

Dynamic Functional Variable Selection for Multimodal mHealth Data

Matthew D. Koslovsky^{*}, Kelley Pettee Gabriel[†], Michael Businelle[‡],
David W. Wetter[§], and Darla Kendzor[‡]

Abstract. Mobile health (mHealth) methods allow researchers to monitor study participants in their natural environments in order to improve health-related outcomes through behavior change. MHealth investigators are interested in understanding the feasibility of supplementing or even replacing actively collected data with passively collected data to reduce participant burden, while also increasing the temporal resolution of the data. In this work, we propose a novel Bayesian dynamic functional variable selection method to explore the relations between multimodal mHealth data collected on different time scales. Specifically, our approach leverages spiked hierarchical species sampling priors to identify critical moments when a participant experiences momentary spikes in a low-frequency outcome, characterize high-frequency outcome trajectories which are related to these critical moments, and cluster the relational trends to explore potential subpopulations. We introduce continuous-time multistate Markov model priors to inform selection based on information learned at previous assessments. We demonstrate the variable selection and clustering performance of our model in various simulation settings motivated by the data structures found in mHealth studies. We then apply our model to multimodal intensive longitudinal data collected in the Pathways between Socioeconomic Status and Behavioral Cancer Risk Factors Study to explore relations between physical activity passively collected with accelerometers and mood actively collected with ecological momentary assessment methods.

Keywords: accelerometry, clustering, ecological momentary assessment, joint modeling, nonparametric Bayes, species sampling priors.

1 Introduction

Behavioral mobile health (mHealth) research refers to the use of mobile phones and other wireless technology to monitor study participants in their natural environments with the goal of understanding and improving health-related outcomes through behavior change. These methods have been applied in various research areas including smoking cessation,

^{*}Department of Statistics, College of Natural Sciences, Colorado State University, Fort Collins, CO, 80521, USA, matt.koslovsky@colostate.edu

[†]Department of Epidemiology, School of Public Health, University of Alabama at Birmingham, Birmingham, AL, 35294, USA, gabrielk@uab.edu

[‡]Department of Family and Preventive Medicine, College of Medicine, University of Oklahoma, Oklahoma City, OK, 73104, USA, Michael-Businelle@ouhsc.edu

[§]Department of Population Health Sciences, College of Medicine, University of Utah, Salt Lake City, UT, 84132, USA, David.Wetter@hci.utah.edu

[‡]Department of Family and Preventive Medicine, College of Medicine, University of Oklahoma, Oklahoma City, OK, 73104, USA, Darla-Kendzor@ouhsc.edu

medication adherence, drug abuse prevention, sports injury, and promotion of healthy behaviors, including physical activity and diet (Businelle et al., 2016; Kendzor et al., 2016; Walsh et al., 2016; Kazemi et al., 2017; Anglada-Martinez et al., 2015; Zapata-Lamana et al., 2020; Businelle et al., 2024). A popular mHealth technique in behavioral health research is the use of ecological momentary assessment (EMA) methods which involve the repeated assessment of individuals in their natural environment to capture behavioral, psychological, and environmental factors that may relate to a behavioral outcome in near real time. In these studies, participants are often prompted to respond to report-based surveys on mobile devices such as smartphones at random or prespecified time points. In addition to EMAs, behavioral health researchers may passively collect various forms of information from wearable sensors (e.g., acceleration/movement, locations, physiological information such as heart rate or skin temperature, and biochemical measures). By fusing or integrating these different data sources, or multimodal data, together, researchers are able to investigate participants' potential risk factors and behavioral outcomes from a third-person observational perspective which promotes ecological validity (Nelson and Allen, 2018; Kumar et al., 2013; Mitchell, 2007; Rehg et al., 2017). Overall, mHealth methods have helped researchers better understand complex psychological and behavioral processes, in addition to enhancing their ability to design, evaluate, and deliver tailored intervention strategies based on a participant's risk profile at critical moments throughout the assessment period (Rehg et al., 2017; Heron and Smyth, 2010; Nahum-Shani et al., 2017; Businelle et al., 2016; Hébert et al., 2020).

mHealth investigators are currently evaluating the feasibility of supplementing or even replacing actively collected intensive longitudinal data (ILD) with passively collected data. Here, ILD refers to repeated measures data intensively collected on individuals to capture complex patterns of change over time (Walls and Schafer, 2006). For example, studies have recently begun to examine the relation between EMAs of negative affect (active data collection) and accelerometer measurements of physical activity (passive data collection), rather than relying upon report-based methods that use a recall period (e.g., past week) to prompt recall (Kim et al., 2018). Passively collecting data would reduce participant burden and can increase the temporal resolution of the data, while still providing insights into related behavioral patterns and momentary risk factors that may initiate intervention delivery (Blaauw et al., 2016; Bertz et al., 2018). Integrated analyses targeting this aim are challenged by the high-dimensionality, large amounts of within- and between-subject heterogeneity, missingness, measurement error, repeated measures structure, and unbalanced, unequally-spaced assessment times that are characteristic of mHealth data. Analysis is further challenged when multimodal behavioral mHealth data (e.g., EMAs and accelerometer data) are collected at different temporal resolutions, which requires novel analytical approaches for information fusion (Kumar et al., 2013).

To explore the relations between passively and actively collected behavioral mHealth ILD on different time scales, we design a Bayesian joint model which links submodels for high-frequency accelerometer data and low-frequency EMA data together with a novel spiked hierarchical nonparametric prior on their joint distribution. Our method

targets the following research objectives: (1) identify critical moments when a participant experiences spikes in the low-frequency outcome; (2) characterize trajectories of the high-frequency measure which are related to these critical moments; and (3) cluster the relational trends within and between participants to suggest potential subpopulations. By addressing these exploratory research aims, our proposed model belongs to the emerging class of dynamic variable selection methods but breaks new ground by performing dynamic *functional* variable selection.

Commonly, researchers take a two-stage approach to model the relations between ILD collected on different time scales in which the low-frequency outcomes are matched with summary measures of the high-frequency data over prespecified assessment windows, or *epochs*, based on their proximity in time. In practice, summary measures may take on various forms from basic summary statistics (e.g., sample means and variances) to parameter estimates from fitted models (e.g., subject-level effects from mixed-effects location-scale models) (Kürüm et al., 2016; Dzubur et al., 2020; de Brito et al., 2020; Cushing et al., 2017). Thereafter, the relations between the low-frequency outcomes and proximal summaries of the high-frequency measures are typically investigated using standard repeated measures modeling techniques, such as generalized estimating equations and generalized linear mixed models (Fitzmaurice et al., 2008). However, inference is highly dependent on the selection of summary measures used in the first stage of the analysis which may not fully capture the complexity of the individual trajectories within each epoch. Additionally, it is well known that two-stage approaches may underestimate model uncertainty and potentially produce biased results (Sayers et al., 2017).

Alternatively, joint modeling techniques provide a flexible framework for handling multimodal longitudinal data and have been shown to properly accommodate model uncertainty and reduce inferential bias in various settings (Tsiatis and Davidian, 2004; Rizopoulos and Lesaffre, 2014; Bigelow and Dunson, 2009). The strength of this framework stems from its flexibility in accommodating various types of outcome structures found in longitudinal data analyses (e.g., continuous, binary, time-to-event, or recurrent events) and specifications of the shared parameters (e.g., individual-level intercept or slope terms or area under the longitudinal trajectory over a window of time) which are used to link outcomes together. Despite their widespread use in biomedical research, joint modeling techniques are often overlooked in mHealth settings where they are particularly relevant (Scherer et al., 2017).

Historically, functional data analysis (FDA) methods have played an integral role in mHealth research for investigating the dynamic relations between risk factors and behavioral outcomes using ILD for intervention evaluation and design (Dziak et al., 2015; Tan et al., 2012; Koslovsky et al., 2018a, 2020; Liang et al., 2023). FDA methods are well-suited for high-dimensional data with unbalanced and unequally-spaced observation times, matching the format of data collected with mHealth methods. They also require few assumptions on the structure of the relations between risk factors and behavioral outcomes. One of the most popular FDA methods applied in mHealth studies is the concurrent functional, or varying-coefficient, regression model which relates a functional response to concurrently observed functional covariates (Ramsay and Silverman, 2002;

Hastie and Tibshirani, 1993). This approach is often used to investigate how a risk factor’s relation with the outcome varies as a function of time, often referred to as a time-varying effect model (Tan et al., 2012). Recent extensions include mixture formulations to investigate subgroups of subjects who respond similarly over time (Dziak et al., 2015) and variable selection techniques to identify significant functional covariates and learn clusters of risk factors that share similar trends (Koslovsky et al., 2020; Goldsmith and Schwartz, 2017). Relatedly, Islam et al. (2018) introduced a two-step variable selection method for varying functional linear models to identify relations between recent past behavior of forearm electromyogram signals with finger/wrist velocities, allowing the functional terms to vary with position. While these methods provide unique insights into the complex relations found in mHealth data streams, they are not designed to explicitly identify critical moments when the data may be related over time, information that is necessary for the evaluation, design, and delivery of tailored intervention strategies for behavior change (Walls, 2013).

Advanced methods for *dynamic* variable selection have recently been proposed to identify moments in which regressors are active or inactive over time (Rockova and McAlinn, 2021; Cassese et al., 2019; Kowal et al., 2019). In particular, Cassese et al. (2019) employ zero-inflated conditionally identically distributed species sampling priors to identify momentary departures or spikes of a process from a baseline state over time and space. Their model belongs to the class of “spiked” nonparametric priors, which have recently gained traction in various applied settings as they provide Bayesian nonparametric inference on clustering patterns while simultaneously performing variable selection (Kim et al., 2009; Koslovsky et al., 2020; Savitsky and Vannucci, 2010; Canale et al., 2017). Despite their strengths, existing methods for dynamic variable selection are limited to individual time series or data collected on the same time scale, do not accommodate functional covariates, and are not designed to identify clusters of temporal trends within and between participants, which limits their use in mHealth research.

In this work, we design a Bayesian joint model to perform dynamic *functional* variable selection on ILD collected on two different time scales. Specifically, our approach leverages spiked hierarchical species sampling priors (sHSSPs) to identify and cluster momentary spikes in a low-frequency outcome that are associated with the functional trajectory of a high-frequency outcome measured within proximal epochs. To capture individuals’ momentary spikes in their low-frequency outcome over time, we embed dynamic latent inclusion indicators which we assume follow a continuous-time multistate Markov model. This approach accommodates unbalanced, unequally-spaced assessment times and potential missingness from device malfunction or nonresponse, characteristic of data collected in mHealth studies. Our method is developed to explore the relations between multimodal ILD collected in the Pathways between Socioeconomic Status and Behavioral Cancer Risk Factors (Pathways) Study. By addressing the research objectives outlined above, we develop the first method for dynamic functional variable selection.

In the following section, we present our dynamic functional variable selection method for ILD collected on two different time scales. In Section 3, we evaluate the dynamic functional variable selection and clustering performance of the proposed method on simulated data. Section 4 applies the proposed method to investigate multimodal mHealth data collected in the Pathways study. We conclude with final remarks in Section 5.

