m(V_{ij})$ across 100 simulations. As a benchmark for comparison, the mean squared bias from complete data using smoothing splines fit with leave-one-out cross validation is 1.6×10^{-3} ; with incomplete data the mean squared biases for the four model fits are 596.3×10^{-3} , 64.9×10^{-3} , 4.7×10^{-3} and 3.4×10^{-3} respectively.

In addition, the top left panel of Figure 4 presents the $\bar{m}(V_{ij})$ curves and the true curve $m(V)$. Note that $\bar{m}(V_{ij})$ are themselves smoothed to diminish Monte Carlo variability. Because the subjects with $Y_{ij} = 1$ are dropping out as the covariate V increases, it is not surprising that the Bayesian P-splines fit assuming independence between observations shows increasing bias as the covariate V increases; and its mean squared bias is the largest (596.3×10^{-3}). Given that $V_{ij} \in (0, 1)$ and $m(V_{ij}) \in (-1.20, 0.66)$ in Table 2, the magnitude of this deviation is approximately 32% of the range of the true $m(V)$, which is considerably large in this setting. The MTM(1) with constant serial dependence has dramatically reduced mean squared bias (64.9×10^{-3}) compared to the naive fit, although it still shows more bias than the other two semiparametric MTM(1) when the covariate value is greater than 0.4. The two MTM(1) fits with linear parametric and semiparametric serial dependence are both quite close to the true curve (because graphically it is difficult to distinguish the curves from the two models with unaided eyes, we do not present the curve from the linear parametric serial dependence model in Figure 4), and their mean squared biases (4.7×10^{-3} and 3.4×10^{-3}) are comparable to the bias (1.6×10^{-3}) from complete data. The MTM(1) with semiparametric serial dependence performs slightly better when the curvatures are large, which is sensible because there is not much space for improvement from the MTM(1) fit with linear parametric serial dependence. The noticeable biases at large curvatures might be due to the sparseness of the binary data and the consequent over-smoothness of the curves ([26,28]).

Overall, the results for serial dependence models are consistent with the results for estimating $m(V_{ij})$ as seen from the top right panel of Figure 4. In this case, the posterior mean estimates from MTM(1) with semiparametric serial dependence are the closest to the true $h(u)$ for the majority of u_{ij} s, except the constant estimates beyond the 90% quantile forced by our specific restriction. We also notice that there still exists some discrepancy between the true $h(u)$ and the estimates produced by the semiparametric serial dependence model; but again this might be explained by the sparseness of the binary data and the situation is even worse for estimating serial dependence parameters than for estimating mean regression parameters. The mean squared biases for constant, linear parametric and semiparametric serial dependence models are 823.2×10^{-3} , 165.1×10^{-3} and 40.4×10^{-3} respectively, with the same ranking as for the case of $m(V_{ij})$. However, since the linear parametric and semiparametric serial dependence models both capture the general pattern of the true $h(u)$, they tend to give similar estimates for $m(V_{ij})$. In contrast, the constant serial dependence model ignores the decreasing trend of serial dependence; hence its estimates for $m(V_{ij})$ still show considerable discrepancy from the true function when the dropout becomes severe, although the bias is much smaller compared with the naive fit. Thus in our simulations it appears that under MAR, properly modeling the serial dependence structure does reduce the bias for estimating the marginal mean.

5.3 Variability

Apart from misspecification bias, we are also concerned about the possibly increased variability due to semiparametric serial dependence modeling. Here we do not evaluate the model-based variability estimates in our simulations. Instead we focus on simple comparisons using Monte Carlo variability estimates and defer the discussion on model-based variability estimates to Section 6.

As shown in the bottom panel of Figure 4, in all model fits the smoothed pointwise mean squared errors (MSE) on $\log_{10}$ scale demonstrate the same patterns as for mean squared biases. Semiparametric MTM(1) with linear parametric and semiparametric serial dependence structures perform similarly but much better than the naive fit and constant serial dependence fit (we do not present the MSE from the linear parametric serial dependence model in Figure 4 due to difficulty in graphical presentation). Note that the MSE from all fits are increasing when the covariate value increases and the dropout becomes severe; this makes sense because the increased variability results from fewer available data points due to dropouts. Thus it turns out that the possibly increased variability by semiparametric regression modeling does not overwhelm the effect of bias reduction in our simulations.

