



Flexible Bayesian modeling for longitudinal binary and ordinal responses

Jizhou Kang¹ · Athanasios Kottas¹

Received: 11 March 2024 / Accepted: 14 October 2024
© The Author(s) 2024

Abstract

Longitudinal studies with binary or ordinal responses are widely encountered in various disciplines, where the primary focus is on the temporal evolution of the probability of each response category. Traditional approaches build from the generalized mixed effects modeling framework. Even amplified with nonparametric priors placed on the fixed or random effects, such models are restrictive due to the implied assumptions on the marginal expectation and covariance structure of the responses. We tackle the problem from a functional data analysis perspective, treating the observations for each subject as realizations from subject-specific stochastic processes at the measured times. We develop the methodology focusing initially on binary responses, for which we assume the stochastic processes have Binomial marginal distributions. Leveraging the logits representation, we model the discrete space processes through continuous space processes. We utilize a hierarchical framework to model the mean and covariance kernel of the continuous space processes nonparametrically and simultaneously through a Gaussian process prior and an Inverse-Wishart process prior, respectively. The prior structure results in flexible inference for the evolution and correlation of binary responses, while allowing for borrowing of strength across all subjects. The modeling approach can be naturally extended to ordinal responses. Here, the continuation-ratio logits factorization of the multinomial distribution is key for efficient modeling and inference, including a practical way of dealing with unbalanced longitudinal data. The methodology is illustrated with synthetic data examples and an analysis of college students' mental health status data.

Keywords Bayesian hierarchical modeling · Continuation-ratio logits · Functional data analysis · Markov chain Monte Carlo · Student-t process

1 Introduction

Recent years have witnessed a rapid growth of longitudinal studies with binary and ordinal responses in several disciplines, including econometrics, and the health and social sciences. In such studies, of primary importance are the probability response curves, i.e., the probabilities of the response categories that evolve dynamically over time. This article aims at developing a hierarchical framework, customized to longitudinal settings, that allows flexible inference for the probability response curves. In addition, the defining charac-

teristic of longitudinal data is that repeated measurements on the same subject induce dependence. Hence, a further objective is to flexibly model lead-lag correlations among repeated measurements.

The development of statistical methods for longitudinal binary and ordinal data stems from models for longitudinal continuous responses, postulating the generalized linear model framework. Analogous to the continuous case, a specific model is formulated under one of three broad approaches pertaining to marginal models, conditional models, or subject-specific models. Marginal models provide alternative modeling options when likelihood-based approaches are difficult to implement. A conditional model describes the distribution of responses conditional on the covariates and also on part of the other components of the responses. In a subject-specific model, the effects of a subset of covariates are allowed to vary randomly from one individual to another. In the absence of predictor variables, functions

✉ Jizhou Kang
jkang37@ucsc.edu

Athanasios Kottas
thanos@soe.ucsc.edu

¹ Department of Statistics, University of California, Santa Cruz, CA, USA

of the observation time are usually used as covariates. We refer to Molenberghs and Verbeke (2006) for a comprehensive review. In Sect. 2.5, we elaborate on the connection of our proposed modeling approach with existing methods.

In this article, we introduce a novel viewpoint for longitudinal binary and ordinal data analysis. We begin with the model construction for longitudinal binary responses. The critical insight that distinguishes our methodology from the majority of the existing literature is functional data analysis. We treat the subjects' measurements as stochastic process realizations at the corresponding time points. The benefits are twofold. First, the models can incorporate unbalanced data from longitudinal studies in a unified scheme; directly inferring the stochastic process provides a well-defined probabilistic model for the missing values. Secondly, we can exploit the power of Bayesian hierarchical modeling for continuous functional data (e.g., Yang et al. 2016). To that end, we adopt the Binomial distribution with the logit link that connects binary responses to continuous signals, which, subject to additive measurement error, are then modeled as (conditionally) independent and identically distributed (i.i.d.) realizations from a Gaussian process (GP) with random mean and covariance function. We place an Inverse-Wishart process (IWP) prior on the covariance function, and conditional on it, use a GP prior for the mean function. Therefore, the two essential ingredients in longitudinal modeling, the trend and the covariance structure, are modeled simultaneously and nonparametrically.

The hierarchical structure allows borrowing of strength across the subjects' trajectories. We apply a specific setting of hyperpriors for the GP and IWP priors, such that marginalizing over them, the latent continuous functions have a Student-t process (TP) prior. The TP enhances the flexibility of the GP (e.g., Shah et al. 2014). It retains attractive GP properties, such as analytic marginal and predictive distributions, and it yields predictive covariance that, unlike the GP, explicitly depends on the observed values. For inferential purposes, we represent the joint posterior distribution in multivariate form through evaluating the functions on the pooled grid, resulting in the common normal-inverse-Wishart conditional conjugacy. In conjunction with the Pólya-Gamma data augmentation technique (Polson et al. 2013), we develop a relatively simple and effective posterior simulation algorithm, circumventing the need for specialized techniques or tuning of Metropolis-Hastings steps.

To extend the model for ordinal responses, we utilize the continuation-ratio logits representation of the multinomial distribution. Such representation features an encoding of an ordinal response with C categories as a sequence of $C - 1$ binary indicators, in which the j -th indicator signifies whether the ordinal response belongs to the j -th category or to one of the higher categories. We show that fitting a multinomial model for the ordinal responses is equivalent to fitting

separately the aforementioned model on the binary indicators. Hence, we can conduct posterior simulation for each response category in a parallel fashion, leading to significant computational efficiency gains in model implementation.

In modern longitudinal studies, it is common that the complete vector of repeated measurements is not collected on all subjects. As a specific example, in ecological momentary assessment (EMA) studies, emotions and behaviors are repeatedly measured for a cohort of participants, through wearable electronic devices (Ruwaard et al. 2018). For instance, in the *StudentLife* study (Wang et al. 2014), researchers monitored the students' mental status through pop-up questionnaires on their smartphones that prompted multiple times at pseudorandom intervals during the study period. Since the data collection process is based on the participants' conscious responding to prompted questions several times a day, non-response is inevitable. Missing values are typically considered to be a nuisance rather than a characteristic of EMA time series. Parametric and nonparametric Bayesian methods have been developed to handle longitudinal data with missingness; see Daniels and Xu (2020) for a review. The common issue is that one has to bear the drawbacks of making either structured or unstructured assumptions to manage missingness. The unstructured approach leads to flexibility, yet it may result in difficulties due to a large number of parameters relative to the sample size. Besides, the majority of the existing literature on longitudinal studies with missingness focuses on the scenario with continuous responses, and the extension to discrete responses is not trivial.

Accordingly, our contributions can be summarized as follows: (i) we model the mean and covariance jointly and nonparametrically, avoiding potential biases caused by a pre-specified model structure; (ii) we unify the toolbox for balanced and unbalanced longitudinal studies; (iii) the model encourages borrowing of strength, preserving systematic patterns that are common across all subject responses; (iv) we develop a computationally efficient posterior simulation method by taking advantage of conditional conjugacy; (v) the model facilitates applications for ordinal responses with a moderate to large number of categories.

The rest of the paper is organized as follows. Section 2 develops the methodology for binary responses, including model formulation, study of model properties, and the computational approach to inference and prediction. Section 3 illustrates the modeling approach through an EMA study that focuses on analyzing students' mental health through binary outcomes. The modeling extension for longitudinal ordinal responses is presented in Sect. 4, including an illustration involving an ordinal outcome from the same EMA study. Finally, Sect. 5 concludes with discussion.

2 The modeling approach for binary responses

Here, we develop the methodology for longitudinal binary responses. The data consist of repeated binary responses on n subjects, with the observation on subject i at time τ_{it} denoted by Y_{it} . The set of repeated outcomes for the i -th subject is collected into a T_i -dimensional vector $\mathbf{Y}_i = (Y_{i1}, \dots, Y_{iT_i})^\top$. The hierarchical model construction is presented in Sect. 2.1. In Sect. 2.2, we discuss model properties related to our inference objectives. Bayesian inference and prediction is developed in Sect. 2.3. In Sect. 2.4, we outline the findings from simulation studies, the details of which are included in the Supplementary Material. Finally, to place our contribution within the literature, we discuss in Sect. 2.5 the proposed model in the context of relevant Bayesian nonparametric approaches.

Regarding notation under our functional data analysis modeling approach, we use the regular letter and its bold form to distinguish between the trajectory of responses over time and its evaluation on a number of time points. We use similar notation for other functional variables, possibly including the time input(s) inside parentheses or as a subscript. Moreover, τ and $\boldsymbol{\tau}$ denote the generic time input and a grid of times, respectively.

2.1 Model specification

We examine the data from a functional data analysis perspective, treating each observed data vector \mathbf{Y}_i as the evaluation of trajectory $Y_i(\tau)$ on grid $\boldsymbol{\tau}_i = (\tau_{i1}, \dots, \tau_{iT_i})^\top$, for $i = 1, \dots, n$. The n trajectories are assumed to be (conditionally) independent realizations from a continuous-time stochastic process. The prior probability model is built on the stochastic process. This approach avoids strong predetermined assumptions on the transition mechanism within the sequence of subject-specific responses in \mathbf{Y}_i , while it is suitable to accommodate repeated measurements regardless of their observational pattern.

The functional data analysis view of longitudinal data dates back at least to Zhao et al. (2004), where it is suggested that functional data analysis tools, such as principal component analysis, can be used to capture periodic structure in longitudinal data. Indeed, Yao et al. (2005) study functional principal component analysis (FPCA) for sparse longitudinal data, a method that can provide effective recovery of the entire individual trajectories from fragmental data. FPCA has been applied in finance (Ingrassia and Costanzo 2005), biomechanics (Donà et al. 2009), and demographic studies (Shamshoian et al. 2020). Its extension to examine sequences of discrete data is studied in Hall et al. (2008).

Our methodology builds from a GP-based hierarchical model for continuous functional data (Yang et al. 2016). Regarding mean-covariance estimation, the model in Yang et al. (2016) can be considered as a Bayesian counterpart of Yao et al. (2005). The hierarchical scheme enables a natural extension to studies with binary responses. We assume that, subject to measurement error, the i -th subject's responses, $Y_{it} \equiv Y_i(\tau_{it})$, depend on the i -th trajectory of the underlying process, evaluated at times τ_{it} , through the following model

$$Y_i(\tau_{it}) \mid Z_i(\tau_{it}), \epsilon_{it} \stackrel{ind.}{\sim} \text{Bin}(1, \varphi(Z_i(\tau_{it}) + \epsilon_{it})), \\ t = 1, \dots, T_i, \quad i = 1, \dots, n,$$

where $\varphi(x) = \exp(x)/\{1 + \exp(x)\}$ denotes the expit function. The error terms are i.i.d. from a white noise process, that is, $\epsilon_{it} \mid \sigma_\epsilon^2 \stackrel{i.i.d.}{\sim} N(0, \sigma_\epsilon^2)$, and independent of the process realizations $Z_i(\cdot)$. The main building block for the model construction is a hierarchical GP prior for the $Z_i(\cdot)$. In particular, given random mean function $\mu(\cdot)$ and covariance kernel $\Sigma(\cdot, \cdot)$, the $Z_i(\cdot)$ are i.i.d. GP realizations, denoted by $Z_i \mid \mu, \Sigma \stackrel{i.i.d.}{\sim} GP(\mu, \Sigma)$, for $i = 1, \dots, n$. The hierarchical GP prior model is completed with nonparametric priors for the mean function and covariance kernel:

$$\mu \mid \Sigma \sim GP(\mu_0, \Sigma/\kappa), \quad \Sigma \sim IWP(\nu, \Psi_\phi), \quad (1)$$

where $GP(\cdot, \cdot)$ and $IWP(\cdot, \cdot)$ denote the GP and IWP prior, respectively. The nonparametric prior reflects the intuition that parametric forms will generally not be sufficiently flexible for the mean and covariance functions.