2 Model

We first introduce the high-frequency outcome model used to estimate the functional trajectories. Next, we describe the low-frequency outcome model which is designed to capture observation-specific spikes from participants' average observations over the assessment period. We then construct a hierarchical prior for the joint distribution of the trajectories and proximal observation-specific effects which links the two models together. Let Y_{ij} represent the low-frequency outcome for the i^{th} participant, $i = 1, \dots, N$, collected at the j^{th} assessment, $j = 1, \dots, n_i$. Motivated by the application data, assume $U_{ijk} \in \{0, 1\}$ represents the k^{th} measurement, $k = 1, \dots, n_{ij}$, of the i^{th} participant's high-frequency binary outcome within the epoch paired with the j^{th} assessment. For example, 1 may represent active behavior, and 0 otherwise. Let t_{ij} represent the time at which observation Y_{ij} is collected, and t_{ijk} represent the time U_{ijk} is measured.

2.1 High-Frequency Outcome Model

To model the binary high-frequency outcome, we assume

$$\text{logit}(P(U_{ijk} = 1 | t_{ijk}, \mathbf{x}_{ijk}, \boldsymbol{\theta}_{ij})) = f_{ij}(t_{ijk}) + \mu_U(t_{ijk}), \quad (1)$$

where $f_{ij}(\cdot)$ represents the functional trajectory for the i^{th} participant at the j^{th} assessment and $\mu_U(t_{ijk})$ represents a general regression term, which may include subject- or population-level trends given baseline or time-dependent covariates, \mathbf{x}_{ijk} . In FDA, functional trajectories are typically modeled with basis functions (e.g., splines, Fourier series, principal components, wavelets, and others), where the type of basis functions used depends on the characteristics of the data (Ramsay and Silverman, 2002). For example, Fourier basis functions are often used for cyclical or periodic functions, wavelets for irregular functions with rapid changes or discontinuities, and spline basis functions for relatively smooth, nonperiodic functions (Morris, 2015). In general, the proposed dynamic functional variable selection framework is flexible to various basis functions as well as fully nonparametric approaches (e.g., Gaussian processes) to model the smooth functions (Shi and Choi, 2011; Rodríguez et al., 2009). Since the number of smooth functions estimated in the model increases with the number of participants and assessments, we preferred a semiparametric approach over nonparametric alternatives due to its computational efficiency. Given the structure of the high-frequency outcome data in the application (i.e., relatively smooth, nonperiodic), we chose to approximate the functional trajectories $f_{ij}(\cdot)$ with B-spline basis functions, where the S -dimensional vector $\boldsymbol{\theta}_{ij}$ represents the corresponding spline coefficients. To reduce the impact of choosing the number and position of the knots on inference, we chose to include a large number of equidistant knots and then penalized the spline coefficients, $\boldsymbol{\theta}_{ij}$, with a horseshoe prior (Carvalho et al., 2010) which encourages smoothness and avoids overfitting. For efficient sampling of the resulting posterior distribution of $\boldsymbol{\theta}_{ij}$, we implemented the data augmentation approach of Makalic and Schmidt (2015). Specifically, we assume $\theta_{ijs} \sim \text{Normal}(0, v_\theta v_{\theta_s})$, for $s = 1, \dots, S$, where the global variances $v_\theta \sim \text{Inverse-Gamma}(1/2, 1/\mathcal{A}_\theta)$, local variances $v_{\theta_s} \sim \text{Inverse-Gamma}(1/2, 1/\vartheta_{\theta_s})$, and auxiliary parameters $\vartheta_{\theta_1}, \dots, \vartheta_{\theta_S}, \mathcal{A}_\theta \sim \text{Inverse-Gamma}(1/2, 1)$. Note that the spline

coefficients will eventually be linked to the observation-specific effects introduced in Section 2.2 via a hierarchical model on their joint distribution.

2.2 Low-Frequency Outcome Model

For the low-frequency outcome, we assume a linear regression model

$$Y_{ij} = \alpha_{ij} + \mu_Y(t_{ij}) + \epsilon_{ij}, \quad (2)$$

where α_{ij} represents an observation-specific effect for the i^{th} participant at the j^{th} assessment, $\mu_Y(t_{ij})$ represents a general regression term, which may include subject- or population-level trends given baseline or time-dependent covariates, \mathbf{x}_{ij} , and ϵ_{ij} is an error term. We propose α_{ij} as an observation-specific intercept term, but the model could easily be adjusted to accommodate observation-specific effects for other covariates.

The interpretation of α_{ij} depends on the specification of $\mu_Y(t_{ij})$. For example, if $\mu_Y(t_{ij}) = \beta_i$, then a non-zero value for α_{ij} is interpreted as a departure in a participant's mean low-frequency outcome at a particular assessment. To identify these critical moments, we assume spike-and-slab priors for α_{ij} (George and McCulloch, 1997; Brown et al., 1998). See Section 2.3 for more details, including the hierarchical model for the joint distribution of α_{ij} and $\boldsymbol{\theta}_{ij}$ that links the low- and high-frequency outcome models together. To complete the low-frequency outcome model specification, we assume $\epsilon_{ij} \sim \text{Normal}(0, \sigma^2)$ with $\sigma^2 \sim \text{Inverse-Gamma}(a_{\sigma^2}, b_{\sigma^2})$.

2.3 Spiked Hierarchical Species Sampling Priors

Species sampling models are a general class of discrete random measures that includes the famed Dirichlet process (Ferguson, 1973) and its generalization the Pitman-Yor process (Pitman and Yor, 1997; Pitman, 1996). Species sampling priors are commonly used to govern clustering allocation due to their computational simplicity (Ray and Mallick, 2006; Suarez and Ghosal, 2016; White and Gelfand, 2020; Das et al., 2021). Hierarchical species sampling priors enable group-based clustering and allow the sharing of clusters across groups (Bassetti et al., 2020). To identify clusters of high-frequency outcome trajectories that are dynamically associated with momentary spikes in the low-frequency outcome within and between participants, we jointly model the observation-specific effects, α_{ij} , and the functional parameters, $\boldsymbol{\theta}_{ij}$, with a hierarchical species sampling prior (HSSP). Specifically for $\Theta_{ij} = (\alpha_{ij}, \boldsymbol{\theta}_{ij})$, we assume

$$\begin{aligned} \Theta_{i1}, \dots, \Theta_{in_i} | p_i &\sim p_i, \quad i = 1, \dots, N, \\ (p_1, \dots, p_N) | p_0 &\sim SSrp(q_1, p_0), \\ p_0 &\sim SSrp(q_0, G_0), \end{aligned} \quad (3)$$

where $SSrp$ is a species sampling random probability, q_1 and q_0 are exchangeable partition probability functions characterizing the process, and G_0 is the base distribution. Common choices of $SSrps$ include the Pitman-Yor (PY) process (Pitman and Yor, 1997), $PY(\rho, \vartheta, G)$, and the Dirichlet Process (DP) (Ferguson, 1973), $DP(\vartheta, G)$, with

discount parameter ϱ , $0 \leq \varrho < 1$, concentration parameter $\vartheta > -\varrho$, and base distribution G . For example by assuming a DP at both levels of the HSSP in (3), the hierarchical DP presented in Teh et al. (2006) is obtained, which we denote as DPDP($\vartheta_0, \vartheta_1, G_0$). See Bassetti et al. (2020) for more details of HSSPs, including those with non-diffuse base distributions, and Ghosal and Van der Vaart (2017), Müller and Quintana (2004), and Hjort et al. (2010) for a more general review of Bayesian nonparametric modeling approaches.

In exploratory research settings, researchers often use spiked nonparametric priors to simultaneously perform variable selection and clustering in a unified framework (Dunson et al., 2008; Kim et al., 2009; Savitsky and Vannucci, 2010; Canale et al., 2017; Cassese et al., 2019; Koslovsky et al., 2020). There are two popular techniques for constructing spiked nonparametric priors (i.e., “inner” and “outer” formulations (Canale et al., 2017)). “Inner” formulations, which assume a non-diffuse base distribution (e.g., the spike-and-slab prior (George and McCulloch, 1997; Brown et al., 1998)), are considered less informative than “outer” formulations and more robust to prior misspecification (Canale et al., 2017). In this work, we construct an “inner” formulation in order to identify moments in which the high-frequency outcome trajectories are associated with momentary spikes in the low-frequency outcome. Specifically, we assume that the base distribution of the HSSP, G_0 , is the product of S normal distributions for θ_{ij} described in Section 2.1 and the spike-and-slab prior for α_{ij} introduced in Section 2.2. Specifically,

$$G_0 = \prod_{s=1}^S N(0, v_{\theta} v_{\theta_s}) \times (N(0, v_{\alpha}) \gamma_{ij} + \delta_0(1 - \gamma_{ij})),$$

where $\gamma_{ij} \in \{0, 1\}$ is a latent inclusion indicator and $\delta_0(\cdot)$ is a Dirac delta function, or point mass, at zero. The latent inclusion indicator represents whether or not a subject is experiencing a spike in negative affect at a given assessment and determines whether the corresponding observation-specific effect, α_{ij} , is set to zero (spike) or is freely estimated in the model (slab). By assuming a normal slab distribution, our method is able to detect positive and negative spikes in the low-frequency over time which are related to unique high-frequency outcome trajectories. Note that our approach is agnostic to the length of the observation window for the high-frequency outcome trajectories and their proximal relation with the corresponding low-frequency outcome.

2.4 Dynamic Selection Priors

The prior specification for the latent inclusion indicators controls the overall sparsity of the model. Common assumptions include beta-binomial priors, logistic regression priors which allow covariate information to inform the prior probability of inclusion, and Markov random field priors which accommodate (un)known graphical structure (Stingo et al., 2010; Li and Zhang, 2010). Since the latent inclusion indicators in the proposed method represent momentary spikes in the low-frequency outcome over time, we propose modeling transitions between spikes dynamically using a continuous-time multistate Markov (MSM) model prior. While MSM models are often applied to mHealth data to model participants as they transition through discrete behavioral states over time

(Koslovsky et al., 2018b; Liang et al., 2021; Liu et al., 2017), they have yet to be used as variable selection priors. Note that we assume a continuous-time MSM model as opposed to a discrete-time MSM model to handle unbalanced and unequally-spaced assessment times characteristic of mHealth data. Specifically, we assume that the probability of γ_{ij} depends on the realization of the previous inclusion indicator $\gamma_{i,j-1}$. For a two-state Markov model, the transition probability matrix is defined as

$$\mathbf{P}(\delta_{ij}) = \begin{array}{cc} & \begin{array}{cc} \text{current} & \text{state} \\ 0 & 1 \end{array} \\ \begin{array}{c} \text{previous} \\ \text{state} \end{array} & \begin{array}{cc} 0 & 1 \\ \left[\begin{array}{cc} P(0,0|\delta_{ij}) & P(0,1|\delta_{ij}) \\ P(1,0|\delta_{ij}) & P(1,1|\delta_{ij}) \end{array} \right] \end{array} \end{array},$$

where $P(0,1|\delta_{ij})$ is interpreted as the probability of transitioning from an inactive latent inclusion indicator ($\gamma_{i,j-1} = 0$) to an active inclusion indicator ($\gamma_{ij} = 1$) given $\delta_{ij} = t_{ij} - t_{i,j-1}$, the change in time between the current and previous assessment. Transition probabilities have closed-form solutions, namely

$$P_{01}(\delta_{ij}) = 1 - P_{00}(\delta_{ij}) = \frac{\lambda}{\lambda + \mu} [1 - \exp(-(\lambda + \mu)\delta_{ij})]$$

$$P_{10}(\delta_{ij}) = 1 - P_{11}(\delta_{ij}) = \frac{\mu}{\lambda + \mu} [1 - \exp(-(\lambda + \mu)\delta_{ij})],$$

where $\mu, \lambda < 0$ represent transition rates (Pinsky and Karlin, 2010). For inference, the steady state probability of experiencing a spike in the low-frequency outcome $P_1 = 1 - P_0 = \frac{\lambda}{\lambda + \mu}$. In practice, μ and λ can be chosen a priori to induce sparsity in the model or freely estimated by specifying a prior distribution that may depend on other covariates.