6 Application to HIV cohort study

In this section, we illustrate our model using the HERS viral load data introduced in Section 1.3. Recall that the analysis objective is to characterize the profile of detectable viral load prevalence before and after HAART initiation and investigate whether there is an HCV coinfection effect associated with this time course, after accounting for baseline and longitudinal CD4 count levels.

Specifically, we assume that the marginal mean of Y_{ij} (i.e., detectability of viral load) follows

$$\text{logit}(\pi_{ij}^M) = \beta_1 X_{1i} + \beta_2 X_{2i} + \beta_3 X_{3ij} + m(t_{ij}), \quad (10)$$

where $X_{1i} \in \{0, 1\}$ is the indicator of being HCV negative at baseline, X_{2i} is the CD4 count at baseline (scaled by 100), X_{3ij} is the CD4 count change (scaled by 100) from baseline at the j th visit, $\beta = (\beta_1, \beta_2, \beta_3)$ are regression coefficients, and t_{ij} is the time in days since reported HAART initiation for the i th individual at the j th visit (scaled by 1000). Note that an interaction term between HCV status and time since HAART initiation can also be included in the marginal mean structure. We fitted this model and the results showed no evidence of differential HCV effects before and after HAART initiation. Therefore the model in (10) is used instead.

The serial dependence follows

$$\text{logit}(\pi_{ij}^C) = \Delta_{ij} + \gamma_{ij} \cdot Y_{i,j-1},$$

and the following regression models for γ_{ij} are used for comparison:

- a. $\gamma_{ij} = 0$,
- b. $\gamma_{ij} = \alpha_0 + \alpha_1 X_{1i}$,
- c. $\gamma_{ij} = \alpha_2 + \alpha_3 X_{1i} + \alpha_4 u_{ij}$,
- d. $\gamma_{ij} = \alpha_5 X_{1i} + h(u_{ij})$,

where $u_{ij} = t_{ij} - t_{i,j-1}$ is the number of days elapsed between adjacent measurements (scaled by 1000) and X_{1i} again is the baseline HCV coinfection status.

We model both $m(\cdot)$ and $h(\cdot)$ using Bayesian P-splines with 5% quantile truncated quadratic bases. Further, the following prior distributions are specified: $\theta_1 \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$, $\theta_2 | \sigma_1 \sim N(\mathbf{0}, \sigma_1^2 \cdot \mathbf{I})$, $\sigma_1 \sim \text{Unif}(0, 10)$; $\varphi_1 \sim N(\mathbf{0}, 10^4 \cdot \mathbf{I})$, $\varphi_2 | \sigma_2 \sim N(\mathbf{0}, \sigma_2^2 \cdot \mathbf{I})$ and $\sigma_2 \sim \text{Unif}(0, 10)$. These priors are selected because of their relative flexibility. They are all proper, yet suitable and vague. We run two MCMC chains and assess convergence using history plots and Gelman Rubin convergence statistics for randomly selected data points on $m(\cdot)$ and $h(\cdot)$ as well as for other parameters. After convergence, pooled posterior samples of size 10,000 are used for inference.

6.1 Parameter estimates for marginal mean

Table 3 presents the summary statistics for posterior distributions of marginal mean regression coefficients. From the 95% credible intervals for β_1 , there is no evidence for an HCV effect on viral load prevalence profile under four different serial dependence models, although the signs for β_1 estimates indicate that the HCV positive patients were more likely to have viral suppression in this HERS population. This is not sensible and the model with independence assumption gives the largest point estimate for this HCV effect. The phenomenon might be explained by the amount of missing data and the slightly differential follow-up times between these HCV groups. The HCV positive group tended to have shorter follow-up (66 months on average) with more missing outcomes than HCV negative group (69 months of follow-up on average). Because patients with detectable viral load were more likely to have missing observations, the simple average of available data (as in the model assuming independence) in the HCV positive group will underestimate the viral load prevalence. The same problem also applies for the HCV negative group, but it is less severe because of the amount of missingness. Overall, this possibly created a beneficial effect of being HCV positive and in the model with independence assumption the estimated effect is probably the most biased.

For baseline CD4 effect, results from all models show that patients with higher baseline CD4 counts were more likely to have viral suppression and posterior mean estimates are approximately the same across the models. Further, patients with larger CD4 changes from baseline also tended to have viral suppression; and here the estimated effect from the model with independence assumption is also the largest. Because CD4 count and viral load are generally negatively correlated, it is not surprising that the patients with lower observed viral load and higher observed CD4 count change were more likely to remain in the study. Therefore, in this scenario the negative relationship between viral load prevalence and CD4 count change will be overestimated by assuming no serial dependence and computing the simple averages of the available outcomes.