We adopt an IWP prior for the covariance kernel, defined such that, on any finite grid $\boldsymbol{\tau} = (\tau_1, \dots, \tau_T)$ with $|\boldsymbol{\tau}|$ points, the projection $\Sigma(\boldsymbol{\tau}, \boldsymbol{\tau})$ follows an inverse-Wishart distribution with mean $\Psi_\phi(\boldsymbol{\tau}, \boldsymbol{\tau})/(\nu - 2)$, denoted by $IW(\nu, \Psi_\phi(\boldsymbol{\tau}, \boldsymbol{\tau}))$. Here, $\Psi_\phi(\cdot, \cdot)$ is a non-negative definite function with parameters ϕ . Note that we use the parameterization from Dawid (1981) for the inverse-Wishart distribution, in particular, ν is the shape parameter and $\nu + |\boldsymbol{\tau}| - 1$ is the degrees of freedom parameter in the more common parameterization. Yang et al. (2016) validate that this parameterization defines an infinite dimensional probability measure whose finite dimensional projection on grid $\boldsymbol{\tau}$ coincides with the inverse-Wishart distribution $IW(\nu, \Psi_\phi(\boldsymbol{\tau}, \boldsymbol{\tau}))$.

The model formulation is completed with prior specification for the hyperparameters. The error variance is assigned an inverse gamma prior, $\sigma_\epsilon^2 \sim IG(a_\epsilon, b_\epsilon)$. We focus primarily on stationary specifications under the prior structure in (1). In particular, we work with mean function, $\mu_0(\tau) \equiv \mu_0$, and isotropic covariance function, Ψ_ϕ , within the Matérn class, a widely used class of covariance functions (Rasmussen and Williams 2006). In general, the Matérn covariance function is specified by a scale parameter σ^2 , a range parameter ρ ,

and a smoothness parameter ι . To encourage smoothness in the probability response curves, we set $\iota = 5/2$, such that the covariance kernel is given by

$$\Psi_\phi(\tau, \tau') = \sigma^2 \left(1 + \frac{\sqrt{5}|\tau - \tau'|}{\rho} + \frac{5|\tau - \tau'|^2}{3\rho^2} \right) \times \exp\left(-\frac{\sqrt{5}|\tau - \tau'|}{\rho}\right),$$

where $\phi = \{\sigma^2, \rho\}$. For hyperparameters μ_0, σ^2, ρ , we take the commonly used choice,

$$\begin{aligned} \mu_0 &\sim N(a_\mu, b_\mu), \quad \sigma^2 \sim \text{Gamma}(a_\sigma, b_\sigma), \\ \rho &\sim \text{Unif}(a_\rho, b_\rho). \end{aligned}$$

Finally, we set $\kappa = (\nu - 3)^{-1}$, such that the continuous-time process for the $Z_i(\cdot)$ is a TP when μ and Σ are marginalized out (see Sect. 2.2 for details). As a consequence, parameter ν controls the tail heaviness of the marginal process, with smaller values of ν corresponding to heavier tails. We place a uniform prior on ν , $\nu \sim \text{Unif}(a_\nu, b_\nu)$, with $a_\nu > 3$ to ensure positive definiteness of Σ/κ .

As discussed in Diggle (1988), the correlation of repeated measurements on the same subject commonly has the following patterns. First, it should decrease with respect to the measurements' separation in time, while remaining positive to indicate the measurements are from the same subject. This feature is encapsulated by the form of the covariance kernel Ψ_ϕ . The IWP prior elicits realizations for which this property holds a priori, while enabling a flexible estimate of the covariance structure with information from the data a posteriori. Second, measurements that are made arbitrarily close in time are subject to imperfect correlation, possibly caused by subsampling of each subject. This feature is represented by the error term in our model. Moreover, the motivation for adding the error term arises from the fact that measurement error is introduced in the estimation of a continuous-time function based on data collected at discrete time points.

In addition to the aforementioned methodological considerations, adding the error term is practically important. Effectively, the error term serves as a nugget to the covariance matrix. It can alleviate numerical problems that may arise from its inversion, a calculation required in the posterior inference procedure. Moreover, adding the error term is common practice in other areas involving GP-based models, including spatial statistics (e.g., Carmack et al. 2012) and computer model emulation (e.g., Andrianakis and Challenor 2012).

Although the probability model is formulated through stochastic process realizations, posterior simulation is based on the corresponding finite dimensional distributions (f.d.d.s.). Consequently, to write the model for the data, we need to represent the likelihood and prior in multivariate forms through

evaluating the functions on finite grids. Denoting $Y_i(\tau_i)$ by \mathbf{Y}_i , $Z_i(\tau_i)$ by \mathbf{Z}_i , and $\epsilon_i = (\epsilon_{i1}, \dots, \epsilon_{iT_i})^\top$, the model for the data can be written as

$$\begin{aligned} \mathbf{Y}_i \mid \mathbf{Z}_i, \epsilon_i &\stackrel{\text{ind.}}{\sim} \prod_{t=1}^{T_i} \text{Bin}(1, \varphi(Z_{it} + \epsilon_{it})), \quad i = 1, \dots, n, \\ \mathbf{Z}_i \mid \mu(\tau_i), \Sigma(\tau_i, \tau_i) &\stackrel{\text{ind.}}{\sim} N(\mu(\tau_i), \Sigma(\tau_i, \tau_i)), \\ \epsilon_i \mid \sigma_\epsilon^2 &\stackrel{\text{ind.}}{\sim} N(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I}). \end{aligned} \tag{2}$$

Notice that the grids $\{\tau_i : i = 1, \dots, n\}$ are not necessarily the same for all subjects. Therefore, the shared GP and IWP prior in (1) need to be evaluated on the pooled grid $\tau = \cup_{i=1}^n \tau_i$. If μ , Σ , and Ψ_ϕ denote $\mu(\tau)$, $\Sigma(\tau, \tau)$, and $\Psi_\phi(\tau, \tau)$, respectively, then

$$\mu \mid \Sigma, \mu_0, \nu \sim N(\mu_0 \mathbf{1}, (\nu - 3)\Sigma), \quad \Sigma \mid \nu, \phi \sim IW(\nu, \Psi_\phi). \tag{3}$$

The hierarchical model formulation for the data in (2) and (3) forms the basis for the posterior simulation algorithm, which is discussed in detail in Sect. 2.3.

2.2 Model properties

To fix ideas for the following discussion, we refer to $Z_i(\tau)$ as the signal process of the binary process $Y_i(\tau)$, and to $Z_i(\tau) = Z_i(\tau) + \epsilon_i(\tau)$ as the latent process of $Y_i(\tau)$. Since the stochastic process is characterized by its f.d.d.s., we shall investigate the random vectors $\mathbf{Y}_\tau = Y_i(\tau)$, $\mathcal{Z}_\tau = Z_i(\tau)$, and $\mathbf{Z}_\tau = Z_i(\tau)$, for a generic grid vector $\tau = (\tau_1, \dots, \tau_T)^\top$. We surpass the subject index i because the subject trajectories are identically distributed. The Supplementary Material includes proofs for the propositions included in this section.

Among the various inference goals in a study that involves longitudinal binary data, estimating the probability response curve and the covariance structure of the repeated measurements are the most important ones. In Proposition 1, we derive the probability response curves and covariance matrix of the binary vector \mathbf{Y}_τ , conditional on the signal vector \mathbf{Z}_τ and error variance σ_ϵ^2 . The probability response curve can be defined generically as $\mathbf{P}_{\mathbf{Y}_\tau} = (\Pr(Y_{\tau_1} = y_{\tau_1} \mid \mathbf{Z}_\tau, \sigma_\epsilon^2), \dots, \Pr(Y_{\tau_T} = y_{\tau_T} \mid \mathbf{Z}_\tau, \sigma_\epsilon^2))^\top$, where y_{τ_t} is either 0 or 1. Without loss of generality, we focus on $\mathbf{P}_{\mathbf{Y}_\tau}$.

Proposition 1 *The probability response curve is given by $\mathbf{P}_{\mathbf{Y}_\tau} = E(\pi(\mathcal{Z}_\tau) \mid \mathbf{Z}_\tau, \sigma_\epsilon^2)$, where $\pi(\mathbf{x})$ denotes the vector operator that applies the expit function to every entry of \mathbf{x} . Regarding the covariance matrix, for $\tau \in \tau$, $\text{Var}(Y_\tau \mid \mathbf{Z}_\tau, \sigma_\epsilon^2) = E(\varphi(\mathcal{Z}_\tau) \mid \mathbf{Z}_\tau, \sigma_\epsilon^2) - E^2(\varphi(\mathcal{Z}_\tau) \mid \mathbf{Z}_\tau, \sigma_\epsilon^2)$, and for $\tau, \tau' \in \tau$, with $\tau' \neq \tau$, $\text{Cov}(Y_\tau, Y_{\tau'} \mid \mathbf{Z}_\tau, \sigma_\epsilon^2) =$*

$\text{Cov}(\varphi(\mathcal{Z}_\tau), \varphi(\mathcal{Z}_{\tau'}) \mid \mathbf{Z}_\tau, \sigma_\epsilon^2)$. The conditional expectations in all of the above expressions are with respect to distribution, $\mathcal{Z}_\tau \mid \mathbf{Z}_\tau, \sigma_\epsilon^2 \sim N(\mathbf{Z}_\tau, \sigma_\epsilon^2 \mathbf{I})$.

The practical utility of Proposition 1 lies on performing posterior inference for the probability response curve and the covariance structure of the binary process, conditioning on the signal process and the noise. With posterior samples of \mathbf{Z}_τ and σ_ϵ^2 , we can simulate \mathcal{Z}_τ from $N(\mathbf{Z}_\tau, \sigma_\epsilon^2 \mathbf{I})$ and numerically compute the corresponding moments in Proposition 1. The entries of \mathcal{Z}_τ are independent, given \mathbf{Z}_τ , and thus simulating \mathcal{Z}_τ is not computationally demanding, even when $|\tau|$ is large.

We next establish a closer connection between the binary process and the signal process. Proposition 2 reveals that the evolution of the binary process over time can be (approximately) expressed as a function of the expectation of the signal process and the total variance. Moreover, the covariance of the binary process is approximately the covariance of the signal process scaled by a factor related to the expectation of the signal.