In summary, the proposed dynamic functional variable selection method described above is designed to explore the relations between passively and actively collected behavioral mHealth ILD on different time scales. Our approach links submodels for high-frequency accelerometer data (Section 2.1) and low-frequency EMA data (Section 2.2) together with a novel spiked hierarchical species sampling prior (Section 2.3) on their joint distribution. To capture individuals' momentary spikes in their low-frequency outcome over time, we embed dynamic latent inclusion indicators which we assume follow a continuous-time multistate Markov model (Section 2.4). As a result, the model is able to identify critical moments that a participant experiences momentary spikes in a low-frequency outcome, characterize high-frequency outcome trajectories which are related to these critical moments, and cluster the relational trends to explore potential subpopulations.

2.5 Posterior Inference

For posterior inference, we implement a Metropolis-Hastings within Gibbs algorithm. The full joint posterior distribution is defined as

$$f(\mathbf{Y}|\mathbf{x}, \boldsymbol{\alpha}, \mathbf{c}, \mathbf{d})f(\mathbf{U}|\mathbf{x}, \mathbf{c}, \mathbf{d}, \boldsymbol{\theta})p(\mathbf{c}, \mathbf{d})p(\boldsymbol{\alpha}|\boldsymbol{\gamma})p(\boldsymbol{\gamma})p(\boldsymbol{\theta}|\boldsymbol{\nu}_\theta, \boldsymbol{\vartheta}_\theta, \mathcal{A}_\theta)p(\boldsymbol{\nu}_\theta)p(\boldsymbol{\vartheta}_\theta)p(\mathcal{A}_\theta).$$

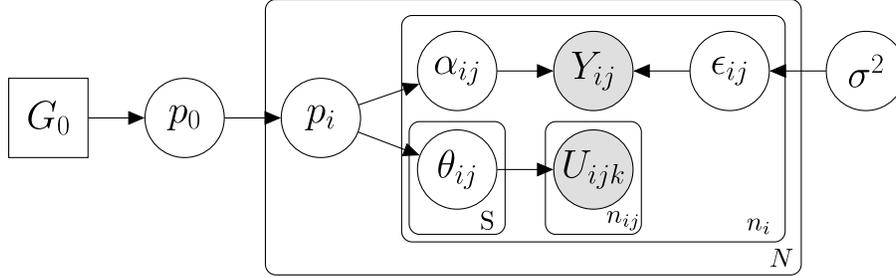


Figure 1: Graphical representation of the proposed Bayesian joint model for dynamic functional variable selection. Auxiliary covariates and hyperparameters are suppressed for clarity.

A graphical representation of a simplified version of our model focusing on the sHSSP formulation for dynamic functional variable selection is provided in Figure 1. The Markov chain Monte Carlo (MCMC) sampler used to implement our model is outlined below in Algorithm 1. A more detailed description of the MCMC steps is provided in the Supplementary Material. After burn-in, the remaining samples obtained from running Algorithm 1 for M iterations are used for inference. To determine a significant departure in the low-frequency outcome (i.e., active inclusion indicator γ_{ij}), the marginal posterior probability of inclusion (MPPI) is empirically estimated by calculating the average of its respective inclusion indicator’s MCMC samples (George and McCulloch, 1997). Typically, covariates are included in the model if their MPPI exceeds 0.50 (Barbieri et al., 2004) or a Bayesian false discovery rate threshold, which controls for multiplicity (Newton et al., 2004). Clusters of spikes in the low-frequency outcome and proximal high-frequency outcome trajectories are determined using sequentially-allocated latent structure optimization to minimize the lower bound of the variation of information loss (Wade and Ghahramani, 2018; Dahl et al., 2021).

3 Simulation Study

In this section, we evaluate the clustering and selection performance of our model on simulated data. To our knowledge, there are no existing methods for dynamic functional variable selection currently developed for comparison, so we compare our model with various parameterizations of the sHSSP, including a Dirichlet process at both levels (sDPDP($\varrho_0, \varrho_1, G_0$)), a Pitman-Yor process at both levels (sPYPY($\varrho_0, \vartheta_0, \varrho_1, \vartheta_1, G_0$)), a Pitman-Yor-Dirichlet process (sPYDP($\varrho_0, \vartheta_0, \vartheta_1, G_0$)), and a Dirichlet-Pitman-Yor process (sDPPY($\vartheta_0, \varrho_1, \vartheta_1, G_0$)).

We simulated $N = 20$ subjects with $n_i = 20$ observations of the low-frequency outcome which were each paired with $n_{ij} = 20$ high-frequency measures. Thus, we generated a total of 400 low-frequency and 8,000 high-frequency observations. For each subject, we first simulated inclusion indicators following a multistate Markov model with $\lambda = \exp(3)$ and $\mu = \exp(2)$ using the `msm` package in R (Jackson, 2011). Observation

Algorithm 1 MCMC Sampler.

```

1: Input data  $\mathbf{Y}, \mathbf{U}, \mathbf{X}_{ij}, \mathbf{X}_{ijk}, \mathbf{T}_{ij}, \mathbf{T}_{ijk}$ 
2: Initialize parameters:  $\mathbf{c}, \mathbf{d}, \boldsymbol{\alpha}, \boldsymbol{\gamma}, \sigma^2, \boldsymbol{\theta}, v_\theta, v_{\theta_s}, \vartheta_{\theta_s}, \mathcal{A}_\theta$ 
3: Specify hyperparameters:  $\vartheta_0, \vartheta_1, \varrho_0, \varrho_1, \mu, \lambda, v_\alpha, a_{\sigma^2}, b_{\sigma^2}$ 
4: for iteration  $m = 1, \dots, M$  do
5:   for  $i = 1, \dots, N$  do
6:     for  $j = 1, \dots, n_i$  do
7:       for  $k = 1, \dots, n_{ij}$  do
8:         Update  $\omega_{ijk} \sim \text{PG}(1, \psi_{ijk})$ , where  $\psi_{ijk} = f_{ij}(t_{ijk}) + \mu_U(t_{ijk})$  via Polson et al. (2013).
9:       end for
10:    end for
11:  end for
12:  for  $i = 1, \dots, N$  do
13:    for  $j = 1, \dots, n_i$  do
14:      Update cluster assignments  $\mathbf{c}_{ij}$  via Bassetti et al. (2020) and Algorithm 8 of Neal (2000).
15:    end for
16:  end for
17:  for  $d = 1, \dots, D$  do
18:    Update dishes  $\mathbf{d}_d$  via Bassetti et al. (2020) and Algorithm 8 of Neal (2000).
19:    Update  $\boldsymbol{\theta}_d, v_{\theta_d}, v_{\theta_{ds}}, \vartheta_{\theta_{ds}}, \mathcal{A}_{\theta_d}$  via Makalic and Schmidt (2015).
20:  end for
21:  Jointly update  $\boldsymbol{\alpha}$  and  $\boldsymbol{\gamma}$  with Between and Within Steps via Savitsky et al. (2011).
22:  Update regression coefficients in  $\mu_U(\cdot)$  and  $\mu_Y(\cdot)$ , if necessary.
23:  Update  $\sigma^2$  from Inverse-Gamma( $\bar{a}_{\sigma^2}, \bar{b}_{\sigma^2}$ ).
24: end for

```

times t_{ij} and t_{ijk} were sampled from a Uniform(0,1) without loss of generality. Low-frequency observations with active latent inclusion indicators (i.e., $\gamma_{ij} = 1$) were generated from $Y_{ij} \sim N(\alpha_{dij} + f(t_{ij}) + \mathbf{x}_{ij}\boldsymbol{\beta}_Y, \sigma^2)$, where $d \in \{1, 2\}$ with equal probability, $f(t_{ij}) = 0.5 + 0.2t_{ij} - (\pi/10)\sin(2\pi t_{ij})$, $\boldsymbol{\beta}_Y$ is a 2-dimensional vector of population-level regression coefficients randomly sampled from ± 0.75 , \mathbf{x}_{ij} represents two covariates simulated from a standard normal distribution, and $\sigma^2 = 0.2$. The cluster specific α_{1ij} and α_{2ij} were set to 1.5 and -1.5 , respectively. The corresponding high-frequency outcomes were generated from $U_{ijk} \sim \text{Bernoulli}(p_{ijk})$, where

$$p_{ijk} = \frac{\exp(f_{dij}(t_{ijk}) + \mathbf{x}_{ijk}\boldsymbol{\beta}_U)}{1 + \exp(f_{dij}(t_{ijk}) + \mathbf{x}_{ijk}\boldsymbol{\beta}_U)},$$

$f_{1ij}(t_{ijk}) = 1.5 - 0.5\cos(2t_{ijk})$, $f_{2ij}(t_{ijk}) = -1.1$, $\boldsymbol{\beta}_U$ is a 2-dimensional vector with elements set to ± 0.75 with equal probability, and \mathbf{x}_{ijk} represents two covariates simulated from a standard normal distribution. Note that by including a time-varying intercept term in the low-frequency model, in addition to the high-frequency model, and accommodating potential covariates at both levels of the model, we evaluate our approach for dynamic functional variable selection in a more general setting than our application study to demonstrate its flexibility in practice. Low-frequency observations with inactive latent inclusion indicators (i.e., $\gamma_{ij} = 0$) were simulated similar to the above with $\alpha_{dij} = 0$, and their corresponding high-frequency outcomes were generated with $f_{1ij}(t_{ijk}) = 0.5\cos(t_{ijk})$ and $f_{2ij}(t_{ijk}) = -1.0 - 1.5\cos(-t_{ijk})$. Thus, each observation-specific effect and corresponding trajectory could belong to one of four unique clusters. Additionally, we investigated the performance of our proposed method with varying

sample sizes. Treating the above simulation as a baseline setting, we varied the overall sample size, $N = \{5, 50\}$, adjusting the number of low-frequency observations to 100 and 1,000, respectively. We also investigated varying numbers of high-frequency measures, $n_{ij} = \{5, 50\}$.