The posterior mean estimates for $m(t_{ij})$ are similar under different serial dependence specifications. Here we only present the results from model (d), i.e., the semiparametric serial dependence model. Figure 5 shows the posterior mean and pointwise 95% credible band for $m(t_{ij})$ on the probability scale. After accounting for other covariates, the viral load prevalence remained approximately constant until about 6 months before the visit with first reported HAART use. Recall that the patients in the HERS only reported whether they had HAART in the past 6 months during the follow-up. Therefore the actual HAART initiation for them happened in a 6-month window before the visits with reported HAART use. From Figure 5, it is apparent that the viral load prevalence dropped dramatically in this 6-month window for HAART initiation, which clearly indicates the effectiveness of HAART on viral suppression in the HERS population. However, after some period of HAART treatment, the viral load prevalence reached its bottom and began to go up. This phenomenon could be explained by the viral load rebound due to related drug resistance and/or poor adherence after extended period of treatment.

As pointed out by a referee, model-based variability estimates (posterior standard deviations) of the marginal mean parameters from misspecified serial dependence models can be invalid under MAR even if the point estimates of those parameters are similar across models. As a crude summary of the model-based variability estimates, an average of the posterior standard deviations for all $m(t_{ij})$ s is calculated for all models. We find that the model assuming independence has the smallest variability estimates for $m(t_{ij})$ (0.19), while those from other serial dependence models are similar (0.22). This is expected because in our case estimating a specific value on $m(t_{ij})$ is similar to estimating an intercept term in the marginal mean and it is well known that the model assuming independence will underestimate the variability for the coefficient of a time-invariant covariate (e.g., the intercept). The same phenomenon is observed for β_1 , the coefficient for HCV effect, where the posterior standard deviation is smaller (0.09) in the model assuming independence than in other serial dependence models (0.13).

6.2 Parameter estimates for serial dependence

Table 4 presents the summary statistics for posterior distributions of the regression coefficients in the serial dependence models and the fitted $h(u_{ij})$ function in the semiparametric serial dependence model is given in Figure 6. Results in all four fitted models indicate large serial dependence in the HERS viral load data. For example, the posterior means of intercept terms in model (b) and (c) as well as the fitted values for $h(u_{ij})$ when the time lags are less than 6 months in model (d) are all above two on the logit scale. There is no evidence that the serial dependence differs by HCV groups. Both model (c) and model (d) suggest that the serial dependence is decreasing as the time lag between adjacent outcomes increases; and the fitted $h(u_{ij})$ curve in Figure 6 is approximately linear, which suggests that model (c) as a parametric serial dependence model might be sufficient to characterize the dependence structure in the HERS viral load data. The rebound of serial dependence in the right tale of the time lag is due to the small sample size as indicated by the increasing width of pointwise 95% credible band.

6.3 Model comparison

Deviance Information Criterion (DIC) is a useful Bayesian model comparison approach ([37]). However, the literature is sparse on recommendations for constructing the DIC for incomplete data ([38]). Under ignorable missingness, an obvious choice would be to construct the DIC based on the observed data likelihood. This approach has been applied in the literature ([39,40]), and can generally be implemented in WinBUGS. Nevertheless, for the first place we should keep in mind that given all the possible choices for model comparison techniques, we can only assess the fit of the full data model to the observed data. The adequacy of the missing data assumption can not be evaluated.

In our HERS example, we compute the DIC based on the observed data likelihood and compare the fits from the four serial dependence models. The last two rows of Table 3 give the DIC values and the effective numbers of parameters for the models. The model assuming independence has a much larger DIC value than other models. Model (c) and model (d) fit similarly but slightly better than model (b). This is consistent with our findings on the relationship between serial dependence and time lag shown in Figure 6.

In summary, the results for our HERS data example show that in practice the models for serial dependence can affect both the point estimates and variability estimates for marginal mean parameters under missingness. In the presence of missing data, caution needs to be taken when building structures for both marginal mean and serial dependence.

7 Conclusion

We have proposed a fully likelihood-based approach to analyzing longitudinal binary process data with missingness. Our model features semiparametric regression formulations in both marginal mean and serial dependence structures. Under MAR, our semiparametric serial dependence model allows more flexible estimation of the within-individual serial dependence, which can benefit the marginal mean estimation as shown in the simulation study. In practice semiparametric models can be used as a tool for exploring the full serial dependence structure.