Proposition 2 Consider the proposed model as described in (2) and denote $\mu(\tau) = \mu$, and $\Sigma(\tau, \tau) = \Sigma$. Then,

$$\Pr(Y_\tau = 1 \mid \mu, \Sigma, \sigma_\epsilon^2) \approx \varphi(E(Z_\tau \mid \mu, \Sigma)) \\ + \frac{\text{Var}(Z_\tau \mid \mu, \Sigma) + \sigma_\epsilon^2}{2} \varphi''(E(Z_\tau \mid \mu, \Sigma)), \quad \forall \tau \in \tau,$$

$$\begin{aligned} \text{Cov}(Y_\tau, Y_{\tau'} \mid \mu, \Sigma, \sigma_\epsilon^2) &\approx \varphi'(E(Z_\tau \mid \mu, \Sigma)) \\ &\times \varphi'(E(Z_{\tau'} \mid \mu, \Sigma)) \text{Cov}(Z_\tau, Z_{\tau'} \mid \mu, \Sigma) \\ &- \frac{1}{4} [\text{Var}(Z_\tau \mid \mu, \Sigma) + \sigma_\epsilon^2] [\text{Var}(Z_{\tau'} \mid \mu, \Sigma) + \sigma_\epsilon^2] \\ &\times \varphi''(E(Z_\tau \mid \mu, \Sigma)) \varphi''(E(Z_{\tau'} \mid \mu, \Sigma)), \quad \forall \tau, \tau' \in \tau. \end{aligned}$$

Here, $\varphi'(x) = \frac{d\varphi(x)}{dx} = \varphi(x)[1 - \varphi(x)]$ and $\varphi''(x) = \frac{d^2\varphi(x)}{dx^2} = \varphi(x)[1 - \varphi(x)][1 - 2\varphi(x)]$.

Our inference results are based on exact expressions, such as the ones in Proposition 1. Nonetheless, the approximate expressions derived in Proposition 2 are practically useful to gain more insight on properties of the binary process, as well as for prior specification. Note that exploring properties of the binary process is not trivial due to the lack of general analytical forms for moments of logit-normal distributions. Hence, a connection with properties of the signal process is useful. For instance, if we specify the covariance for the signal process to decrease as a function of separation in time, an analogous structure will hold (approximately) for the binary process.

The previous discussion focuses on studying the f.d.d.s of the binary process given the signal process. Therefore, it is important to investigate the marginal f.d.d.s of the signal

process. We show that, under the specification $\kappa = (\nu - 3)^{-1}$, the f.d.d.s. of the signal process correspond to a multivariate Student-t (MVT) distribution, and thus the signal process is a TP. We first state the definition of the MVT distribution and the TP (see, e.g., Shah et al. 2014). Notice that we use the covariance matrix as a parameter for the MVT distribution, instead of the more common parameterization based on a scale matrix.

Definition 1 The random vector $\mathbf{Z} \in \mathbb{R}^n$ is MVT distributed, denoted $\mathbf{Z} \sim MVT(\nu, \mu, \Psi)$, if it has density

$$\frac{\Gamma(\frac{\nu+n}{2})}{[(\nu-2)\pi]^{n/2}\Gamma(\frac{\nu}{2})} |\Psi|^{-1/2} \left(1 + \frac{(\mathbf{Z} - \mu)^T \Psi^{-1} (\mathbf{Z} - \mu)}{\nu - 2} \right)^{-\frac{\nu+n}{2}}$$

where $\nu > 2$ is the degrees of freedom parameter, $\mu \in \mathbb{R}^n$, and Ψ is an $n \times n$ symmetric, positive definite matrix. Under this parameterization, $E(\mathbf{Z}) = \mu$ and $\text{Cov}(\mathbf{Z}) = \Psi$.

Consider a process $Z(\tau)$ formulated through mean function $\mu(\tau)$, a non-negative kernel function $\Psi(\tau, \tau)$, and parameter $\nu > 2$, such that its f.d.d.s correspond to the MVT distribution with mean vector and covariance matrix induced by $\mu(\tau)$ and $\Psi(\tau, \tau)$, respectively. Then, $Z(\tau)$ follows a TP, denoted by $Z(\tau) \sim TP(\nu, \mu(\tau), \Psi(\tau, \tau))$.

Marginalizing over μ and Σ in (2) and (3), the implied distribution for \mathbf{Z}_τ is MVT, with degrees of freedom parameter ν (with $\nu > 3$ in our context), mean vector $\mu_0 \mathbf{1}$, and covariance matrix $\Psi_\phi = \Psi_\phi(\tau, \tau)$. We thus obtain the following result for the signal process.

Proposition 3 Under the model formulation in (2) and (3), the signal process follows marginally a TP, that is, $Z \sim TP(\nu, \mu_0, \Psi_\phi)$.

Proposition 3 is beneficial in terms of both computation and interpretation. Without a constraint on κ , as in Yang et al. (2016), the marginal distribution of \mathbf{Z}_τ does not have analytical form. Hence, for prediction at new time points, one has to sample from an IWP and a GP, which is computationally intensive, especially for a dense grid. In contrast, we can utilize the analytical form of the TP predictive distribution to develop a predictive inference scheme that resembles that of GP-based models (see Sect. 2.3). Moreover, the result highlights the model property that the degrees of freedom parameter ν controls how heavy tailed the process is. Smaller values of ν correspond to heavier tails. As ν gets larger, the tails resemble Gaussian tails. Additionally, ν controls the dependence between Z_τ and $Z_{\tau'}$, which are jointly MVT distributed, with smaller values indicating higher dependence. Such interpretation of parameter ν facilitates the choice of its hyperprior.

The local behavior of stochastic process realizations is crucial for interpolation. Under the longitudinal setting, continuous, or perhaps differentiable, signal process trajectories are typically anticipated. Evidently, the observed data can not visually inform the smoothness of signal process realizations. Rather, such smoothness should be captured in the prior specification that incorporates information about the data generating mechanism. For weakly stationary processes, mean square continuity is equivalent to the covariance function being continuous at the origin (Stein 1999). And, the process is ι -times mean square differentiable if and only if the 2ι -times derivative of the covariance function at the origin exists and is finite. Under our model, the signal process follows a TP marginally. Its covariance structure is specified by the Matérn covariance function with smoothness parameter ι . Referring to the behavior of the Matérn class of covariance functions at the origin, we obtain the following result for the mean square continuity and differentiability of the signal process.

Proposition 4 Consider the proposed model with marginal signal process $Z \sim TP(v, \mu_0, \Psi_\phi)$, where Ψ_ϕ belongs to the Matérn family of covariance functions with smoothness parameter ι . Then, the signal process is mean square continuous and $\lfloor \iota \rfloor$ -times mean square differentiable.

The results in this section study several properties that are useful in model implementation. Indeed, the practical utility of such model properties with respect to prior specification and posterior inference is discussed in the next section.

2.3 Prior specification and posterior inference

The model described in Sect. 2.1 contains parameters $\{\sigma_\epsilon^2, \mu_0, \sigma^2, \rho, v\}$ whose prior hyperparameters need to be specified. We develop a default specification strategy that relies on the model properties explored in Sect. 2.2.

First, we set the prior for μ_0 such that the prior expected probability response curve does not favor any category, and the corresponding prior uncertainty bands span a significant portion of the unit interval. For instance, this can be achieved with prior $\mu_0 \sim N(0, 100)$ which yields prior expected probability of positive response of about 1/2 across τ . In general, we would not expect to have available prior information about the variance and correlation structure of the unobserved signal process, which are controlled by parameters σ^2 and ρ . However, Proposition 2 suggests an approximate relationship between the covariance structure of the binary process and the signal process, and we can thus specify the corresponding priors similarly to GP-based models. In particular, we select the uniform prior for the range parameter ρ such that the correlation between Z_τ and $Z_{\tau'}$ decreases to 0.05 when the difference between τ and τ' is within a pre-specified subset of

the observation time window. For instance, for the data analysis in Sect. 3 where the total observation window comprises 72 days, we used a $Unif(3, 12)$ prior for ρ , which implies that the aforementioned correlation decreases to 0.05 when the time difference ranges from 7 to 31 days. The hyperprior for v is $Unif(a_v, b_v)$. We specify $a_v > 3$ to reflect the constraint for $\Sigma/(v - 3)$ to be a well-defined covariance matrix, and b_v large enough such that the tail behavior of the marginal TP is hard to distinguish from that of a GP. For instance, a default choice is $a_v = 4$ and $b_v = 30$.

We follow Fong et al. (2010) to specify the prior for $\sigma_\epsilon^2 \sim IG(a_\epsilon, b_\epsilon)$. Integrating out σ_ϵ^2 , the measurement error ϵ is marginally distributed as a univariate Student-t distribution with location parameter 0, scale parameter b_ϵ/a_ϵ , and degrees of freedom parameter $2a_\epsilon$. For a predetermined measurement error range $(-R, R)$ with degree of freedom v , we can use the relationship $\pm t_{1-(1-q)/2}^v \sqrt{b_\epsilon/a_\epsilon} = \pm R$ to obtain $a_\epsilon = v/2$ and $b_\epsilon = R^2 v / [2(t_{1-(1-q)/2}^v)^2]$, where t_q^v is the q -th percentile of a Student-t distribution with v degrees of freedom.

Proceeding to posterior inference, we develop an MCMC algorithm based on (2) and (3). We introduce layers of latent variables, beginning with $\xi_{it} \sim PG(1, 0)$ for every observation Y_{it} , where $PG(a, b)$ denotes the Pólya-Gamma distribution with shape parameter a and tilting parameter b (Polson et al. 2013). Denote the collection of Pólya-Gamma variables for each subject by $\xi_i = (\xi_{i1}, \dots, \xi_{iT_i})^\top$. Also, introduce $Z_{it} = Z_{it} + \epsilon_{it}$, and let $Z_i = (Z_{i1}, \dots, Z_{iT_i})^\top$. Recall that $\tau = \cup_{i=1}^n \tau_i$ is the pooled grid. Denote the evaluations on the pooled grid by $\tilde{Z}_i = Z_i(\tau)$ and let $Z_i^* = \tilde{Z}_i \setminus Z_i$. That is, $Z_i^* = Z_i(\tau_i^*)$, where $\tau_i^* = \tau \setminus \tau_i$ is the set of grid points at which the i -th trajectory misses observations. Then, the hierarchical model for the data $\{Y_{it} : t = 1, \dots, T_i, i = 1, \dots, n\}$ can be expressed as

$$\begin{aligned} Y_{it} | Z_{it}, \xi_{it} &\stackrel{ind.}{\sim} \mathcal{B}(Z_{it}, \xi_{it}), \\ \xi_{it} &\stackrel{i.i.d.}{\sim} PG(1, 0), \quad t = 1, \dots, T_i, \\ Z_i | Z_i, \sigma_\epsilon^2 &\stackrel{ind.}{\sim} N(Z_i, \sigma_\epsilon^2 \mathbf{I}_{T_i}), \\ \tilde{Z}_i = (Z_i, Z_i^*)^\top | \mu, \Sigma &\stackrel{i.i.d.}{\sim} N(\mu, \Sigma), \quad i = 1, \dots, n, \\ \sigma_\epsilon^2 &\sim IG(a_\epsilon, b_\epsilon), \quad \mu | \mu_0, \Sigma, v \sim N(\mu_0 \mathbf{1}, (v - 3)\Sigma), \\ \mu_0 &\sim N(a_\mu, b_\mu), \quad \Sigma | v, \Psi_\phi \sim IW(v, \Psi_\phi), \\ \Psi_\phi &= \Psi_\phi(\tau, \tau), \quad \phi = \{\sigma^2, \rho\}, \quad \sigma^2 \sim \text{Gamma}(a_\sigma, b_\sigma), \\ \rho &\sim Unif(a_\rho, b_\rho), \quad v \sim Unif(a_v, b_v). \end{aligned}$$