In a second simulation study, we investigated the robustness of our proposed approach to model misspecification of the high-frequency trajectories by introducing corrupted data. Specifically, we simulated data similar to the above, but we flipped 5% and 10% of the high-frequency binary observations, mimicking a systematic error in the measurement device. We evaluated the models at various numbers of high-frequency observations, $n_{ij} = \{20, 50\}$.

We ran each of the MCMC algorithms on 30 replicated data sets for 1,000 iterations, treating the first 500 iterations as burn-in. We generated B-spline basis functions with five degrees of freedom to model $f(\cdot)$ and $f_{ij}(\cdot)$, where the Q -dimensional vector ζ represents the corresponding spline coefficients for $f(\cdot)$. Each chain was initiated with observations allocated to the same cluster with $\theta_{ij} = \mathbf{0}$ and $\alpha_{ij} = 0$. Regression coefficients β_Y , β_U , and ζ were also initialized at zero. All horseshoe variances and corresponding auxiliary parameters were initialized at one. We set the hyperparameters for the inclusion indicators $\lambda = \exp(-1.5)$ and $\mu = \exp(1.5)$ to represent ~ 0.05 steady state probability of experiencing a spike in the low-frequency outcome at a given moment. We assumed $\zeta_q \sim \text{Normal}(0, v_\zeta v_{\zeta_q} \sigma^2)$, where the global variances $v_\zeta \sim \text{Inverse-Gamma}(1/2, 1/\mathcal{A}_\zeta)$, local variances $v_{\zeta_q} \sim \text{Inverse-Gamma}(1/2, 1/\vartheta_{\zeta_q})$, and auxiliary parameters $\vartheta_{\zeta_1}, \dots, \vartheta_{\zeta_Q}, \mathcal{A}_\zeta \sim \text{Inverse-Gamma}(1/2, 1)$. We assumed the covariate regression coefficients $\beta_Y, \beta_U \sim \text{Normal}(0, v_\beta)$ with variance v_β . Additionally, we set $v_\beta = v_\alpha = 5$, which places a 95% prior probability of included regression coefficients between ± 4.4 . We specified the shape and scale parameters of the inverse-gamma distribution for σ^2 , $a_{\sigma^2} = b_{\sigma^2} = 1$. To complete the hyperparameter specification, we assumed the concentration parameters $\vartheta_0 = \vartheta_1 = 1$. For the PY priors, the discount parameters were set to 0.5 when appropriate.

Each of the models were evaluated in terms of their selection and clustering performance. Variable selection performance was evaluated via sensitivity (SENS), specificity (SPEC), and Matthew's correlation coefficient (MCC), defined as

$$\begin{aligned} SENS &= \frac{TP}{FN + TP}, & SPEC &= \frac{TN}{FP + TN}, \\ MCC &= \frac{TP \times TN - FP \times FN}{\sqrt{(TP + FP)(TP + FN)(TN + FP)(TN + FN)}}, \end{aligned}$$

where TN , TP , FN , and FP represent the true negatives, true positives, false negatives, and false positives, respectively. Active spikes were determined if their MPPI exceeded 0.50 (Barbieri et al., 2004). Clustering performance was evaluated using the variation of information (VI), a measure of distance between two clusterings ranging from 0 to $\log B$, where B is the number of items to cluster and lower values imply better clustering (Meilă, 2003). Additionally, we evaluated model fit for each of the models with the widely applicable information criterion (WAIC) for both the high-frequency binary

	SENS	SPEC	MCC	VI	WAIC _Y	WAIC _U
sDPDP	0.916	0.715	0.709	0.615	2253.6	8004.0
sDPPY	0.955	0.669	0.742	0.559	2593.6	7992.7
sPYDP	0.876	0.505	0.485	0.618	2596.1	8006.2
sPYPY	0.893	0.504	0.521	0.570	2512.8	7995.7

Table 1: Results of the proposed dynamic functional variable selection model with various sHSSP specifications for the $N = 20$, $n_i = 20$, and $n_{ij} = 20$ setting. Results are averaged over 30 replicate data sets.

outcomes (WAIC_U) and low-frequency continuous outcomes (WAIC_Y) (Watanabe and Opper, 2010). WAIC is an extension of the deviance information criterion (DIC) for Bayesian models that uses the entire posterior distribution to estimate the pointwise out-of-sample prediction accuracy of the model, and it is asymptotically equivalent to leave-one-out cross-validation (Watanabe and Opper, 2010). Lower values indicate better model fit. See the Supplementary Material for convergence assessment of the models.

Results of the simulation study with 20 subjects and 20 high-frequency measures are presented in Table 1. Overall, the models had relatively high sensitivity (SENS > 0.85) with moderate specificity (SPEC > 0.50). We found that the sDPPY model obtained the highest sensitivity, resulting in the highest overall performance with respect to MCC, and the best clustering performance compared to the other model specifications (VI = 0.559). The sDPDP model obtained the highest specificity overall (SPEC = 0.715). Clustering performance remained the same or improved with larger sample sizes for all models (Table 2). Similar trends in the overall variable selection performance were found for the sDPDP and sPYDP models. The sPYPY and sDPPY models were more robust to sample size in terms of selection performance. We found that the selection and clustering results were more sensitive to the number of high-frequency outcomes generated than the sample size (Table 3). For most models, we observed a reduction (improvement) in selection and clustering performance with respect to all metrics as the number of high-frequency measures decreased (increased). This was not surprising, as there was very little information available to identify the unique high-frequency trajectories in the $n_{ij} = 5$ setting. Notably, the $n_{ij} = 50$ setting showed almost perfect clustering performance for all models.

In the second simulation study, we observed that our proposed method with different sHSSPs is relatively robust to moderate amounts of measurement error in the high-frequency outcomes (Table 4). We observed a marginal decrease in overall selection and clustering performance for all models, except the sPYPY model, which maintained or even improved selection performance.

In each of the simulation scenarios, the models performed similarly with respect to WAIC for the high-frequency binary outcomes. While we did observe some marginal differences in fit for the low-frequency continuous outcomes across model specifications, we did not find any systematic trends with the exception that the sPYPY model performed the best in terms of WAIC_Y in the presence of measurement error (Table 4).

	N	SENS	SPEC	MCC	VI	WAIC _Y	WAIC _U
sDPDP		0.850	0.472	0.452	0.785	946.8	2030.5
sDPPY	5	0.971	0.577	0.760	0.607	729.1	2015.2
sPYDP		0.903	0.449	0.480	0.776	927.1	2026.9
sPYPY		0.896	0.586	0.551	0.691	867.5	2020.2
sDPDP		0.904	0.674	0.654	0.573	5704.0	19981.7
sDPPY	50	0.958	0.531	0.751	0.547	5675.5	19957.7
sPYDP		0.954	0.696	0.753	0.578	4898.8	19983.1
sPYPY		0.974	0.546	0.745	0.548	5647.0	19959.7

Table 2: Results of the proposed dynamic functional variable selection model with various sHSSP specifications for the $N = 5$ and 50, $n_i = 20$, and $n_{ij} = 20$ setting. Results are averaged over 30 replicate data sets.

	n_{ij}	SENS	SPEC	MCC	VI	WAIC _Y	WAIC _U
sDPDP		0.731	0.464	0.249	1.116	2647.4	2097.8
sDPPY	5	0.918	0.397	0.438	0.981	3056.4	2082.4
sPYDP		0.759	0.492	0.318	1.113	2592.2	2095.2
sPYPY		0.931	0.521	0.521	1.021	3422.0	2084.8
sDPDP		0.929	0.642	0.700	0.159	2499.9	19781.7
sDPPY	50	0.978	0.771	0.857	0.154	2079.6	19779.1
sPYDP		0.884	0.540	0.515	0.173	2104.6	19782.5
sPYPY		0.927	0.776	0.741	0.155	2107.3	19779.9

Table 3: Results of the proposed dynamic functional variable selection model with various sHSSP specifications for the $N = 20$, $n_i = 20$, and $n_{ij} = 5$ and 50 setting. Results are averaged over 30 replicate data sets.

	% Error	SENS	SPEC	MCC	VI	WAIC _Y	WAIC _U
sDPDP		0.904	0.676	0.655	0.724	2522.8	8129.2
sDPPY	5	0.937	0.560	0.675	0.683	2192.4	8112.0
sPYDP		0.834	0.345	0.257	0.719	4185.0	8130.2
sPYPY		0.967	0.719	0.811	0.678	2120.1	8116.6
sDPDP		0.872	0.548	0.504	0.836	3321.8	8316.6
sDPPY	10	0.919	0.537	0.593	0.787	2873.5	8291.1
sPYDP		0.862	0.488	0.443	0.838	2685.9	8316.4
sPYPY		0.954	0.494	0.653	0.788	2420.5	8297.6

Table 4: Results of the proposed dynamic functional variable selection model with various sHSSP specifications for the $N = 20$, $n_i = 20$, and $n_{ij} = 20$ setting with 5% and 10% error for the high-frequency binary observations. Results are averaged over 30 replicate data sets.

	$v_\alpha = 1$	$\vartheta_0 = 0.1$	$\vartheta_1 = 0.1$	$\varrho_0 = 0.25$	$\varrho_1 = 0.25$
SENS	0.984	0.944	0.938	0.953	0.908
SPEC	0.707	0.538	0.597	0.710	0.556
MCC	0.836	0.683	0.692	0.756	0.564
VI	0.562	0.567	0.572	0.559	0.584
WAIC _Y	2707.8	2458.5	2539.6	2179.2	2412.9
WAIC _U	7995.9	7994.6	7997.0	7993.4	8000.3
	$v_\alpha = 10$	$\vartheta_0 = 10$	$\vartheta_1 = 10$	$\varrho_0 = 0.75$	$\varrho_1 = 0.75$
SENS	0.856	0.983	0.941	0.969	0.937
SPEC	0.507	0.709	0.518	0.708	0.650
MCC	0.452	0.851	0.679	0.811	0.704
VI	0.562	0.570	0.562	0.555	0.561
WAIC _Y	3226.3	3341.5	2712.4	2811.6	2825.2
WAIC _U	7995.9	8004.6	7993.9	7996.2	8001.4

Table 5: Simulated Data: Sensitivity results for the proposed model with sPYPY using various hyperparameter specifications. Results are averaged over 30 replicated data sets.

3.1 Sensitivity Analysis

To assess the model’s sensitivity to hyperparameter settings, we set each of the hyperparameters to default values and then evaluated the effect of manipulating each term on selection and clustering performance. For the default parameterization, we set the hyperparameters for the prior inclusion indicators, γ , to $\mu = \exp(1.5)$ and $\lambda = \exp(-1.5)$, reflecting a ~ 0.05 steady state prior probability of experiencing a spike in the low-frequency outcome at a given assessment. The default values for the variance of the normal distribution for the slab of α_{ij} were each fixed at 5, and v_β was set similarly. The hyperparameters for the concentration parameters $\vartheta_0 = \vartheta_1 = 1$, and the discount parameters were set to 0.5. We ran the joint model on the 30 replicated data sets generated in the first simulation scenario.