Several aspects warrant further research. First, extension to higher order Markov models for serial dependence is necessary. In addition, random effects can be included in the dependence structure in order to account for the dependence induced by unexplained between-subject variability ([41,42,43]). Second, in reality when MAR is violated and missing data is not ignorable, it is useful to generalize our model and provide sensitivity analysis using either selection model ([44]) or mixture model framework ([45,46]). Third, with continuous measurement times, the timing of observations can be informative in the sense that it is related to the outcome process ([27,47,48]). It may be worthwhile to extend our model to accommodate the informative observation times. Lastly, further potential applications can include time series, spatial statistics and molecular genetics. Our methods are readily applicable to single or replicated binary time series data when the analysis objective is to explain covariate effects ([49]). For spatial binary data, our model shares certain characteristics with autologistic regression models. For example, the Markov assumption in our model is similar to Markov random fields defined in the autologistic models ([50]). In molecular genetics our model could be useful for the problem of identifying multiple change-points ([51]) and extreme regions of gene expression on chromosomes in human tumors ([52]).

Acknowledgments

We are grateful to Jeffrey Blume, Mike Daniels, Constantine Gatsonis, Patrick Heagerty, Tony Lancaster and the referees for helpful comments. This research was supported by NIH grants R01-AI-50505 and R01-HL-79457. Data for the HERS were collected under CDC cooperative agreements U64-CCU106795, U64-CCU206798, U64-CCU306802 and U64-CCU506831.

References

1. Breslow NE, Clayton DG. Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association*. 1993; 88:9–25.
2. Liang KY, Zeger SL. Longitudinal data analysis using generalized linear models. *Biometrika*. 1986; 73:13–22.
3. Diggle, PJ.; Heagerty, P.; Liang, K-Y.; Zeger, SL. *Analysis of Longitudinal Data*. Oxford University Press; New York: 2002.
4. Severini TA, Staniswalis JG. Quasi-likelihood estimation in semiparametric models. *Journal of the American Statistical Association*. 1994; 89:501–511.
5. Wild CJ, Yee TW. Additive extensions to generalized estimating equation methods. *Journal of the Royal Statistical Society Series B*. 1996; 58:711–725.
6. Lin X, Carroll RJ. Nonparametric function estimation for clustered data when the predictor is measured without/with error. *Journal of the American Statistical Association*. 2000; 95:520–534.
7. Lin X, Carroll RJ. Semiparametric regression for clustered data. *Biometrika*. 2001; 88:1179–1865.
8. Lin X, Carroll RJ. Semiparametric regression for clustered data using generalized estimating equations. *Journal of the American Statistical Association*. 2001; 96:1045–1056.