Here, $\mathcal{B}(Z_{it}, \xi_{it}) \propto \exp\{(Y_{it} - 0.5) Z_{it} - 0.5 \xi_{it} Z_{it}^2\}$ denotes the probability mass function of Y_{it} conditional on both sets of latent variables, Z_{it} and ξ_{it} . Based on Theorem 1 from Polson et al. (2013), marginalizing out ξ_{it} from $\mathcal{B}(Z_{it}, \xi_{it})$, we obtain the Binomial distribution for Y_{it} conditional on

latent variable \mathcal{Z}_{it} , that is, $Y_{it} \mid \mathcal{Z}_{it} \stackrel{ind.}{\sim} \text{Bin}(1, \varphi(\mathcal{Z}_{it}))$. Hence, the joint posterior density of all model parameters can be written as

$$\begin{aligned} & p(\{\mathcal{Z}_i\}_{i=1}^n, \{\xi_i\}_{i=1}^n, \{\tilde{\mathbf{Z}}_i\}_{i=1}^n, \boldsymbol{\mu}, \boldsymbol{\Sigma}, \sigma_\epsilon^2, \\ & \mu_0, \sigma^2, \rho, \nu \mid \{\mathbf{Y}_i\}_{i=1}^n) \\ & \propto \prod_{i=1}^n \{p(\mathbf{Y}_i \mid \mathcal{Z}_i, \xi_i)p(\xi_i)p(\mathcal{Z}_i \mid \mathbf{Z}_i, \sigma_\epsilon^2) \\ & p(\mathbf{Z}_i^* \mid \mathbf{Z}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})p(\mathbf{Z}_i \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})\} \\ & \times p(\boldsymbol{\mu} \mid \mu_0, \boldsymbol{\Sigma}, \nu)p(\boldsymbol{\Sigma} \mid \sigma^2, \rho, \nu)p(\sigma_\epsilon^2)p(\mu_0) \\ & p(\sigma^2)p(\rho)p(\nu). \end{aligned} \quad (4)$$

The introduction of the latent variables enables a Gibbs sampling scheme with conditionally conjugate updates. Denote generically by $p(\boldsymbol{\theta} \mid \cdot)$ the posterior full conditional for parameter $\boldsymbol{\theta}$. Notice that $p(\mathcal{Z}_i, \xi_i \mid \cdot) \propto p(\mathbf{Y}_i \mid \mathcal{Z}_i, \xi_i)p(\xi_i)p(\mathcal{Z}_i \mid \mathbf{Z}_i, \sigma_\epsilon^2)$, which matches the Bayesian logistic regression structure in Polson et al. (2013). Therefore, $p(\mathcal{Z}_i \mid \cdot)$ and $p(\xi_i \mid \cdot)$ can be sampled directly. Factorizing the prior of $\tilde{\mathbf{Z}}_i$ as $p(\tilde{\mathbf{Z}}_i \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) = p(\mathbf{Z}_i^* \mid \mathbf{Z}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})p(\mathbf{Z}_i \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})$, results in $p(\mathbf{Z}_i^*, \mathbf{Z}_i \mid \cdot) \propto p(\mathbf{Z}_i^* \mid \mathbf{Z}_i, \boldsymbol{\mu}, \boldsymbol{\Sigma})p(\mathbf{Z}_i \mid \boldsymbol{\mu}, \boldsymbol{\Sigma})p(\mathcal{Z}_i \mid \mathbf{Z}_i, \sigma_\epsilon^2)$. This forms yields ready updates for \mathbf{Z}_i^* and \mathbf{Z}_i using GP-based predictive sampling. All other model parameters can be sampled using standard updates. The details of the MCMC algorithm are given in the Supplementary Material.

We have linked the probability response curve and covariance structure of the binary process $Y_i(\tau)$ to the corresponding signal process $Z_i(\tau)$. To estimate the signal process, we obtain posterior samples for $\mathbf{Z}_i^+ = Z_i(\tau^+)$, where $\tau^+ \supset \tau$ is a finer grid than the pooled grid. Denote $\check{\tau} = \tau^+ \setminus \tau$ as the time points where none of the subjects have observations, and let $\check{\mathbf{Z}}_i = Z_i(\check{\tau})$. Using the marginal TP result from Proposition 3,

$$\begin{pmatrix} \check{\mathbf{Z}}_i \\ \check{\mathbf{Z}}_i \end{pmatrix} \sim MVT \left(\nu, \begin{pmatrix} \mu_{0\tau} \\ \mu_{0\check{\tau}} \end{pmatrix}, \begin{pmatrix} \Psi_{\tau, \tau} & \Psi_{\tau, \check{\tau}} \\ \Psi_{\check{\tau}, \tau} & \Psi_{\check{\tau}, \check{\tau}} \end{pmatrix} \right),$$

where $\mu_{0\cdot} = \mu_0 \mathbf{1}_{|\cdot|}$, and $\Psi_{\cdot, \cdot}$ denotes the covariance function evaluation $\Psi_\phi(\cdot, \cdot)$. Next, based on the conditionals of the MVT distribution (Shah et al. 2014),

$$\check{\mathbf{Z}}_i \mid \check{\mathbf{Z}}_i \sim MVT \left(\nu + |\tau|, \check{\mu}_{i\check{\tau}}, \frac{\nu + S_{i\tau} - 2}{\nu + |\tau| - 2} \check{\Psi}_{\check{\tau}, \check{\tau}} \right), \quad (5)$$

with $\check{\mu}_{i\check{\tau}} = \Psi_{\check{\tau}, \tau} \Psi_{\tau, \tau}^{-1} (\check{\mathbf{Z}}_i - \mu_{0\tau}) + \mu_{0\check{\tau}}$, $S_{i\tau} = (\check{\mathbf{Z}}_i - \mu_{0\tau})^\top \Psi_{\tau, \tau}^{-1} (\check{\mathbf{Z}}_i - \mu_{0\tau})$ and $\check{\Psi}_{\check{\tau}, \check{\tau}} = \Psi_{\check{\tau}, \check{\tau}} - \Psi_{\check{\tau}, \tau} \Psi_{\tau, \tau}^{-1} \Psi_{\tau, \check{\tau}}$. Using (5), given each posterior sample for $\check{\mathbf{Z}}_i$, μ_0 , ϕ and ν , we can complete the posterior realization for the signal process over the finer grid. As discussed in Sect. 2.2, we can then

obtain full posterior inference for functionals of the binary process.

The predictive distribution of the signal process also illustrates the information borrowed across subjects. For the i -th subject, the grid, τ^+ , where predictions are made can be partitioned as $\tau_i \cup \tau_i^* \cup \check{\tau}$, where $\tau_i^* = \tau \setminus \tau_i$ represents the grid points where subject i does not have observations, while at least one of the other subjects have observations. Then, we first predict $Z_i(\tau_i^*)$ conditioning on $Z_i(\tau_i)$ by the GP predictive distribution, and next predict $Z_i(\check{\tau})$ conditioning on $Z_i(\tau_i)$ and $Z_i(\tau_i^*)$ by the TP predictive distribution. Comparing with the GP, (5) suggests the TP is scaling the predictive covariance by the factor $\frac{\nu + S_{i\tau} - 2}{\nu + |\tau| - 2}$. Note that $S_{i\tau}$ is distributed as the sum of squares of $|\tau|$ independent $MVT_1(\nu, 0, 1)$ random variables and hence $E(S_{i\tau}) = |\tau|$. Accordingly, if we have made good interpolation prediction, the predictive covariance for extrapolation of $Z_i(\check{\tau})$ is expected to scale down and vice versa. Comparing with predicting both $Z_i(\tau_i^*)$ and $Z_i(\check{\tau})$ conditioning on $Z_i(\tau_i)$ through the GP predictive distribution, our model allows using information across subjects to adjust the individual trajectory's credible interval.

Another crucial benefit of modeling the signal process as a TP emerges when we consider making predictions at $\check{\tau}$, the grid points where none of the subjects have observations. Under the hierarchical GP prior in Yang et al. (2016), for which the marginal is not generally a TP, such predictions would require the conditional distribution $\boldsymbol{\Sigma}_{\check{\tau}, \check{\tau}} \mid \boldsymbol{\Sigma}_{\tau, \tau}$ from their joint inverse-Wishart distribution, which is not analytically available. We circumvent this issue by marginalizing out $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$. The predictions are then based on the conditional $\check{\mathbf{Z}}_i \mid \check{\mathbf{Z}}_i$ from their joint multivariate t distribution, which is the MVT distribution in (5). Hence, for prediction on a grid denser than the pooled grid τ , the marginal TP specification for the signal process is a practically important model feature.

2.4 Synthetic data examples

We assess the model by applying it to carefully designed simulation scenarios that reflect our main contributions. The full details are provided in the Supplementary Material. Here, we briefly discuss the simulation study setting and summarize the main findings.

For the three sets of simulation studies we considered, the longitudinal binary responses are generated from the following generic process:

$$\begin{aligned} & Y_i(\tau_i) \mid Z_i(\tau_i) \stackrel{ind.}{\sim} \text{Bin}(1, \eta(Z_i(\tau_i))), \\ & \tau_i = (\tau_{i1}, \dots, \tau_{iT_i}), \quad i = 1, \dots, n, \\ & Z_i(\tau_i) = f(\tau_i) + \omega_i + \epsilon_i \quad \epsilon_i \stackrel{i.i.d.}{\sim} N(\mathbf{0}, \sigma_\epsilon^2 \mathbf{I}), \end{aligned} \quad (6)$$

where $\eta(\cdot)$ is a link function mapping \mathbb{R} to $(0, 1)$, $f(\tau)$ is a signal function, and ω_i is a realization from a mean zero continuous stochastic process that depicts the temporal covariance within the i -th subject.

The first set of simulation studies focuses on evaluating the effectiveness of the proposed model in capturing the fluctuation of the temporal trend. We consider different link function, signal function, and temporal covariance structure combinations, and we simulate unbalanced data with different sparsity levels. The results demonstrate that, despite the data generating process and the sparsity level, the model can recover not only the subject's probability response curve, but also the underlying continuous signal function.

The objective of the second set of simulation studies is to explore the performance of the proposed model in estimating the within subject covariance structure. To this end, we examine a number of possible choices for generating the ω_i in (6), which imply covariance structures that are not of the same form as the covariance kernel of the model. The results reveal that the model can recover the true covariance between the signal variables, $(Z_i(\tau_{it}), Z_i(\tau_{it}'))$, and the binary responses, $(Y_i(\tau_{it}), Y_i(\tau_{it}'))$, thus providing empirical evidence for the robustness of the covariance kernel choice.