The results of the sensitivity analysis are presented in Table 5. We observed that the proposed joint model was insensitive to hyperparameter specification in terms of clustering. We found that the selection results were not affected by an increase in the prior variance for the observation-specific effects, v_α . However, as v_α decreased, we observed improved selection performance, potentially an artifact of the data generation process. We found a marginal positive association between the hyperparameter specification of the sHSSP and WAIC_Y. However, the fit of the high-frequency binary outcome as measured by WAIC_U was not affected by hyperparameter specification.

4 Application

In this section, we illustrate the performance of the proposed joint model on data collected in the Pathways between Socioeconomic Status and Behavioral Cancer Risk Factors (Pathways) Study, which was a 7-day prospective observational study designed to characterize proximal predictors of health behaviors using mHealth methods in a racially and ethnically diverse community sample of adults. Full details of the Pathways design including recruitment, eligibility, data collection/preprocessing are found elsewhere (Kendzor et al., 2016). After baseline measures were collected, participants

were provided with a mobile phone, which they were asked to carry over a 7-day period. Within the assessment period, each participant was randomly prompted to complete EMAs four times per day during self-reported waking hours. Each assessment collected information on participants’ current affective status, environment, behaviors, loneliness, and stress. Additionally, each participant was fitted with an ActiGraph GT3X triaxial accelerometer to passively collect information on their ambulatory movement (i.e., time spent sedentary or active) during the assessment period.

In this analysis, we use the Pathways data to investigate the relations between negative affect actively collected with EMAs (low-frequency outcome) and accelerometer data passively captured minute-by-minute with accelerometers in proximal 30-minute epochs (high-frequency outcome). Negative affect was calculated as the average of a participant’s responses to five EMA items (i.e., “I feel irritable, frustrated/angry, sad, worried, and miserable”). Time spent sedentary has been linked to a variety of adverse health outcomes, including cancer (Friedenreich et al., 2021), sleep disturbance/insomnia (Yang et al., 2017), major depressive disorder (Schuch et al., 2017), all cause mortality, cardiovascular disease, diabetes (Rezende et al., 2014; Young et al., 2016), and dementia (Yan et al., 2020). There is growing evidence that prolonged sedentary time is a risk factor for mental health (Giurgiu et al., 2020). Due to differences in movement patterns across ages, it is important to consider age-specific criteria when processing accelerometer data (Migueles et al., 2017). In contrast to thresholds based on vertical count axis data, very little work has been done to establish and validate intensity-based thresholds based on vector magnitude count (triaxial) data used to differentiate sedentary from active time. To investigate the relation between negative affect and sedentary behavior, physical activity vector magnitudes were calculated using three-dimensional movement data and thresholded at 150 counts per minute (cpm), where $\text{cpm} < 150$ indicates sedentary behavior and $\text{cpm} \geq 150$ indicates physical activity at any intensity category, following Peterson et al. (2015) and the age-group criteria recommendations for adults (i.e., 19–59 years) defined in Migueles et al. (2017). To account for non-compliance, EMAs with missing negative affect responses or one hour periods of zero cpm vector magnitudes preceding the proximal EMA were removed from this analysis. Overall, 222 participants were investigated in this study, comprising over 4,500 EMAs and roughly 140,000 accelerometer minutes of data. On average (SD), 20.8 (5.6) EMAs and 625.4 (166.9) high-frequency accelerometer measures per participant were available for inference.

4.1 Model and Prior Specification

We use the following model and prior specifications for the application study. For the low-frequency outcome model (Section 2.2) in (2), we assume $\mu_Y(t_{ij}) = \beta_i$, where β_i represents the participant-specific mean negative affect. We assume the slab variance for the α_{ij} prior $v_\alpha = 5$, and the hyperparameters for the error term $a_{\sigma^2} = b_{\sigma^2} = 1$. In the high-frequency outcome model (Section 2.1) for physical activity, we approximate the smooth functions in (1) using B-spline basis functions with five degrees of freedom. Recall that the spline coefficients are penalized with horseshoe priors to encourage smoothness. For the sHSSP (Section 2.3), we assume a spiked hierarchical Dirichlet process (sDPDP) with $\vartheta_0 = \vartheta_1 = 1$. For the multistate Markov model prior assigned to the

dynamic inclusion indicators (Section 2.4), we set the hyperparameters $\lambda = \exp(-3)$ and $\mu = \exp(1.5)$, reflecting a ~ 0.01 steady state prior probability of experiencing a spike in negative affect at a given assessment. The MCMC algorithm was run for 5,000 iterations, treating the first 1,000 iterations as burn-in and was initiated with observations allocated to the same cluster with $\theta_{ij} = \mathbf{0}$ and $\alpha_{ij} = 0$. Participant-specific means, β , were also initialized at zero. All horseshoe variances and corresponding auxiliary parameters were initialized at one. Traceplots of the total number of clusters in the model and the number of clusters in the model with active spikes indicated good mixing and overall convergence. (See the Supplementary Material for more details). The average effective sample size across active parameters in the model was 2,500. Additionally, we initialized the model with subjects in their own cluster and corresponding cluster coefficients sampled from a standard normal distribution. Between the two MCMC chains initiated at different values, we observed potential scale reduction factors, \hat{R} , for α , σ^2 , θ , and β below 1.1 (Gelman et al., 1992), further demonstrating that the MCMC procedure was working properly and the chains converged. Selection of spikes in momentary negative affect were determined using the median model approach (i.e., $\text{MPPI} \geq 0.50$) (Barbieri et al., 2004). Clusters were determined using sequentially-allocated latent structure optimization to minimize the lower bound of the variation of information loss (Wade and Ghahramani, 2018; Dahl et al., 2017).

4.2 Results

Using our Bayesian dynamic functional variable selection method, we identified 439 moments (i.e., $\text{MPPI} \geq 0.5$) in which subjects departed from their average negative affect over the 7-day assessment window. We observed good separation between active and non-active observation-specific effects, where the average MPPI for active and non-active terms was around 0.83 and 0.08, respectively. A plot of the MPPIs for each observation is provided in the Supplementary Material. Figure 2 presents the observed negative affect and estimated momentary spikes versus time for a subset of participants. These plots illustrate the ability of our model to dynamically identify critical moments in which participants depart from their mean negative affect. In post-hoc analyses, we did not find strong evidence of temporal trends in the magnitude of the departures in negative affect over the assessment window. Additionally, most of the clusters were uniformly distributed across the assessment window, with exceptions for relatively small clusters.

Each of the active observation-specific effects in the low-frequency outcome model for negative affect was clustered with a unique trajectory in the high-frequency outcome model for activity. We identified 15 clusters with active observation-specific effects using the sDPDP model with an average (SD) of 39.9 (36.7) observations in each cluster. We observed a variety of cluster compositions with some comprised of observations from a small subset of participants and others with observations from nearly half of the participants. A plot of the clustered trajectories for the estimated probability of being active over the 30-minute epoch prior to the proximal EMA is presented in Figure 3. For comparison, the time scales are shifted to range from 0 to 30 minutes for each physical activity trajectory, however the actual time of measurement may have occurred at any

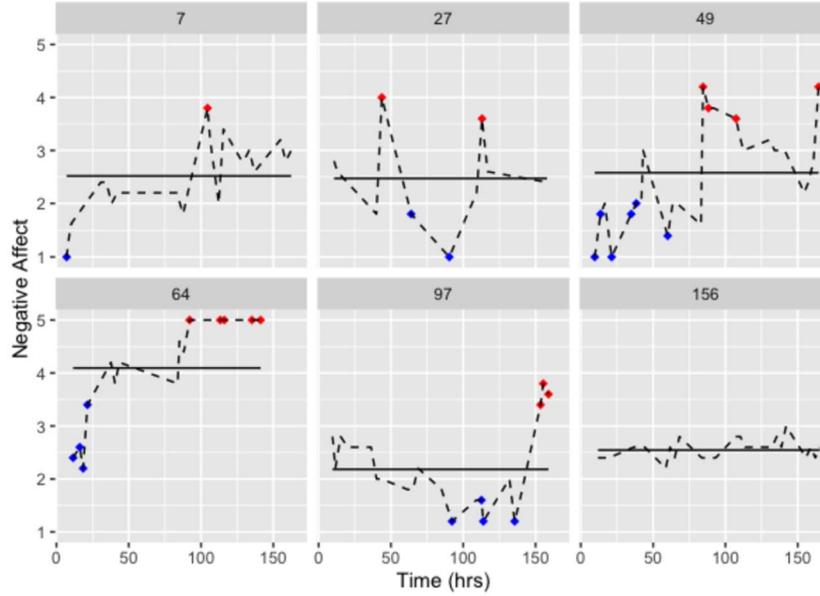


Figure 2: Application Study: Observed momentary negative affect (dashed lines) for a subset of participants with corresponding participant-specific means (solid line). Red (Blue) diamonds indicate moments in which participants experienced a positive (negative) spike in negative affect compared to their mean which were identified as associated with proximal physical activity trajectories by our model.

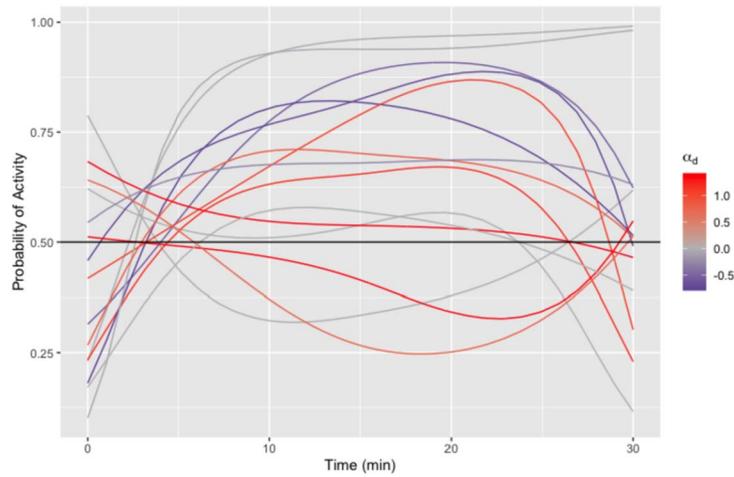


Figure 3: High-Frequency Outcome Trajectories: Clusters of the time-varying probability of being active over the 30-minute (min) epoch prior to a proximal EMA. Darker purple indicates trajectory clusters associated with negative spikes in momentary negative affect and brighter red indicates positive spikes in negative affect. Active clusters with 5 or more observations are plotted for ease of exposition.