9. Welsh AH, Lin X, Carroll RJ. Marginal longitudinal nonparametric regression: Locality and efficiency of spline and kernel methods. *Journal of the American Statistical Association*. 2002; 97:482–493.
10. Wang N, Carroll RJ, Lin X. Efficient semiparametric marginal estimation for longitudinal/clustered data. *Journal of the American Statistical Association*. 2005; 100:147–157.
11. Fitzmaurice GM, Laird NM. A likelihood-based method for analysing longitudinal binary responses. *Biometrika*. 1993; 80:141–151.
12. Heagerty PJ. Marginalized transition models and likelihood inference for longitudinal categorical data. *Biometrics*. 2002; 58:342–351. [PubMed: 12071407]
13. Azzalini A. Logistic regression for autocorrelated data with application to repeated measures. *Biometrika*. 1994; 81:767–775.
14. Little, RJA.; Rubin, DB. *Statistical Analysis with Missing Data*. 2. John Wiley & Sons; New York: 2002.
15. Barndorff-Nielsen, OE.; Cox, DR. *Inference and Asymptotics*. Chapman & Hall; 1994.
16. Kurland BF, Heagerty PJ. Marginalized transition models for longitudinal binary data with ignorable and non-ignorable drop-out. *Statistics in Medicine*. 2004; 23:2673–2695. [PubMed: 15316952]
17. Heagerty PJ, Zeger SL. Lorelogram: A regression approach to exploring dependence in longitudinal categorical responses. *Journal of the American Statistical Association*. 1998; 93:150–162.
18. Mayer KH, Hogan JW, Smith D, Klein RS, Schuman P, Margolick JB, Korkontzelou C, Farzedegan H, Vlahov D, Carpenter CC. Clinical and immunologic progression in HIV-infected US women before and after the introduction of highly active antiretroviral therapy. *Journal of Acquired Immune Deficiency Syndromes*. 2003; 33:614–624. [PubMed: 12902807]
19. Monga HK, Rodriguez-Barradas MC, Breaux K, Khattak K, Troisi CL, Velez M, Yoffe B. Hepatitis c virus infection-related morbidity and mortality among patients with human immunodeficiency virus infection. *Clinical Infectious Diseases*. 2001; 33:240–247. [PubMed: 11418885]
20. Sterling TR, Vlahov D, Astemborski J, Hoover DR, Margolick JB, Quinn TC. Initial plasma HIV-1 RNA levels and progression to AIDS in women and men. *New England Journal of Medicine*. 2001; 344:720–725. [PubMed: 11236775]
21. Eilers PHC, Marx BD. Flexible smoothing with B-splines and penalties (with discussion). *Statistical Science*. 1996; 11:89–121.
22. Ruppert D, Carroll RJ. Spatially adaptive penalties for spline fitting. *Australian and New Zealand Journal of Statistics*. 2000; 42:205–223.
23. Ruppert, D.; Wand, MP.; Carroll, RJ. *Semiparametric Regression*. Cambridge University Press; Cambridge: 2003.
24. Zeger SL, Diggle PJ. Semi-parametric models for longitudinal data with application to CD4 cell numbers in HIV seroconverters. *Biometrics*. 1994; 50:689–699. [PubMed: 7981395]
25. Zhang D, Lin X, Raz J, Sowers M. Semi-parametric stochastic mixed models for longitudinal data. *Journal of the American Statistical Association*. 1998; 93:710–719.
26. Lin X, Zhang D. Inference in generalized additive mixed models using smoothing splines. *Journal of the Royal Statistical Society Series B*. 1999; 61:381–400.
27. Lin DY, Ying Z. Semiparametric and nonparametric regression analysis of longitudinal data (with discussion). *Journal of the American Statistical Association*. 2001; 96:103–126.
28. Zhang D. Generalized linear mixed models with varying coefficients for longitudinal data. *Biometrics*. 2004; 60:8–15. [PubMed: 15032768]
29. Ruppert D. Selecting the number of knots for penalized splines. *Journal of Computational and Graphical Statistics*. 2002; 11:735–757.
30. Brumback B, Ruppert D, Wand MP. Comment on variable selection and function estimation in additive nonparametric regression using data-based prior by Shively, Kohn and Wood. *Journal of the American Statistical Association*. 1999; 94:794–797.