In both cases, we examine simplified versions of the model for comparison. The simplified models are constructed by modeling either the mean structure or the covariance structure parametrically in the two sets of simulation studies, respectively. Demonstrating that the proposed model outperforms its parametric backbones, we highlight the practical utility of the nonparametric modeling for the mean and covariance structure.

To further illustrate the practical benefits of the functional data analysis perspective, in the third simulation study, we consider a scenario where observations are made irregularly. Through formal model comparison, we show that the proposed model outperforms a traditional approach under the GLMM framework.

2.5 Connections with existing literature

Our methodology is broadly related with certain Bayesian nonparametric methods. The proposed model is related to a particular class of conditional models, known as transition models, which induce the aging effect by allowing past values to explicitly affect the present observation, usually through autoregressive dynamics. Di Lucca et al. (2013) studied a class of non-Gaussian autoregression models for continuous responses, which can be extended to handle binary longitudinal outcomes by treating them as a discretized version of the continuous outcomes. DeYoreo and Kottas (2018) developed a nonparametric density regression model for ordinal regression relationships that evolve in discrete time. Compared with the proposed methodology, these models are more

flexible in terms of the binary response distribution. However, it is demanding to handle higher than first-order dynamics, and there is no natural way to treat missing data under a discrete time autoregressive framework, hindering applications for unbalanced longitudinal studies.

The proposed model is more closely related to subject-specific models, where the responses are assumed to be independent conditioning on subject-specific effects. The main approach has been to construct models for longitudinal binary responses building from Bayesian nonparametric models for longitudinal continuous data, developed under the mixed effects framework, utilizing Dirichlet process mixture models (e.g., Li et al. 2010; Ghosh and Hanson 2010; Quintana et al. 2016) or additive GP models (e.g., Cheng et al. 2019). For instance, embedding a Dirichlet process mixture of normals prior as the probability model for the latent variables, Jara et al. (2007) and Tang and Duan (2012) consider binary responses, and Kunihama et al. (2019) handle mixed-scale data comprising continuous and binary responses. The proposed model differs in the way of treating subject-specific effects, and it arguably offers benefits in terms of computational efficiency.

There is a growing trend of adopting functional data analysis tools in longitudinal data modeling. These methods specify observations as linear combinations of functional principal components (FPCs), with the FPCs represented as expansions of a pre-specified basis. Bayesian methods include Jiang et al. (2020) for continuous responses, and Van Der Linde (2009) for binary and count responses. Challenges include inference which is sensitive to the basis choice, and a complex orthogonality constraint on the FPCs. Recently, Matuk et al. (2022) proposed an approach that can serve as foundation for generalized FPC analysis of sparse and irregular binary responses. Nonetheless, our model involves a more parsimonious formulation, including the structure with the GP and TP predictive distributions.

3 Application with binary responses: *Studentlife* data

3.1 Data for analysis

Studentlife (Wang et al. 2014) is a study that integrates automatic sensing data and an EMA component to probe students' mental health status and to study its relationship with students' academic performance and behavior trends. The data were collected by a smartphone app carried by 48 students over a 10-week term at Dartmouth College. The dataset is available from the R package "studentlife" (Fryer et al. 2022).

We focus on a subset of the data that corresponds to assessing the students' emotional status. In the *Studentlife* study, the assessment of emotion is conducted by the Photographic

Affect Meter (PAM), a tool for measuring affect in which users select from a wide variety of photos the one which best suits their current mood (Pollak et al. 2011). The PAM survey is deployed to the mobile app and prompts everyday during the study period. The participants either respond to the survey, or ignore it, introducing missingness. The outcome of the survey contains two attributes, the PAM valence and the PAM arousal. They are scores of -2 to 2 (excluding 0) that measure the subject's extent of displeasure to pleasure or state of activation ranging from low to high, respectively. We dichotomize the valence and arousal scores by their sign, representing the positive values by 1. In this section, we focus on analyzing the change of binary valence and arousal responses to evaluate students' affects as the term progresses.

The data were collected during the spring 2013 term at Dartmouth college. We set the study period according to the official academic calendar, from the first day of classes (March 25, 2013) to the end of the final exam period (June 4, 2013), resulting in a total of 72 days. We exclude subjects with less than 12 responses, resulting in 45 students. The longitudinal recordings of valence or arousal of the i -th student are denoted by $Y_i(\tau_i)$, for $i = 1, \dots, 45$, where the student-specific grid points are a subset of $\tau = (0, 1, \dots, 71)^\top$, representing the days on which the measurements are recorded. Several special events occurred during the study period, and we are particularly interested in investigating the change of students' affects on the time intervals around these events. Specifically, the events and corresponding periods are: (i) Days following the Boston marathon bombing (April 15, 2013 to April 17, 2013); (ii) The Green Key (a spring festival at Dartmouth) period (May 17, 2013 to May 18, 2013); (iii) The Memorial Day long weekend (May 25, 2013 to May 27, 2013); (iv) The final examination period (May 31, 2013 to June 3, 2013).

We retrieve the data for the specific responses and study period from the R package “studentlife” that contains the database for the entire study. Over all observations, the percentage of missing values is 31.1%. There are slightly more missing responses at the beginning and toward the end of the study. In light of the structure of EMA studies, the missing pattern for each subject can be viewed as random. We elaborate on the missing-at-random assumption in the Supplementary Material.

We further explore the correlations between the binary responses within a week. We split the whole observation sequence into batches representing a week, and empirically calculate the Pearson and the tetrachoric correlation coefficient for each pair of time and distance combinations. The results, plotted in Fig. 1, suggest that the correlation of the students' response to valence and arousal decreases slowly in time.

3.2 Analysis and results

We fit the proposed model for the binary valence and arousal responses separately. We specify the prior for the model parameters by the procedure mentioned in Sect. 2.3. (Results from prior sensitivity analysis are presented in the Supplementary Material.) Posterior inference results are based on 5000 MCMC samples obtained every 4 iterations from a chain of 50000 iterations with a 30000 burn-in period (which is conservative).

We first examine in Fig. 2 the probability response curves, defined as the probability of obtaining positive valence or arousal as a function of time. For the valence, the happiness level drops as the term begins and increases when the term ends. The Boston marathon bombing may have had a minor effect on the valence. We observe local peaks around the Green Key festival and the Memorial Day holiday. As the students finish their exams, there is a trend toward happiness. As for arousal, it is relatively stable at the beginning of the term, and fluctuates as the term progresses. There is a drop in activation level after the Boston marathon bombing and during the final exam period, while the activation level reaches a local maximum at around the Green Key festival and the Memorial Day holiday.

Moreover, we assess the student's emotional status on specific days. According to Russell (1980), various states of emotional status can be represented by points located at the two dimensional mood coordinate space spanned by valence for the horizontal dimension and arousal for the vertical dimension. Moods such as excitement, distress, depression, and contentment, are represented by points in the quadrants of the space. For each observation, we can map the corresponding pairs of probabilities for positive valence and arousal onto the unit square in the mood space. In Fig. 3, the density heatmap is obtained by the posterior samples of positive probabilities for a new student of the same cohort, while the posterior means of the in-sample positive probabilities are marked by crosses. Panels (a) and (b) suggest the students are mostly excited at the festival and holiday. Moving from panel (c) to panel (d), we observe that the happiness level increases and the activation level decreases towards the end of the exam period.

We also obtain the posterior point and 95% interval estimate for the covariance kernel of the signal process, which is displayed in Fig. 4. It is noteworthy that there is a similar decreasing trend for the two distinct binary responses of valence and arousal. The practical range, defined as the distance at which the correlation is 0.05, has an estimated mean of 20.99 for valence and 22.97 for arousal.

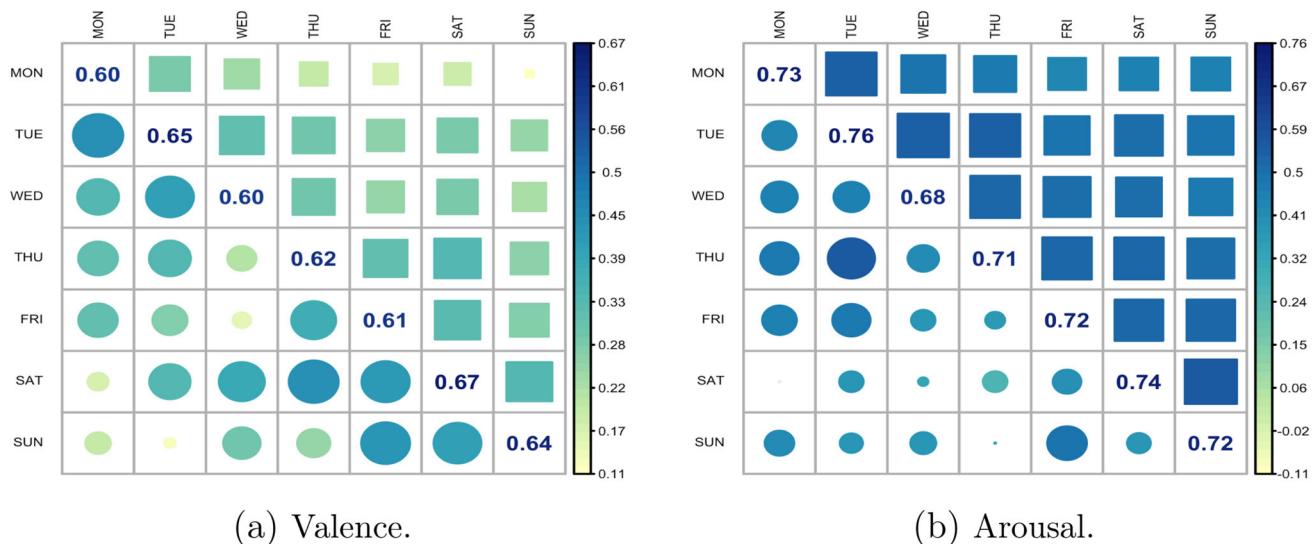


Fig. 1 *Studentlife* data. Empirical estimate of the correlation coefficients between binary responses within a week. In each panel, the upper triangle and the lower triangle are for the Pearson and the tetrachoric correlation coefficient, respectively