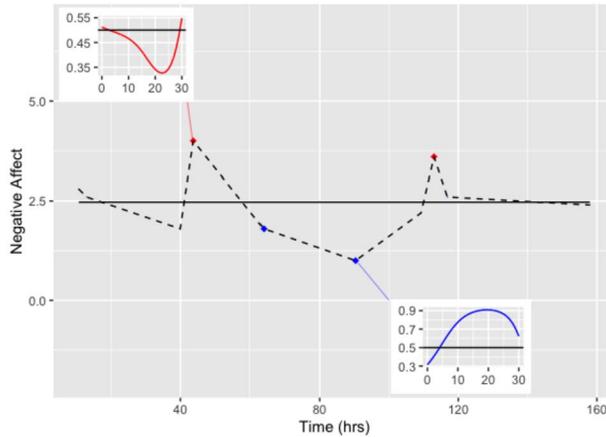


Figure 4: Participant-Level Dynamic Trajectory Selection: Observed momentary negative affect (dashed line) for a select participant compared to their average negative affect (solid line) over time in hours (hrs). Inset plots capture trajectories of the probability of being active for corresponding spikes in negative affect over the 30-minute epoch prior to the proximal EMA.

point over the 7-day assessment window. We found that trajectories associated with a decrease in proximal negative affect tended to have a higher probability of activity over the 30-minute epoch compared to trajectories associated with a positive spike. Most of the trajectories associated with a positive spike hovered around a 0.50 probability of activity over the 30-minute period.

To demonstrate the dynamic relation between momentary negative affect and proximal physical activity trajectories at the participant level, we investigated a select participant’s behavioral trends over the assessment period. Figure 4 depicts the observed negative affect trajectory (dashed line) of participant 27 over the course of the assessment period compared to their estimated average (solid line). We showcase two departures from their average negative affect, one on the second day of the assessment period and another on the fourth. On the second day, the individual experienced a positive spike in negative affect from their mean. This spike was associated with an activity trajectory that dropped from a 0.5 to 0.3 probability of being active in the first 25 minutes of the epoch and then increased to roughly a 0.5 probability of being active just before the proximal EMA. On the fourth day, a negative spike in negative affect was associated with an activity trajectory that increased from a 0.3 probability of being active to 0.9 by the 20th minute, before decreasing to 0.7 prior to the assessment.

Our joint modeling approach is designed to leverage the high temporal resolution of multimodal technology-based assessment strategies to provide insights on the dynamic relation between physical activity and negative affect. Our analysis revealed dynamic relations between accelerometer-measured physical activity over 30 minutes before EMA-measured negative affect. Findings indicated that higher activity level trajectories prior

to the proximal EMA were typically associated with subsequent ratings of participants' negative affect that were lower than their average. We also found that clusters were often comprised of observations on the same participant or small subpopulations of participants. These results emphasize the potential utility of developing personalized just-in-time adaptive interventions focused on negative affect management, where moments of increased negative affect are identified and targeted with activity prompts and/or other motivational or educational messaging via smartphone, smartwatch, or other technologies.

4.3 Sensitivity Analysis

HSSPs provide a flexible framework for nonparametric clustering at multiple levels without having to specify the number of clusters a priori. As such, one of the strengths of this class of priors is its ability take on characteristics of the assumed $SSrp$ at each level of the model. A consequence of this flexibility is that results may be sensitive to the assumed hierarchical structure. To explore the model's sensitivity to the specified sHSSP on the application study results, we fit a similar joint model for dynamic functional variable selection, but we varied the $SSrp$ specifications at both levels of the model. Holding $\vartheta_0 = \vartheta_1 = 1$, we investigated a Pitman-Yor process at both levels ($sPYPY(\varrho_0, \vartheta_0, \varrho_1, \vartheta_1, G_0)$, $\varrho_0 = \varrho_1 = 0.5$), a Pitman-Yor-Dirichlet process ($sPYDP(\varrho_0, \vartheta_0, \vartheta_1, G_0)$, $\varrho_0 = 0.5$), and a Dirichlet-Pitman-Yor process ($sDPPY(\vartheta_0, \varrho_1, \vartheta_1, G_0)$, $\varrho_1 = 0.5$).

Compared to the sDPDP model in Section 4, the sDPPY and sPYPY models identified more spikes in momentary negative affect (i.e., 865 and 560, respectively), whereas the sPYDP model only identified 305 using the median model approach. As expected, the sDPPY, sPYDP, and sPYPY models identified more active clusters (i.e., 45, 24, and 48, respectively) than the sDPDP model. A majority of the observations were clustered into similar groups across model specifications. See Figure 5 for a heatmap of the ad-

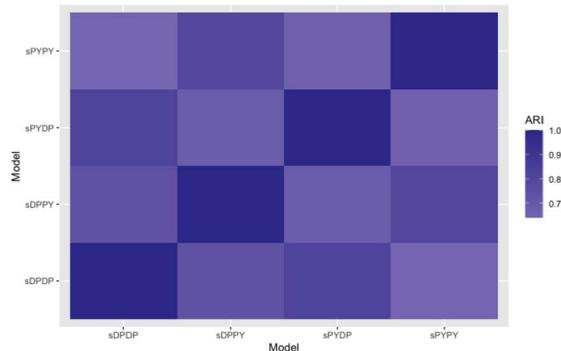


Figure 5: Clustering Allocation Sensitivity: Heatmap of the adjusted Rand index (ARI) for the pairwise comparison of clustering allocation by the spiked hierarchical Dirichlet process by the spiked hierarchical Dirichlet process (sDPDP), Dirichlet-Pitman-Yor process (sDPPY), Pitman-Yor-Dirichlet process (sPYDP), and the hierarchical Pitman-Yor process (sPYPY) models.

justed Rand index, which measures cluster similarity for each pairwise combination of model specifications and adjusts for the expected number of chance agreements (Hubert and Arabie, 1985). The adjusted Rand index ranges from 0 to 1, with 1 representing that the clustering allocations are the same. Overall, we observed that models with similar prior specifications at the first level of the hierarchy tended to cluster more similarly as well. The sDPDP and sPYDP models were the most similar in terms of clustering overlap (ARI = 0.82), and the sPYDP and sPYPY models were the most dissimilar (ARI = 0.66). Lastly, we compared the fit of each model using WAIC for both the high- and low-frequency outcomes. We found that the sDPDP model best fit the low-frequency continuous outcomes (WAIC_Y = 193449) with relatively similar performance for the sPYDP model (WAIC_Y = 210768). The sDPY and sPYPY models obtained a poorer fit (WAIC_Y = 313687 and 318409, respectively). For the high-frequency binary outcomes, we found a similar fit for the sDPDP, sPYDP, and sDPY models (WAIC_U = 1162113, 1152028, and 1159852, respectively) with the sPYPY model obtaining the best fit (WAIC_U = 1069216). We note that the other prior distribution hyperparameter settings of the model were consistent with typical specifications in the literature, and the choice of these values had minimal effects on inference in the application study.

5 Conclusions

In this work, we developed the first method for dynamic functional variable selection. We applied the proposed method to investigate the relations between two actively and passively assessed health-related variables, a fundamental research objective in behavioral mHealth studies (Walls, 2013). Motivated by data collected in the Pathways study, our approach is designed to accommodate multimodal intensive longitudinal data collected on different time scales. To achieve this, we employed spiked hierarchical species sampling priors to identify critical moments that a participant departs from their average negative affect that were associated with physical activity trajectories observed in proximal epochs. Our hierarchical approach is designed to deliver unique insights into the dynamic behavioral patterns shared within and between subjects, critical information for the design and evaluation of personalized intervention strategies using mHealth methods. While the results highlight the importance of understanding the dynamic relation between physical activity trajectories and proximal negative affect, we recommend using our method in exploratory settings for hypothesis generation and conducting confirmatory studies before generalizing results.

The proposed method is designed for continuous low-frequency outcomes and binary high-frequency outcomes given the structure of the application data set. However, it can be easily adjusted to accommodate continuous (binary) high- (low-) frequency outcomes via the connection between the Pólya-Gamma data augmentation technique of Polson et al. (2013) and a Gaussian likelihood function. As such, our approach is easily modified to handle a variety of settings in which multivariate intensive longitudinal data are collected on different time scales. In the FDA literature space, sparsity or variable selection often refers to identifying portions of the domain that retain relevant information for a functional trajectory, balancing between representativeness and

parsimony (Berrendero et al., 2016), or capture significant differences between different groups or unobserved clusters (Vitelli, 2019; Pini and Vantini, 2017). In this work, the proposed method is designed to identify the critical moments that participants’ depart from their baseline average of the low-frequency outcome which are associated with proximal high-frequency outcome trajectories. A future extension would be to develop the model to identify intervals of time in which the participant departs or remains at their low-frequency baseline average. In extreme high-frequency settings (i.e., n_{ij} large), the likelihood contribution of the data is dominated by the high-frequency model when determining cluster allocation. As a result, the model may achieve good clustering performance but fail to correctly identify active/inactive observation-specific effects due to poor mixing. While our method was relatively robust to measurement error in the simulation study, future work may explore robustification methods, such as the coarsened-posterior method, which in addition to providing robust inference, may improve mixing (Miller and Dunson, 2018; Geyer and Thompson, 1995).

References

- Anglada-Martinez, H., Riu-Viladoms, G., Martin-Conde, M., Rovira-Illamola, M., Sotoca-Momblona, J., and Codina-Jane, C. (2015). “Does mHealth increase adherence to medication? Results of a systematic review.” *International Journal of Clinical Practice*, 69(1): 9–32. 2
- Barbieri, M. M., Berger, J. O., et al. (2004). “Optimal predictive model selection.” *The Annals of Statistics*, 32(3): 870–897. MR2065192. doi: <https://doi.org/10.1214/009053604000000238>. 9, 11, 16
- Bassetti, F., Casarin, R., and Rossini, L. (2020). “Hierarchical species sampling models.” *Bayesian Analysis*, 15(3): 809–838. MR4132651. doi: <https://doi.org/10.1214/19-BA1168>. 6, 7, 10
- Berrendero, J. R., Cuevas, A., and Torrecilla, J. L. (2016). “Variable selection in functional data classification: a maxima-hunting proposal.” *Statistica Sinica*, 619–638. MR3497763. 21
- Bertz, J. W., Epstein, D. H., and Preston, K. L. (2018). “Combining ecological momentary assessment with objective, ambulatory measures of behavior and physiology in substance-use research.” *Addictive Behaviors*, 83: 5–17. 2
- Bigelow, J. L. and Dunson, D. B. (2009). “Bayesian semiparametric joint models for functional predictors.” *Journal of the American Statistical Association*, 104(485): 26–36. MR2663031. doi: <https://doi.org/10.1198/jasa.2009.0001>. 3
- Blaauw, F. J., Schenk, H. M., Jeronimus, B. F., van der Krieke, L., de Jonge, P., Aiello, M., and Emerencia, A. C. (2016). “Let’s get Physiquel—An intuitive and generic method to combine sensor technology with ecological momentary assessments.” *Journal of Biomedical Informatics*, 63: 141–149. 2
- Brown, P. J., Vannucci, M., and Fearn, T. (1998). “Multivariate Bayesian variable selection and prediction.” *Journal of the Royal Statistical Society: Series B (Statistical*