31. Crainiceanu CM, Ruppert D, Wand MP. Bayesian analysis for penalized spline regression using WinBUGS. *Journal of Statistical Software*. 2005; 14(14):1–24.
32. Spiegelhalter, D.; Thomas, A.; Best, N.; Lunn, D. WinBUGS Version 1.4 User Manual. Medical Research Council Biostatistics Unit; Cambridge: 2003.
33. Hastie T, Tibshirani R. Generalized additive models. *Statistical Science*. 1986; 1:297–310.
34. Pepe M, Anderson G. A cautionary note on inference for marginal regression models with longitudinal data and general correlated response data. *Communications in Statistics B*. 1994; 23:939–951.
35. Gelman A. Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*. 2005; 1(2):1–19.
36. Brooks SP, Gelman A. Alternative methods for monitoring convergence of iterative simulations. *Journal of Computational and Graphical Statistics*. 1998; 7:434–455.
37. Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society Series B*. 2002; 64:583–639.
38. Daniels, MJ.; Hogan, JW. *Monographs on Statistics and Applied Probability*. Vol. 101. Chapman & Hall/CRC; New York: 2007. *Missing Data in Longitudinal Studies: Strategies for Bayesian Modeling and Sensitivity Analysis*. in press
39. Pourahmadi M, Daniels MJ. Dynamic conditionally linear mixed models for longitudinal data. *Biometrics*. 2002; 58:225–231. [PubMed: 11890319]
40. Ilk O, Daniels MJ. Marginalized transition random effects models for multivariate longitudinal binary data. *Canadian Journal of Statistics*. 2007; 35:105–123.
41. Heagerty PJ, Zeger SL. Marginalized multilevel models and likelihood inference. *Statistical Science*. 2000; 15:1–26.
42. Miglioretti DL, Heagerty PJ. Marginal modeling of multilevel binary data with time-varying covariates. *Biostatistics*. 2004; 5:381–398. [PubMed: 15208201]
43. Schildcrout JS, Heagerty PJ. Marginalized models for moderate to long series of longitudinal binary response data. *Biometrics*. 2006; 62:1111–1120. doi:10.1111/j.1541-0420.2006.00680.x
44. Diggle PJ, Kenward MG. Informative dropout in longitudinal data analysis (with discussion). *Applied Statistics*. 1994; 43:49–93.
45. Little RJA. Pattern-mixture models for multivariate incomplete data. *Journal of the American Statistical Association*. 1993; 88:125–134.
46. Hogan JW, Laird NM. Mixture models for the joint distribution of repeated measures and event times. *Statistics in Medicine*. 1997; 16:239–257. [PubMed: 9004395]
47. Lipsitz SR, Fitzmaurice GM, Ibrahim JG, Gelber R, Lipshultz S. Parameter estimation in longitudinal studies with outcome-dependent follow-up. *Biometrics*. 2002; 58:621–630. [PubMed: 12229997]
48. Lin H, Scharfstein DO, Rosenheck RA. Analysis of longitudinal data with irregular, outcome-dependent follow-up. *Journal of the Royal Statistical Society Series B*. 2004; 66:791–813.
49. Zeger SL, Irizarry RA, Peng RD. On time series analysis of public health and biomedical data. *Annual Review of Public Health*. 2006; 27:57–79.
50. Banerjee, S.; Carlin, BP.; Gelfand, AE. *Monographs on Statistics and Applied Probability*. Vol. 101. Chapman & Hall/CRC; New York: 2004. *Hierarchical Modeling and Analysis for Spatial Data*.
51. Albert PS, Hunsberger SA, Hu N, Taylor PR. Identifying multiple changepoints in heterogeneous binary data with an application to molecular genetics. *Biostatistics*. 2004; 5:515–529. [PubMed: 15475416]
52. Levin AM, Ghosh D, Cho KR, Kardia SLR. A model-based scan statistic for identifying extreme chromosomal regions of gene expression in human tumors. *Bioinformatics*. 2005; 21:2867–2874. [PubMed: 15814559]

Figure 1.

Nine individual viral load profiles from the HERS data; the time scale is the days since reported HAART initiation; the black circles are the data from HCV positive patients, the black triangles are the data from HCV negative patients.

Figure 2. Empirical estimates of the detectable viral load prevalence profile following HAART initiation in the HERS (fitted by moment-based smoothing splines with leave-one-out cross validation).

Figure 3.

The association between binary outcomes and follow-up times in the HERS: the proportion of observations (within individuals) with viral load > 50 copies/mL as a function of follow-up times; the solid line is the smoothing splines fit for the function.

Figure 4.

Top left: smoothed averages of posterior means for $m(V_{ij})$; top right: smoothed averages of posterior means for $\gamma_{ij} = h(u_{ij})$; bottom: smoothed pointwise mean squared errors (on $\log_{10}$ scale) for $m(V_{ij})$; the dark solid lines are true $m(V)$ or $h(u)$ functions; (a) Bayesian P-splines fit with $\gamma_{ij} = 0$; (b) semiparametric MTM(1) with $\gamma_{ij} = \alpha$; (c) semiparametric MTM(1) with $\gamma_{ij} = \alpha_1 + \alpha_2 \cdot u_{ij}^{1/2}$ (both smoothed averages of posterior means for $m(V_{ij})$ and pointwise mean squared errors in the linear parametric serial dependence model are graphically close to those from the semiparametric serial dependence model, thus we do not present them here); (d) semiparametric MTM(1) with $\gamma_{ij} = h(u_{ij}^{1/2})$ fitted by P-splines; on the top of the top right panel are sample quantiles by 10% for u_{ij} s in complete data.

Figure 5. Fitted HIV viral load prevalence profile ($m(t_{ij})$ on probability scale) from a semiparametric serial dependence model for the HERS data; the solid line is posterior mean estimates; the dotted lines are 95% pointwise credible band; the gray lines are the 6-month window for HAART initiation.

Figure 6.

Fitted $h(u_{ij})$ versus time lags from a semiparametric serial dependence model for the HERS data; the solid line is posterior mean estimates; the dotted lines are 95% pointwise credible band.