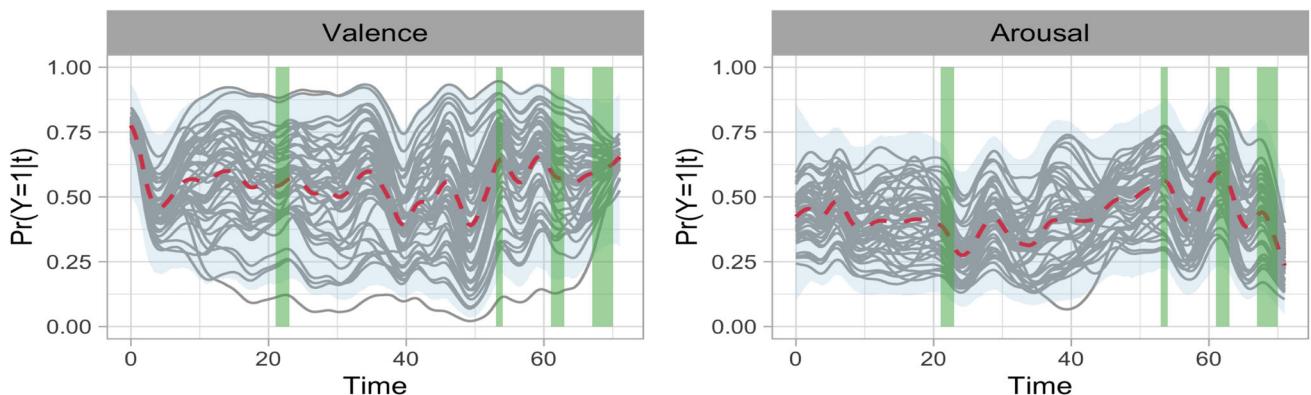


Fig. 2 *Studentlife* data. Posterior mean (dashed line) and 95% interval estimate (shaded region) of the probability response curve for an out-of-sample subject. The posterior mean estimates of probability response

curves for in-sample subjects are given by the solid lines. The vertical shaded regions correspond to the four special time periods (see Sect. 3.1)

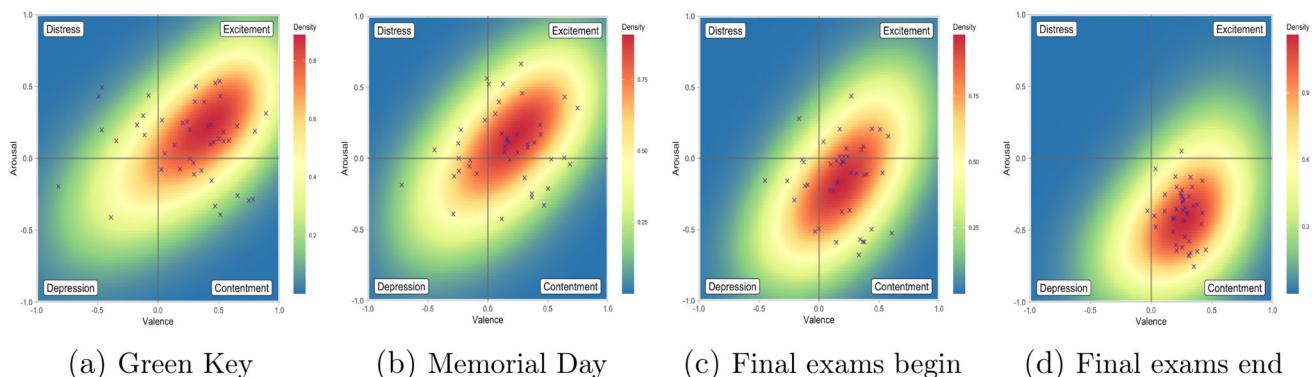


Fig. 3 *Studentlife* data. Posterior density estimate of an out-of-sample subject's valence and arousal probability over the mood coordinate space on four specific days. In each panel, the crosses represent the

posterior means of the in-sample subjects' valence and arousal probability mapped to the mood coordinate space

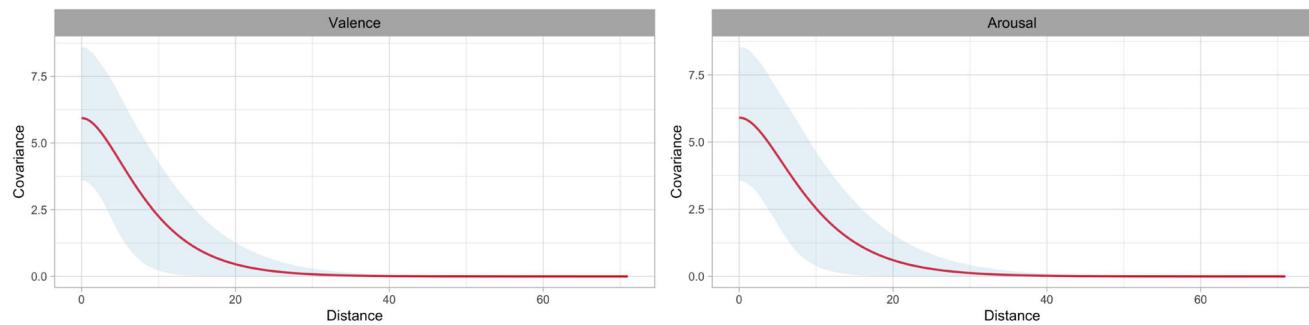


Fig. 4 *Studentlife* data. Posterior mean (solid line) and 95% interval estimate of the signal process covariance kernel

3.3 Performance comparisons

For comparison with a traditional approach, we consider an analysis of the data under the GLMM setting. In particular, we assume the model

$$Y_{it} \mid \mathcal{Z}_{it} \stackrel{\text{ind.}}{\sim} \text{Bin}(1, \varphi(\mathcal{Z}_{it})),$$

$$\mathcal{Z}_{it} = \tilde{\tau}_{it}^\top \beta + \sum_{k=1}^K S_{itk} b_k + \mu_i + \epsilon_{it},$$

$$t = 1, \dots, T_i, \quad i = 1, \dots, n,$$

where $\tilde{\tau}_{it} = (1, \tau_{it})^\top$, β is the vector of fixed effects, and $\epsilon_{it} \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_\epsilon^2)$ is the measurement error. To allow flexibility in modeling the time effect, we consider cubic B-spline basis functions with $K = 9$ knots that separate naturally the observed interval by week; S_{itk} is the k -th basis associated with time, with parameter $b_k \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_b^2)$. Finally, $\mu_i \stackrel{\text{i.i.d.}}{\sim} N(0, \sigma_\mu^2)$ are subject-specific random effects. The model is implemented using the integrated nested Laplace approximation (INLA) approach (Rue et al. 2009) with the “INLA” package in R (Rue et al. 2017). We used the default choices provided by the R package for the prior on β (a flat prior), and for the values of the variance terms, σ_ϵ^2 , σ_b^2 , and σ_μ^2 .

We perform model comparison using two different metrics: the posterior predictive loss criterion which combines a goodness-of-fit term, $G(\mathcal{M})$, and a penalty term, $P(\mathcal{M})$, for model complexity (Gelfand and Ghosh 1998); and, the continuous ranked probability score (CRPS), defined in terms of predictive cumulative distribution functions (Gneiting and

Raftery 2007). Both criteria can be calculated from the posterior samples for model parameters, and both favor the model with a smaller value. Table 1 summarizes the results. For the valence response, both criteria favor the proposed model. As for the arousal response, the proposed model provides a more accurate fit to the data, while being penalized more than the GLMM with respect to model complexity. Nonetheless, our model is favored in terms of total posterior predictive loss, as well as by the CRPS criterion.

4 Model for ordinal responses

4.1 The extended model

We extend the model developed in Sect. 2.1 to handle ordinal responses. Suppose the observation on subject i at time τ_{it} , denoted by Y_{it} , takes C possible categories. We can equivalently encode the response as a vector with binary entries $\mathbf{Y}_{it} = (Y_{i1t}, \dots, Y_{iCt})$, such that $Y_{it} = j$ is equivalent to $Y_{ijt} = 1$ and $Y_{ikt} = 0$ for any $k \neq j$. We assume a multinomial response distribution for \mathbf{Y}_{it} , factorized in terms of binomial distributions,

$$\begin{aligned} & \text{Mult}(\mathbf{Y}_{it} \mid m_{it}, \omega_{i1t}, \dots, \omega_{iCt}) \\ &= \prod_{j=1}^{C-1} \text{Bin}(Y_{ijt} \mid m_{ijt}, \varphi(Z_{ijt} + \epsilon_{ijt})) \end{aligned} \quad (7)$$

Table 1 *Studentlife* data. Summary of comparison between the proposed model and the generalized linear mixed effects model (GLMM) using two different criteria

	Response	Model	Posterior predictive loss $G(\mathcal{M})$	$P(\mathcal{M})$	$G(\mathcal{M}) + P(\mathcal{M})$	CRPS
Valence	Proposed	428.09	475.31	903.40	0.19	
		456.09	475.83	931.92	0.20	
Arousal	Proposed	457.62	496.63	954.25	0.20	
		476.17	492.28	968.45	0.21	

The values in bold correspond to the model favored by the particular criterion.

where $m_{it} = \sum_{j=1}^C Y_{ijt} \equiv 1$, $m_{i1t} = m_{it}$, and $m_{ijt} = m_{it} - \sum_{k=1}^{j-1} Y_{ikt}$. This factorization bridges the gap between binary and ordinal responses. Similar to the model for binary responses, we adopt a functional data analysis perspective on $\{Z_{ijt}\}$, modeling them separately through the hierarchical framework developed in Sect. 2.1. That is, $Z_{ij}(\tau) \mid \mu_j, \Sigma_j \stackrel{i.i.d.}{\sim} GP(\mu_j, \Sigma_j)$, for $i = 1, \dots, n$, and $\mu_j \mid \Sigma_j \stackrel{ind.}{\sim} GP(\mu_{0j}, (v_j - 3)\Sigma_j)$, $\Sigma_j \stackrel{ind.}{\sim} IW(\nu_j, \Psi_{\phi_j})$, where $\phi_j = \{\sigma_j^2, \rho_j\}$, for $j = 1, \dots, C - 1$. The error terms are modeled as $\epsilon_{ijt} \mid \sigma_{\epsilon_j}^2 \stackrel{ind.}{\sim} N(0, \sigma_{\epsilon_j}^2)$. Hence, the hierarchical model for the data can be expressed as

$$\begin{aligned} \mathbf{Y}_i | \{\mathbf{Z}_{ij}\}, \{\epsilon_{ij}\} &\stackrel{ind.}{\sim} \prod_{t=1}^{T_i} \prod_{j=1}^{C-1} \text{Bin}(Y_{ijt} \mid m_{ijt}, \\ &\quad \varphi(Z_{ijt} + \epsilon_{ijt})), \quad i = 1, \dots, n, \\ \mathbf{Z}_{ij} \mid \mu_j(\tau_i), \Sigma_j(\tau_i, \tau_i) &\stackrel{ind.}{\sim} N(\mu_j(\tau_i), \Sigma_j(\tau_i, \tau_i)), \quad (8) \\ \epsilon_{ij} \mid \sigma_{\epsilon_j}^2 &\stackrel{ind.}{\sim} N(\mathbf{0}, \sigma_{\epsilon_j}^2 \mathbf{I}), \\ \mu_j \mid \mu_{0j}, \Sigma_j, v_j &\stackrel{ind.}{\sim} N(\mu_{0j} \mathbf{1}, (v_j - 3)\Sigma_j); \\ \Sigma_j \mid v_j, \Psi_j &\stackrel{ind.}{\sim} IW(\nu_j, \Psi_j), \quad j = 1, \dots, C - 1 \end{aligned}$$

where $\mathbf{Y}_i = (\mathbf{Y}_{i1}, \dots, \mathbf{Y}_{iT_i})^\top$, $\mathbf{Z}_{ij} = (Z_{ij1}, \dots, Z_{ijT_i})^\top$, $\epsilon_{ij} = (\epsilon_{ij1}, \dots, \epsilon_{ijT_i})^\top$, and the collection of the functional evaluations on the pooled grid τ are denoted by the corresponding bold letter.