- Methodology*), 60(3): 627–641. MR1626005. doi: <https://doi.org/10.1111/1467-9868.00144>. 6, 7
- Businelle, M. S., Ma, P., Kendzor, D. E., Frank, S. G., Vidrine, D. J., and Wetter, D. W. (2016). “An ecological momentary intervention for smoking cessation: Evaluation of feasibility and effectiveness.” *Journal of Medical Internet Research*, 18(12): e321. 2
- Businelle, M. S., Perski, O., Hébert, E. T., and Kendzor, D. E. (2024). “Mobile Health Interventions for Substance Use Disorders.” *Annual Review of Clinical Psychology*, 20(1). doi: <https://doi.org/10.1146/annurev-clinpsy-080822-042337>. 2
- Canale, A., Lijoi, A., Nipoti, B., and Prünster, I. (2017). “On the Pitman–Yor process with spike and slab base measure.” *Biometrika*, 104(3): 681–697. MR3694590. doi: <https://doi.org/10.1093/biomet/asx041>. 4, 7
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). “The horseshoe estimator for sparse signals.” *Biometrika*, 97(2): 465–480. MR2650751. doi: <https://doi.org/10.1093/biomet/asq017>. 5
- Cassese, A., Zhu, W., Guindani, M., and Vannucci, M. (2019). “A Bayesian nonparametric spiked process prior for dynamic model selection.” *Bayesian Analysis*, 14(2): 553–572. MR3934097. doi: <https://doi.org/10.1214/18-BA1116>. 4, 7
- Cushing, C. C., Mitchell, T. B., Bejarano, C. M., Walters, R. W., Crick, C. J., and Noser, A. E. (2017). “Bidirectional associations between psychological states and physical activity in adolescents: A mHealth pilot study.” *Journal of Pediatric Psychology*, 42(5): 559–568. 3
- Dahl, D. B., Day, R., and Tsai, J. W. (2017). “Random partition distribution indexed by pairwise information.” *Journal of the American Statistical Association*, 112(518): 721–732. MR3671765. doi: <https://doi.org/10.1080/01621459.2016.1165103>. 16
- Dahl, D. B., Johnson, D. J., Müller, P., and Dahl, M. D. B. (2021). “Package ‘salso’” 9
- Das, K., Ghosh, P., and Daniels, M. J. (2021). “Modeling multiple time-varying related groups: A dynamic hierarchical Bayesian approach with an application to the Health and Retirement Study.” *Journal of the American Statistical Association*, 1–11. MR4270003. doi: <https://doi.org/10.1080/01621459.2021.1886105>. 6
- de Brito, J. N., Loth, K. A., Tate, A., and Berge, J. M. (2020). “Associations between parent self-reported and accelerometer-measured physical activity and sedentary time in children: Ecological momentary assessment study.” *JMIR mHealth and uHealth*, 8(5): e15458. 3
- Dunson, D. B., Herring, A. H., and Engel, S. M. (2008). “Bayesian selection and clustering of polymorphisms in functionally related genes.” *Journal of the American Statistical Association*, 103(482): 534–546. MR2523991. doi: <https://doi.org/10.1198/016214507000000554>. 7
- Dziak, J. J., Li, R., Tan, X., Shiffman, S., and Shiyko, M. P. (2015). “Modeling inten-

- sive longitudinal data with mixtures of nonparametric trajectories and time-varying effects." *Psychological Methods*, 20(4): 444. 3, 4
- Dzubur, E., Ponnada, A., Nordgren, R., Yang, C.-H., Intille, S., Dunton, G., and Hedeker, D. (2020). "MixWILD: A program for examining the effects of variance and slope of time-varying variables in intensive longitudinal data." *Behavior Research Methods*, 52(4): 1403–1427. 3
- Ferguson, T. S. (1973). "A Bayesian analysis of some nonparametric problems." *The Annals of Statistics*, 209–230. MR0350949. 6
- Fitzmaurice, G., Davidian, M., Verbeke, G., and Molenberghs, G. (2008). *Longitudinal data analysis*. CRC press. MR1500110. 3
- Friedenreich, C. M., Ryder-Burbidge, C., and McNeil, J. (2021). "Physical activity, obesity and sedentary behavior in cancer etiology: Epidemiologic evidence and biologic mechanisms." *Molecular Oncology*, 15(3): 790–800. 15
- Gelman, A., Rubin, D. B., et al. (1992). "Inference from iterative simulation using multiple sequences." *Statistical Science*, 7(4): 457–472. 16
- George, E. I. and McCulloch, R. E. (1997). "Approaches for Bayesian variable selection." *Statistica Sinica*, 339–373. 6, 7, 9
- Geyer, C. J. and Thompson, E. A. (1995). "Annealing Markov chain Monte Carlo with applications to ancestral inference." *Journal of the American Statistical Association*, 90(431): 909–920. MR1341319. doi: <https://doi.org/10.2307/1390763>. 21
- Ghosal, S. and Van der Vaart, A. (2017). *Fundamentals of Nonparametric Bayesian Inference*, volume 44. Cambridge University Press. MR3587782. doi: <https://doi.org/10.1017/9781139029834>. 7
- Giurgiu, M., Koch, E. D., Plotnikoff, R. C., Ebner-Priemer, U. W., and Reichert, M. (2020). "Breaking up sedentary behavior optimally to enhance mood." *Medicine & Science in Sports & Exercise*, 52(2): 457–465. 15
- Goldsmith, J. and Schwartz, J. E. (2017). "Variable selection in the functional linear concurrent model." *Statistics in Medicine*, 36(14): 2237–2250. MR3660128. doi: <https://doi.org/10.1002/sim.725>. 4
- Hastie, T. and Tibshirani, R. (1993). "Varying-coefficient models." *Journal of the Royal Statistical Society: Series B (Methodological)*, 55(4): 757–779. MR1229881. 3
- Hébert, E. T., Ra, C. K., Alexander, A. C., Helt, A., Moisiuc, R., Kendzor, D. E., Vidrine, D. J., Funk-Lawler, R. K., and Businelle, M. S. (2020). "A Mobile Just-in-Time Adaptive Intervention for Smoking Cessation: Pilot Randomized Controlled Trial." *Journal of Medical Internet Research*, 22(3): e16907. 2
- Heron, K. E. and Smyth, J. M. (2010). "Ecological momentary interventions: Incorporating mobile technology into psychosocial and health behaviour treatments." *British Journal of Health Psychology*, 15(1): 1–39. 2
- Hjort, N. L., Holmes, C., Müller, P., and Walker, S. G. (2010). *Bayesian nonparamet-*

- rics*, volume 28. Cambridge University Press. MR2722988. doi: <https://doi.org/10.1017/CB09780511802478.002>. 7
- Hubert, L. and Arabie, P. (1985). “Comparing partitions.” *Journal of Classification*, 2(1): 193–218. 20
- Islam, M. N., Stallings, J., Staicu, A.-M., Crouch, D., Pan, L., and Huang, H. (2018). “Functional Variable Selection for EMG-based Control of a Robotic Hand Prosthetic.” *arXiv preprint arXiv:1805.03098*. 4
- Jackson, C. (2011). “Multi-state models for panel data: The msm package for R.” *Journal of Statistical Software*, 38(1): 1–28. 9
- Kazemi, D. M., Borsari, B., Levine, M. J., Li, S., Lamberson, K. A., and Matta, L. A. (2017). “A systematic review of the mHealth interventions to prevent alcohol and substance abuse.” *Journal of Health Communication*, 22(5): 413–432. 2
- Kendzor, D. E., Shuval, K., Gabriel, K. P., Businelle, M. S., Ma, P., High, R. R., Cuate, E. L., Poonawalla, I. B., Rios, D. M., Demark-Wahnefried, W., et al. (2016). “Impact of a mobile phone intervention to reduce sedentary behavior in a community sample of adults: A quasi-experimental evaluation.” *Journal of Medical Internet Research*, 18(1): e19. 2, 14
- Kim, J., Marcusson-Clavertz, D., Togo, F., and Park, H. (2018). “A practical guide to analyzing time-varying associations between physical activity and affect using multi-level modeling.” *Computational and Mathematical Methods in Medicine*, 2018. 2
- Kim, S., Dahl, D. B., and Vannucci, M. (2009). “Spiked Dirichlet process prior for Bayesian multiple hypothesis testing in random effects models.” *Bayesian Analysis (Online)*, 4(4): 707. MR2570085. doi: <https://doi.org/10.1214/09-BA426>. 4, 7
- Koslovsky, M. D., Hébert, E. T., Businelle, M. S., and Vannucci, M. (2020). “A Bayesian time-varying effect model for behavioral mHealth data.” *The Annals of Applied Statistics*, 14(4): 1878–1902. MR4194252. doi: <https://doi.org/10.1214/20-AOAS1402>. 3, 4, 7
- Koslovsky, M. D., Hébert, E. T., Swartz, M. D., Chan, W., Leon-Novelo, L., Wilkinson, A. V., Kendzor, D. E., and Businelle, M. S. (2018a). “The time-varying relations between risk factors and smoking before and after a quit attempt.” *Nicotine and Tobacco Research*, 20(10): 1231–1236. 3
- Koslovsky, M. D., Swartz, M. D., Chan, W., Leon-Novelo, L., Wilkinson, A. V., Kendzor, D. E., and Businelle, M. S. (2018b). “Bayesian variable selection for multistate Markov models with interval-censored data in an ecological momentary assessment study of smoking cessation.” *Biometrics*, 74(2): 636–644. MR3825350. doi: <https://doi.org/10.1111/biom.12792>. 8
- Kowal, D. R., Matteson, D. S., and Ruppert, D. (2019). “Dynamic shrinkage processes.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(4): 781–804. MR3997101. doi: <https://doi.org/10.1111/rssb.12325>. 4
- Kumar, S., Nilsen, W. J., Abernethy, A., Atienza, A., Patrick, K., Pavel, M., Riley,

- W. T., Shar, A., Spring, B., Spruijt-Metz, D., et al. (2013). “Mobile health technology evaluation: The mHealth evidence workshop.” *American Journal of Preventive Medicine*, 45(2): 228–236. 2
- Kürüm, E., Li, R., Shiffman, S., and Yao, W. (2016). “Time-varying coefficient models for joint modeling binary and continuous outcomes in longitudinal data.” *Statistica Sinica*, 26(3): 979. MR3559939. 3
- Li, F. and Zhang, N. R. (2010). “Bayesian variable selection in structured high-dimensional covariate spaces with applications in genomics.” *Journal of the American Statistical Association*, 105(491): 1202–1214. MR2752615. doi: <https://doi.org/10.1198/jasa.2010.tm08177>. 7
- Liang, M., Koslovsky, M. D., Hébert, E. T., Businelle, M. S., and Vannucci, M. (2023). “Functional concurrent regression mixture models using spiked Ewens-Pitman attraction priors.” *Bayesian Analysis*, 1(1): 1–29. 3
- Liang, M., Koslovsky, M. D., Hébert, E. T., Kendzor, D. E., Businelle, M. S., and Vannucci, M. (2021). “Bayesian continuous-time hidden Markov models with covariate selection for intensive longitudinal data with measurement error.” *Psychological Methods*. 8
- Liu, Y.-Y., Moreno, A., Li, S., Li, F., Song, L., and Rehg, J. M. (2017). “Learning continuous-time hidden markov models for event data.” In *Mobile Health*, 361–387. Springer. 8
- Makalic, E. and Schmidt, D. F. (2015). “A simple sampler for the horseshoe estimator.” *IEEE Signal Processing Letters*, 23(1): 179–182. 5, 10
- Meilă, M. (2003). “Comparing clusterings by the variation of information.” In *Learning Theory and Kernel Machines*, 173–187. Springer. 11
- Migueles, J. H., Cadenas-Sanchez, C., Ekelund, U., Delisle Nyström, C., Mora-Gonzalez, J., Löf, M., Labayen, I., Ruiz, J. R., and Ortega, F. B. (2017). “Accelerometer data collection and processing criteria to assess physical activity and other outcomes: A systematic review and practical considerations.” *Sports Medicine*, 47(9): 1821–1845. 15
- Miller, J. W. and Dunson, D. B. (2018). “Robust Bayesian inference via coarsening.” *Journal of the American Statistical Association*. MR4011766. doi: <https://doi.org/10.1080/01621459.2018.1469995>. 21
- Mitchell, H. B. (2007). *Multi-sensor Data Fusion: An Introduction*. Springer Science & Business Media. 2
- Morris, J. S. (2015). “Functional regression.” *Annual Review of Statistics and Its Application*, 2: 321–359. 5
- Müller, P. and Quintana, F. A. (2004). “Nonparametric Bayesian data analysis.” *Statistical Science*, 95–110. MR2082149. doi: <https://doi.org/10.1214/08834230400000017>. 7