Table 1

Contingency tables for bivariate binary outcomes with dropouts.

Completers				
		Y_2		
		0	1	
Y_1	0	n_{00}^c	n_{01}^c	n_{0+}^c
	1	n_{10}^c	n_{11}^c	n_{1+}^c
		n_{+0}^c	n_{+1}^c	r
Dropouts				
		Y_2		
		0	1	
Y_1	0	?	?	n_{0+}^d
	1	?	?	n_{1+}^d
		$N - r$		

Table 2

Data generation for misspecification bias simulations.

<i>Marginal mean</i>	$\text{logit}\{E(Y_{ij} V_{ij})\} = m(V_{ij})$ $V_{ij} = t_{ij} \stackrel{i.i.d}{\sim} \text{Unif}(0, 1)$ $m(V) = \exp(-2V) \cdot [\sin\{10(V + 0.1)\} + 1] - 1.2$ $i = 1, \dots, N; j = 1, \dots, n_i$
<i>Serial dependence</i>	$\text{logit}\{E(Y_{ij} Y_{i,j-1} = y_{i,j-1})\} = \Delta_{ij} + \gamma_{ij} \cdot y_{i,j-1}$ $\gamma_{ij} = h(u_{ij})$ $u_{ij} = t_{ij} - t_{i,j-1}$ $h(u) = \text{dnorm}_{0.35, 0.15}(\sqrt{u}) \cdot I(\sqrt{u} > 0.35) + \text{dnorm}_{0.35, 0.15}(\sqrt{0.35}) \cdot I(\sqrt{u} \leq 0.35) + 0.5 \text{dnorm}_{a,b}(\cdot)$ <p>$\text{dnorm}_{a,b}(\cdot)$ is a normal density function with mean a and standard deviation b; $I(\cdot)$ is the indicator function.</p>
<i>Dropout model</i>	$R_{ij} = 1 \text{ if } Y_{ij} \text{ is observed, o.w. } R_{ij} = 0 \text{ for } l \geq j$ $R_{i1} = 1 \text{ with probability } 1$ $\text{pr}(R_{ij} = 1 R_{i1} = \dots = R_{i,j-1} = 1, \mathbf{Y}_i)$ $= 1 \text{ if } Y_{i,j-1} = 1,$ $= 0.2 \text{ if } Y_{i,j-1} = 0, j = 2, \dots, n_i.$ <p>approximate dropout rate=55%</p>
<i>Other details</i>	<p>sample size: $N = 500$</p> <p>100 simulated data sets per scenario cluster size: $n_i = 5$ for complete data</p>

Table 3

Posterior mean, posterior standard deviation (in parentheses), 95% credible interval (in brackets) for regression coefficients in the marginal mean and model comparisons for the serial dependence models of the HERS data

Parameters	Model (a)	Model (b)	Model (c)	Model (d)
HCV status (β_1)	0.065(0.094)	0.020(0.129)	0.021(0.131)	0.033(0.128)
	[−0.120, 0.250]	[−0.236, 0.272]	[−0.231, 0.274]	[−0.216, 0.288]
Baseline CD4 (β_2)	−0.214(0.023)	−0.202(0.029)	−0.205(0.031)	−0.208(0.029)
	[−0.259, −0.169]	[−0.257, −0.145]	[−0.268, −0.147]	[−0.266, −0.152]
CD4 change (β_3)	−0.330(0.024)	−0.286(0.028)	−0.284(0.029)	−0.289(0.028)
	[−0.378, −0.283]	[−0.339, −0.232]	[−0.342, −0.228]	[−0.345, −0.234]
DIC	3064.580	2693.350	2690.510	2690.540
Effective number of parameters (p_D)	12.680	13.962	14.535	16.402

Table 4

Posterior mean, posterior standard deviation (in parentheses) and 95% credible interval (in brackets) for regression coefficients in the serial dependence models for the HERS data

Parameters	Model (b)	Model (c)	Model (d)
Intercept	2.262(0.170) [1.943, 2.593]	2.805(0.307) [2.197, 3.388]	— —
HCV status	0.110(0.244) [−0.386, 0.583]	0.082(0.245) [−0.397, 0.567]	0.044(0.256) [−0.452, 0.553]
Time lag between adjacent outcomes	— —	−2.636(1.251) [−5.114, −0.131]	— —