The structure in (7) is referred to as the continuation-ratio logits representation of the multinomial distribution (Tutz 1991). In the context of Bayesian nonparametric modeling, it has been used as the kernel of nonparametric mixture models for cross-sectional ordinal regression (Kang and Kottas 2022).

Examining model properties reveals the practical utility of the continuation-ratio logits structure. The factorization in (7) allows us to examine the probability response curves and the within subject covariance structure in the same fashion as for binary responses. Specifically, the continuation-ratio logit for response category j is the logit of the conditional probability of response j , given that the response is j or higher. As a consequence, for any finite grid $\tau = (\tau_1, \dots, \tau_T)^\top$, the probability response curves are given by

$$\begin{aligned} \mathbf{P}_{j\tau} &= (\Pr(Y_{\tau_1} = j \mid \mathbf{Z}_\tau, \sigma_\epsilon^2), \dots, \Pr(Y_{\tau_T} = j \mid \mathbf{Z}_\tau, \sigma_\epsilon^2))^\top \\ &= E\left(\pi_{j\tau} \mid \mathbf{Z}_{j\tau}, \sigma_{\epsilon_j}^2\right) \prod_{k=1}^{j-1} E\left((1 - \pi_{k\tau}) \mid \mathbf{Z}_{k\tau}, \sigma_{\epsilon_k}^2\right), \quad (9) \end{aligned}$$

where $\pi_{j\tau} = (\varphi(\mathcal{Z}_{j1}), \dots, \varphi(\mathcal{Z}_{jT}))^\top$ and $\mathcal{Z}_{j\tau} \mid \mathbf{Z}_{j\tau}, \sigma_{\epsilon_j}^2 \sim N(\mathbf{Z}_{j\tau}, \sigma_{\epsilon_j}^2 \mathbf{I}_T)$, for $j = 1, \dots, C$. To avoid redundant expressions, we include the term $\pi_{C\tau}$ and set it always equal to 1. As for the covariance structure, we study the joint probability of the repeated measurements on the same subject at time τ and τ' taking category j and j' . Exploiting the conditional independence structure across the categories,

$$\begin{aligned} &\Pr(Y_\tau = j, Y_{\tau'} = j' \mid \{\mathbf{Z}_{j\tau}\}, \{\sigma_{\epsilon_j}^2\}) \\ &= \begin{cases} E(\pi_{j\tau} \pi_{j\tau'} \mid \mathbf{Z}_{j\tau}, \sigma_{\epsilon_j}^2) \\ \quad \times \prod_{k \neq j} E[(1 - \pi_{k\tau})(1 - \pi_{k\tau'}) \mid \mathbf{Z}_{k\tau}, \sigma_{\epsilon_k}^2] & j = j' \\ E[\pi_{j\tau} (1 - \pi_{j\tau'}) \mid \mathbf{Z}_{j\tau}, \sigma_{\epsilon_j}^2] \\ \quad \times E[(1 - \pi_{j'\tau}) \pi_{j'\tau'} \mid \mathbf{Z}_{j'\tau}, \sigma_{\epsilon_{j'}}^2] \\ \quad \times \prod_{k \neq j, j'} E[(1 - \pi_{k\tau})(1 - \pi_{k\tau'}) \mid \mathbf{Z}_{k\tau}, \sigma_{\epsilon_k}^2] & j \neq j' \end{cases} \quad (10) \end{aligned}$$

Hence, we can explore the covariance of the two ordinal responses $\mathbf{Y}_\tau, \mathbf{Y}_{\tau'}$ by studying the pairwise covariance for each entry.

The continuation-ratio logits structure is also key to efficient model implementation. It implies a sequential mechanism, such that the ordinal response is determined through a sequence of binary outcomes. Starting from the lowest category, each binary outcome indicates whether the ordinal response belongs to that category or to one of the higher categories. This mechanism inspires a novel perspective on the model implementation. That is, we can re-organize the original data set containing longitudinal ordinal responses to create $C - 1$ data sets with longitudinal binary outcomes. Then, fitting model (8) to the original data set is equivalent to fitting the model of Sect. 2.1 separately on the $C - 1$ re-organized data sets. The procedure is elaborated below.

Denote the set of all possible subject and time indices by \mathcal{I}_1 , that is, $\mathcal{I}_1 = \{(i, t) : i = 1, \dots, n, t = 1, \dots, T_i\}$. To build the first re-organized data set with binary outcomes, we create binary indicators $Y_{it}^{(1)}$, such that $Y_{it}^{(1)} = 1$ if $Y_{i1t} = 1$ and $Y_{it}^{(1)} = 0$ if $Y_{i1t} = 0$. The first data set is then $\mathcal{D}_1 = \{Y_{it}^{(1)} : (i, t) \in \mathcal{I}_1\}$. Moving to the second data set, we first filter out the observations that are already categorized into the smallest scale, and denote the remaining indices set by $\mathcal{I}_2 = \mathcal{I}_1 \setminus \{(i, t) : Y_{i1t} = 1\}$. This is the set of indices with original ordinal responses belonging to categories higher than or equal to the second smallest scale. Then, we create new binary indicators $Y_{it}^{(2)}$, such that $Y_{it}^{(2)} = 1$ if $Y_{i2t} = 1$, and $Y_{it}^{(2)} = 0$ if $Y_{i2t} = 0$. The second data set is obtained as $\mathcal{D}_2 = \{Y_{it}^{(2)} : (i, t) \in \mathcal{I}_2\}$. The process is continued until we obtain the $(C - 1)$ -th data set, $\mathcal{D}_{C-1} = \{Y_{it}^{(C-1)} : (i, t) \in \mathcal{I}_{C-1}\}$, where \mathcal{I}_{C-1} is the indices set such that the original ordinal responses belong to either category $C - 1$ or C . Notice that every re-organized data set \mathcal{D}_j , for $j = 1, \dots, C - 1$, contains longitudinal binary

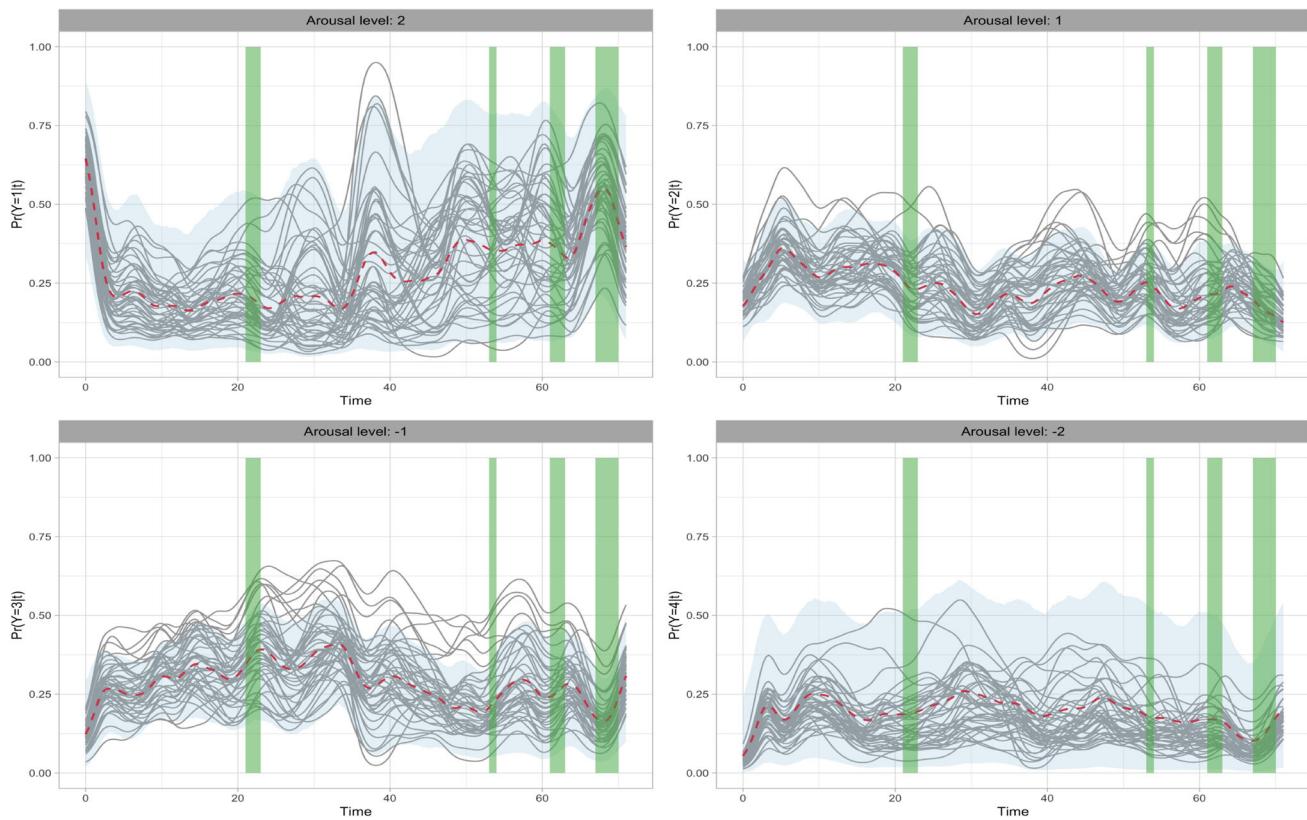


Fig. 5 Four levels arousal score data. Posterior mean (dashed line) and 95% interval estimate (shaded region) of probability response curve for an out-of-sample subject. The posterior mean estimates for the probability response curves of in-sample subjects are given by the solid lines. The vertical shaded regions correspond to the four special time periods (see Sect. 3.1)

outcomes for which the model of Sect. 2.1 is directly applicable. Provided the priors placed on each ordinal response category's parameters are independent, it is straightforward to verify that fitting separately the model for binary responses to the re-organized data sets $\{\mathcal{D}_j : j = 1, \dots, C - 1\}$ is equivalent to fitting model (8) to the original data set. We formalize the conclusion in the following proposition.

Proposition 5 *Fitting the ordinal responses model in (8) is equivalent to fitting the model for binary responses separately, $C - 1$ times to the data sets $\{\mathcal{D}_j : j = 1, \dots, C - 1\}$.*

Based on Proposition 5, the posterior simulation algorithm for the ordinal responses model can be parallelized and implemented on separate cores. In applications where the number of response categories is moderate to large, such a parallel computing scheme is especially beneficial. Also, since the binary responses model serves as the backbone for modeling ordinal responses, the prior specification strategy and the posterior simulation method described in Sect. 2.3 can be readily extended to model (8). Finally, from (9) and (10), it is clear that the posterior samples obtained from the $C - 1$ separate models suffice to obtain full posterior inference for the ordinal response process.

4.2 Data illustration

As an illustration example, we consider the PAM arousal score on the original scale, which is obtained from the same EMA study discussed in Sect. 3. PAM arousal is a -2 to 2 (excluding 0) score. We examine the same cohort of students on the same study period as described in Sect. 3.1. Over all observations, the distribution of arousal scores involves 16.6% for level -2 , 27.7% for level -1 , 12.6% for level 1 , and 12% for level 2 , while 31.1% of the observations are missing.