- Nahum-Shani, I., Smith, S. N., Spring, B. J., Collins, L. M., Witkiewitz, K., Tewari, A., and Murphy, S. A. (2017). “Just-in-time adaptive interventions (JITAIs) in mobile health: Key components and design principles for ongoing health behavior support.” *Annals of Behavioral Medicine*, 52(6): 446–462. 2
- Neal, R. M. (2000). “Markov chain sampling methods for Dirichlet process mixture models.” *Journal of Computational and Graphical Statistics*, 9(2): 249–265. MR1823804. doi: <https://doi.org/10.2307/1390653>. 10
- Nelson, B. W. and Allen, N. B. (2018). “Extending the passive-sensing toolbox: Using smart-home technology in psychological science.” *Perspectives on Psychological Science*, 13(6): 718–733. 2
- Newton, M. A., Noueir, A., Sarkar, D., and Ahlquist, P. (2004). “Detecting differential gene expression with a semiparametric hierarchical mixture method.” *Biostatistics*, 5(2): 155–176. 9
- Peterson, N. E., Sirard, J. R., Kulbok, P. A., DeBoer, M. D., and Erickson, J. M. (2015). “Inclinometer Validation and Sedentary Threshold Evaluation in University Students.” *Research in Nursing & Health*, 38(6): 492. 15
- Pini, A. and Vantini, S. (2017). “Interval-wise testing for functional data.” *Journal of Nonparametric Statistics*, 29(2): 407–424. MR3635020. doi: <https://doi.org/10.1080/10485252.2017.1306627>. 21
- Pinsky, M. and Karlin, S. (2010). *An Introduction to Stochastic Modeling*. Academic press. MR3293984. doi: <https://doi.org/10.1016/B978-0-12-381416-6.00001-0>. 8
- Pitman, J. (1996). “Some developments of the Blackwell-MacQueen urn scheme.” *Lecture Notes-Monograph Series*, 245–267. MR1481784. doi: <https://doi.org/10.1214/lnms/1215453576>. 6
- Pitman, J. and Yor, M. (1997). “The two-parameter Poisson-Dirichlet distribution derived from a stable subordinator.” *The Annals of Probability*, 855–900. MR1434129. doi: <https://doi.org/10.1214/aop/1024404422>. 6
- Polson, N. G., Scott, J. G., and Windle, J. (2013). “Bayesian inference for logistic models using Pólya–Gamma latent variables.” *Journal of the American Statistical Association*, 108(504): 1339–1349. MR3174712. doi: <https://doi.org/10.1080/01621459.2013.829001>. 10, 20
- Ramsay, J. O. and Silverman, B. W. (2002). *Applied Functional Data Analysis: Methods and Case Studies*, volume 77. Springer. MR1910407. doi: <https://doi.org/10.1007/b98886>. 3, 5
- Ray, S. and Mallick, B. (2006). “Functional clustering by Bayesian wavelet methods.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 68(2): 305–332. MR2188987. doi: <https://doi.org/10.1111/j.1467-9868.2006.00545.x>. 6
- Rehg, J. M., Murphy, S. A., and Kumar, S. (2017). *Mobile Health*. Springer. 2

- Rezende, L. F. M. d., Rodrigues Lopes, M., Rey-López, J. P., Matsudo, V. K. R., and Luiz, O. d. C. (2014). “Sedentary behavior and health outcomes: An overview of systematic reviews.” *PloS one*, 9(8): e105620. 15
- Rizopoulos, D. and Lesaffre, E. (2014). “Introduction to the special issue on joint modelling techniques.” *Statistical Methods in Medical Research*, 23(1): 3–10. MR3190684. doi: <https://doi.org/10.1177/0962280212445800>. 3
- Rockova, V. and McAlinn, K. (2021). “Dynamic variable selection with spike-and-slab process priors.” *Bayesian Analysis*, 16(1): 233–269. MR4194280. doi: <https://doi.org/10.1214/20-BA1199>. 4
- Rodríguez, A., Dunson, D. B., and Gelfand, A. E. (2009). “Bayesian nonparametric functional data analysis through density estimation.” *Biometrika*, 96(1): 149–162. MR2482141. doi: <https://doi.org/10.1093/biomet/asn054>. 5
- Savitsky, T. and Vannucci, M. (2010). “Spiked Dirichlet process priors for Gaussian process models.” *Journal of Probability and Statistics*, 2010. MR2745498. doi: <https://doi.org/10.1155/2010/201489>. 4, 7
- Savitsky, T., Vannucci, M., and Sha, N. (2011). “Variable selection for nonparametric Gaussian process priors: Models and computational strategies.” *Statistical Science: A Review Journal of the Institute of Mathematical Statistics*, 26(1): 130. MR2849913. doi: <https://doi.org/10.1214/11-STS354>. 10
- Sayers, A., Heron, J., Smith, A. D., Macdonald-Wallis, C., Gilthorpe, M., Steele, F., and Tilling, K. (2017). “Joint modelling compared with two stage methods for analysing longitudinal data and prospective outcomes: A simulation study of childhood growth and BP.” *Statistical Methods in Medical Research*, 26(1): 437–452. MR3592734. doi: <https://doi.org/10.1177/0962280214548822>. 3
- Scherer, E. A., Ben-Zeev, D., Li, Z., and Kane, J. M. (2017). “Analyzing mHealth engagement: Joint models for intensively collected user engagement data.” *JMIR mHealth and uHealth*, 5(1): e1. 3
- Schuch, F., Vancampfort, D., Firth, J., Rosenbaum, S., Ward, P., Reichert, T., Bagatini, N. C., Bgeginski, R., and Stubbs, B. (2017). “Physical activity and sedentary behavior in people with major depressive disorder: A systematic review and meta-analysis.” *Journal of Affective Disorders*, 210: 139–150. 15
- Shi, J. Q. and Choi, T. (2011). *Gaussian process regression analysis for functional data*. CRC press. MR2932935. 5
- Stingo, F. C., Chen, Y. A., Vannucci, M., Barrier, M., and Mirkes, P. E. (2010). “A Bayesian graphical modeling approach to microRNA regulatory network inference.” *The Annals of Applied Statistics*, 4(4): 2024. MR2829945. doi: <https://doi.org/10.1214/10-AOAS360>. 7
- Suarez, A. J. and Ghosal, S. (2016). “Bayesian clustering of functional data using local features.” *Bayesian Analysis*, 11(1): 71–98. MR3447092. doi: <https://doi.org/10.1214/14-BA925>. 6

- Tan, X., Shiyko, M. P., Li, R., Li, Y., and Dierker, L. (2012). “A time-varying effect model for intensive longitudinal data.” *Psychological Methods*, 17(1): 61. 3, 4
- Teh, Y. W., Jordan, M. I., Beal, M. J., and Blei, D. M. (2006). “Hierarchical Dirichlet processes.” *Journal of the American Statistical Association*, 101(476): 1566–1581. MR2279480. doi: <https://doi.org/10.1198/016214506000000302>. 7
- Tsiatis, A. A. and Davidian, M. (2004). “Joint modeling of longitudinal and time-to-event data: An overview.” *Statistica Sinica*, 809–834. MR2087974. 3
- Vitelli, V. (2019). “A novel framework for joint sparse clustering and alignment of functional data.” *arXiv preprint arXiv:1912.00687*. MR4706179. doi: <https://doi.org/10.1080/10485252.2023.2206499>. 21
- Wade, S. and Ghahramani, Z. (2018). “Bayesian cluster analysis: Point estimation and credible balls (with discussion).” *Bayesian Analysis*, 13(2): 559–626. MR3807860. doi: <https://doi.org/10.1214/17-BA1073>. 9, 16
- Walls, T. A. (2013). “Intensive longitudinal data.” *The Oxford Handbook of Quantitative Methods: Statistical Analysis*, 2: 432–440. 4, 20
- Walls, T. A. and Schafer, J. L. (2006). *Models for intensive longitudinal data*. Oxford University Press. MR2334216. doi: <https://doi.org/10.1093/acprof:oso/9780195173444.001.0001>. 2
- Walsh, J. C., Corbett, T., Hogan, M., Duggan, J., and McNamara, A. (2016). “An mHealth intervention using a smartphone app to increase walking behavior in young adults: A pilot study.” *JMIR mHealth and uHealth*, 4(3): e5227. 2
- Watanabe, S. and Opper, M. (2010). “Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory.” *Journal of Machine Learning Research*, 11(12): MR2756194. 12
- White, P. A. and Gelfand, A. E. (2020). “Multivariate functional data modeling with time-varying clustering.” *TEST*, 1–17. MR4297269. doi: <https://doi.org/10.1007/s11749-020-00733-z>. 6
- Yan, S., Fu, W., Wang, C., Mao, J., Liu, B., Zou, L., and Lv, C. (2020). “Association between sedentary behavior and the risk of dementia: A systematic review and meta-analysis.” *Translational Psychiatry*, 10(1): 1–8. 15
- Yang, Y., Shin, J. C., Li, D., and An, R. (2017). “Sedentary behavior and sleep problems: A systematic review and meta-analysis.” *International Journal of Behavioral Medicine*, 24(4): 481–492. 15
- Young, D. R., Hivert, M.-F., Alhassan, S., Camhi, S. M., Ferguson, J. F., Katzmarzyk, P. T., Lewis, C. E., Owen, N., Perry, C. K., Siddique, J., et al. (2016). “Sedentary behavior and cardiovascular morbidity and mortality: A science advisory from the American Heart Association.” *Circulation*, 134(13): e262–e279. 15
- Zapata-Lamana, R., Lalanza, J. F., Losilla, J.-M., Parrado, E., and Capdevila, L. (2020). “MHealth technology for ecological momentary assessment in physical activity research: A systematic review.” *PeerJ*, 8: e8848. 2