To implement model (8), we follow the procedure outlined above Proposition 5. We re-organize the original data into separate data sets $\{\mathcal{D}_j : j = 1, \dots, 3\}$, each of them containing the binary responses indicating whether the arousal scores are at level j or a higher level. Then, the proposed model is fitted to the three data sets in parallel.

The primary inference focus is on the change of arousal scores as the term progresses, which is depicted by the probability response curve of each response level. We display posterior point and interval estimates of $\mathbf{P}_{j\tau}$ (defined in (9)) in Fig. 5. The probability of the highest arousal level drops dramatically as the term begins, indicating that the excite-

ment of a new quarter may vanish within a week. The Boston marathon bombing slightly triggers higher probability for moderately low to low arousal level. There is a drop of the probability for moderately high to high arousal level after the Green Key festival and the Memorial Day holiday. The exams may have a significant impact on the arousal level. We observe peaks of arousal at the beginning of the final exam period, and also the middle of the term, which corresponding to the midterm exam period. Since the students are taking different courses, the midterm exam times vary, resulting in some curves with lead or lag peaks compared to the majority. This pattern is not clear in the analysis of binary arousal scores. Hence, examining the finer ordinal scale enables us to discover subtle changes of the students activation states. We have also investigated the temporal covariance structure of the ordinal responses, with details presented in the Supplementary Material.

5 Discussion

We have developed a novel Bayesian hierarchical model for analyzing longitudinal binary data. We approach the problem from a functional data analysis perspective, resulting in a method that is suitable for either regularly or irregularly spaced longitudinal data. The modeling approach achieves flexibility and computational efficiency in full posterior inference. With regard to the former, the key model feature is the joint and nonparametric modeling of the mean and covariance structure. As illustrated by the data application, our approach enables interpretable inference with coherent uncertainty quantification, and provides improvement over the GLMM approach. The model formulation enables a natural extension to incorporate ordinal responses, which is accomplished by leveraging the continuation-ratio logits representation of the multinomial distribution. This representation leads to a factorization of the multinomial model into separate binomial models, on which the modeling approach for binary responses can be applied. The computational benefit is retained, since we can utilize parallel computing across response categories.

The proposed methodology for modeling longitudinal binary and ordinal responses can be elaborated in different directions. We have focused on stationary specifications for the hierarchical GP prior. Nonstationary model components can be incorporated through a time-varying mean function μ_0 and/or a nonstationary covariance kernel Ψ_ϕ . Moreover, longitudinal studies typically have predetermined covariates associated with each subject, or time-varying covariates corresponding to each observation. The predetermined covariates can be incorporated in the model through the prior placed on the mean function of the signal process. Using the functional linear model may be a possible strategy

for the more challenging task of accounting for time-varying covariates.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s11222-024-10525-2>.

Acknowledgements The research was supported in part by the National Science Foundation under award SES 1950902. The authors wish to thank two reviewers for several useful comments.

Author Contributions All authors contributed to the conception and design. JK contributed to the methodology, coding and writing the original draft. AK contributed to supervision, revision, and editing.

Data availability No datasets were generated or analysed during the current study.

Declarations

Conflict of interest The authors have no conflict of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Andrianakis, I., Challenor, P.G.: The effect of the nugget on Gaussian process emulators of computer models. *Computat. Stat. Data Anal.* **56**, 4215–4228 (2012)
- Carmack, P.S., Spence, J.S., Schucany, W.R., Gunst, R.F., Lin, Q., Haley, R.W.: A new class of semiparametric semivariogram and nugget estimators. *Comput. Stat. Data Anal.* **56**, 1737–1747 (2012)
- Cheng, L., Ramchandran, S., Vatanen, T., Lietzén, N., Lahesmaa, R., Vehtari, A., Lähdesmäki, H.: An additive Gaussian process regression model for interpretable non-parametric analysis of longitudinal data. *Nat. Commun.* **10**, 1798 (2019)
- Daniels, M.J., Xu, D.: Bayesian Methods for Longitudinal Data with Missingness. In: Lesaffre, E., Baio, G., Boulanger, B. (eds.) *Bayesian Methods in Pharmaceutical Research*. Chapman and Hall/CRC, pp. 185–205 (2020)
- Dawid, A.P.: Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika* **68**, 265–274 (1981)
- DeYoreo, M., Kottas, A.: Modeling for dynamic ordinal regression relationships: an application to estimating maturity of rockfish in California. *J. Am. Stat. Assoc.* **113**, 68–80 (2018)
- Di Lucca, M.A., Guglielmi, A., Müller, P., Quintana, F.A.: A simple class of Bayesian nonparametric autoregression models. *Bayesian Anal.* **8**, 63–88 (2013)

Diggle, P.J.: An approach to the analysis of repeated measurements. *Biometrics* **44**, 959–971 (1988)

Donà, G., Preatoni, E., Cobelli, C., Rodano, R., Harrison, A.J.: Application of functional principal component analysis in race walking: an emerging methodology. *Sports Biomech.* **8**, 284–301 (2009)

Fong, Y., Rue, H., Wakefield, J.: Bayesian inference for generalized linear mixed models. *Biostatistics* **11**, 397–412 (2010)

Fryer, D., Nguyen, H., Orban, P.: Studentlife: Tidy Handling and Navigation of the Student-Life Dataset. R package version 1.1.0 (2022)

Gelfand, A.E., Ghosh, S.K.: Model choice: a minimum posterior predictive loss approach. *Biometrika* **85**, 1–11 (1998)

Ghosh, P., Hanson, T.: A semiparametric Bayesian approach to multivariate longitudinal data. *Aust. N. Z. J. Stat.* **52**, 275–288 (2010)

Gneiting, T., Raftery, A.E.: Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**, 359–378 (2007)

Hall, P., Müller, H.-G., Yao, F.: Modelling sparse generalized longitudinal observations with latent Gaussian processes. *J. Roy. Stat. Soc. B* **70**, 703–723 (2008)

Ingrassia, S., Costanzo, G.D.: Functional principal component analysis of financial time series. In: Vichi, M., Monari, P., Mignani, S., Montanari, A. (eds.) *New developments in classification and data analysis*. Springer, Berlin, Heidelberg (2005)

Jara, A., José García-Zattera, M., Lesaffre, E.: A Dirichlet process mixture model for the analysis of correlated binary responses. *Adv. Mix. Mod. Comput. Stat. Data Anal.* **51**, 5402–5415. (2007)

Jiang, L., Zhong, Y., Elrod, C., Natarajan, L., Knight, R., Thompson, W. K.: BayesTime: Bayesian functional principal components for sparse longitudinal data. [arXiv:2012.00579](https://arxiv.org/abs/2012.00579) (2020)

Kang, J., Kottas, A.: Structured mixture of continuation-ratio logits models for ordinal regression, [arXiv:2211.04034](https://arxiv.org/abs/2211.04034) (2022)

Kunihama, T., Halpern, C.T., Herring, A.H.: Non-parametric Bayes models for mixed scale longitudinal surveys. *J. R. Stat. Soc. Ser. C* **68**, 1091–1109 (2019)

Li, Y., Lin, X., Müller, P.: Bayesian inference in semiparametric mixed models for longitudinal data. *Biometrics* **66**, 70–78 (2010)

Matuk, J., Herring, A. H., Dunson, D. B.: Bayesian functional principal component analysis using relaxed mutually orthogonal processes, [arXiv:2205.12361](https://arxiv.org/abs/2205.12361) (2022)

Molenberghs, G., Verbeke, G.: Models for discrete longitudinal data, Springer Series in Statistics, Springer-Verlag (2022)

Pollak, J. P., Adams, P., Gay, G.: PAM: A photographic affect meter for frequent, in situ measurement of affect. In: Proceedings of the SIGCHI conference on human factors in computing systems (2011)

Polson, N.G., Scott, J.G., Windle, J.: Bayesian inference for logistic models using Polya-Gamma latent variables. *J. Am. Stat. Assoc.* **108**, 1339–1349 (2013)

Quintana, F.A., Johnson, W.O., Waetjen, L.E., Gold, E.B.: Bayesian nonparametric longitudinal data analysis. *J. Am. Stat. Assoc.* **111**, 1168–1181 (2016)

Rasmussen, C.E., Williams, C.K.I.: *Gaussian Processes for Machine Learning*. The MIT Press, Cambridge (2006)

Rue, H., Martino, S., Chopin, N.: Approximate Bayesian inference for latent Gaussian models using integrated nested laplace approximations (with discussion). *J. R. Stat. Soc. B* **71**, 319–392 (2009)

Rue, H., Riebler, A. I., Sørbye, S. H., Illian, J. B., Simpson, D. P., Lindgren, F. K.: Bayesian computing with INLA: A review, *Ann. Rev. Stat. Appl.* 395–421 (2017)

Russell, J.A.: A circumplex model of affect. *J. Pers. Soc. Psychol.* **39**, 1161–1178 (1980)

Ruwaard, J., Kooistra, L., Thong, M.: *Ecological Momentary Assessment in Mental Health Research: A Practical Introduction, with Examples in R*, APH Mental Health, 1st (build 2018-11-26) edition (2018)

Shah, A., Wilson, A., Ghahramani, Z.: Student-t processes as alternatives to Gaussian processes. In: Kaski, S., Corander, J. (eds.) *Proceedings of the seventeenth international conference on artificial intelligence and statistics*, volume 33 of *proceedings of machine learning research*, Reykjavik, Iceland: PMLR (2014)

Shamshoian, J., Şentürk, D., Jeste, S., Telesca, D.: Bayesian analysis of longitudinal and multidimensional functional data. *Biostatistics* **23**, 558–573 (2020)

Stein, M.L.: *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Series in Statistics, Springer (1999)

Tang, N.-S., Duan, X.-D.: A semiparametric Bayesian approach to generalized partial linear mixed models for longitudinal data. *Comput. Stat. Data Anal.* **56**, 4348–4365 (2012)

Tutz, G.: Sequential models in categorical regression. *Comput. Stat. Data Anal.* **11**, 275–295 (1991)

Van Der Linde, A.: A Bayesian latent variable approach to functional principal components analysis with binary and count data. *AStA Adv. Stat. Anal.* **93**, 307–333 (2009)

Wang, R., Chen, F., Chen, Z., Li, T., Harari, G., Tignor, S., Zhou, X., Ben-Zeev, D., Campbell, A. T.: StudentLife: Assessing mental health, academic performance and behavioral trends of college students using smartphones. In: *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing, UbiComp '14*. Association for Computing Machinery New York, NY, USA (2014)

Yang, J., Zhu, H., Choi, T., Cox, D.D.: Smoothing and mean-covariance estimation of functional data with a Bayesian hierarchical model. *Bayesian Anal.* **11**, 649–670 (2016)

Yao, F., Müller, H.-G., Wang, J.-L.: Functional data analysis for sparse longitudinal data. *J. Am. Stat. Assoc.* **100**, 577–590 (2005)

Zhao, X., Marron, J.S., Wells, M.T.: The functional data analysis view of longitudinal data. *Stat. Sin.* **14**, 789–808 (2004)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.