Published in final edited form as:

Psychometrika. 2018 September ; 83(3): 733–750. doi:10.1007/s11336-017-9594-5.

DETECTING MULTIPLE RANDOM CHANGEPOINTS IN BAYESIAN PIECEWISE GROWTH MIXTURE MODELS

Eric F. Lock,

Division of Biostatistics, School of Public Health, University of Minnesota

Nidhi Kohli, and

Department of Educational Psychology, College of Education & Human Development, University of Minnesota

Maitreyee Bose

Division of Biostatistics, School of Public Health, University of Minnesota

Abstract

Piecewise growth mixture models (PGMM) are a flexible and useful class of methods for analyzing segmented trends in individual growth trajectory over time, where the individuals come from a mixture of two or more latent classes. These models allow each segment of the overall developmental process within each class to have a different functional form; examples include two linear phases of growth, or a quadratic phase followed by a linear phase. The *change point (knot)* is the time of transition from one developmental phase (segment) to another. Inferring the location of the changepoint(s) is often of practical interest, along with inference for other model parameters. A random changepoint allows for individual differences in the transition time within each class. The primary objectives of our study are: (1) to develop a PGMM using a Bayesian inference approach that allows the estimation of *multiple* random changepoints within each class; (2) to develop a procedure to empirically *detect the number of random changepoints* within each class; and (3) to empirically investigate the bias and precision of the estimation of the model parameters, including the random changepoints, via a simulation study. We have developed the user-friendly package *BayesianPGMM* for R to facilitate the adoption of this methodology in practice. We describe an application to mouse-tracking data for a visual recognition task.

Keywords

Bayesian; longitudinal data; Markov chain Monte Carlo; mixture model; non-linear random effects models; piecewise function

1. Introduction

Longitudinal statistical models with piecewise functions (i.e., spliced lines or *splines*) are extremely useful for analyzing data that demonstrate a segmented growth process, where the

overall functional form incorporates two or more distinct functional forms with transitions from one phase to another over time (Kohli et al., 2015b). For instance, the development and/or decline in cognitive ability over time follows a piecewise trajectory (e.g., Bradway and Thompson (1962); Dominicus et al. (2008); Hall et al. (2001); Muniz Terrera et al. (2011)). The appeal of piecewise growth models is due to their flexibility to accommodate different kinds of functional forms with respect to each developmental phase (e.g., Cudeck and Klebe (2002)). To elaborate, a linear-linear piecewise function can be used where the respective trajectory in the first and second developmental phase is linear, or a quadratic-linear or exponential-linear piecewise function can be used when the trajectory in the first developmental phase has some curvature and the rate of change in the second developmental phase is constant.

An important parameter of a piecewise function is the *change point (knot)*, the time of transition from one developmental phase (segment) to another. The changepoint can be specified a priori or estimated as part of the model fitting procedure. Inferring the location of the changepoint(s) is often of practical interest, along with inference for other model parameters. It is possible to estimate a single changepoint or multiple changepoints under the assumption that transition times are common for a population of individuals (i.e., fixed changepoint), or that they vary from individual to individual (i.e., random changepoint) allowing for individual differences where the segments of the growth trajectories join (e.g., Fearnhead and Liu (2011); Lai and Albert (2014); Morrell et al. (1995); Wang and McArdle (2008)).

In many research settings, the data may consist of individuals from two or more unobserved subpopulations (i.e., classes). For example, in studies related to the Head Start programs, it has been found that children with two or more years of program participation have slower achievement growth, on average, than children with only one year of program participation (Kreisman, 2003). Analysis of data with latent classes requires the extension of growth models to include a categorical latent grouping variable. These statistical models are known as growth mixture models. The statistical framework of piecewise growth models can be extended to enable the identification of unknown classes, where individuals within a given class share similar mean segmented developmental trajectory with unknown changepoint(s) (e.g., Kohli et al. (2013, 2015a,b)). A piecewise growth mixture model (PGMM) can incorporate widely different changepoint locations and functional forms within each class; a PGMM with one class is simply referred to as a piecewise growth model (PGM). For example, Zhao and Banerjee (2012) considered a PGMM with a single, unknown random changepoint to examine log prostate-specific antigen (PSA) monthly data after prostate cancer surgery, with race and cancer stage at diagnosis as covariates.

The framework of PGMM is flexible and has a lot of utility in substantive research. However, fitting these models is complicated and time intensive. This is so because PGMMs are intrinsically nonlinear growth mixture models and the estimation of the mean or location of the changepoint(s) (a nonlinear parameter), along with the variance of the changepoint(s) is computationally very challenging. Additionally, the parameter estimation for these models may either employ frequentist inference procedures (e.g., Hall et al. (2000); Kohli et al. (2015a); Naumova et al. (2001)) or Bayesian inference procedures (e.g., Carlin et al. (1992);

Dominicus et al. (2008)). Wang and McArdle (2008), and Wang and Fang (2009) detailed the comparison of both inference approaches within the context of a single class piecewise growth model. Kohli et al. (2015b) compared the Bayesian-MCMC approach with the ML-EM approach for estimating PGMM with an unknown changepoint that is class specific but does *not* vary for individuals within a class. The results from their study showed Bayesian estimation generally performed better than frequentist ML-EM estimation, in terms of both the accuracy of the estimated model parameters and computational feasibility.

In all of these previous studies the researchers hypothesized and *prefixed* the number of unknown changepoint locations, i.e., the number of changepoints were specified in advance. There is no existing methodological study that empirically *detects* the number of changepoints, i.e., considers the number of changepoints as unknown and to be inferred from the data, within a unified framework for inference. This is limiting for many applications. Piecewise studies of educational data typically assume one changepoint (Sullivan et al., 2016; Kohli et al., 2015c; Kieffer, 2012); however, it is plausible that many learning trajectories will have at least two changepoints: one preceding a period of accelerated growth (an “a-ha” moment), and another preceding a period of decelerated growth (a “saturation point”) (Gallistel et al., 2004). Multiple changepoints are also plausible for many physical growth processes. It is generally agreed that human height growth occurs in at least three linear phases (Karlberg, 1987), but that these phases can be altered in malnourished subpopulations (Karlberg et al., 1994). For these and other applications a flexible inferential framework that allows for an arbitrary number of latent changepoints, as well as individual variation and population heterogeneity in the form of latent classes, is needed.

This leads to the primary objective of this article, that is, (1) to develop a PGMM using a Bayesian inference approach that allows the estimation of *multiple* random changepoints within each class; (2) to develop a procedure to empirically *detect* the number of random changepoints within each class; and (3) to empirically investigate the bias and precision of the estimation of the model parameters, including the random changepoints, via a simulation study.

The organization of this article is as follows. In Section 2, we introduce a general probabilistic model for a PGMM with random number of changepoints. In Section 3 we describe a generally applicable Bayesian approach to estimate the parameters of the model in Section 2. In Section 4 we describe an application to mouse-tracking data for a visual language processing experiment, using the Bayesian method described in Section 3. In Section 5 we describe a comprehensive simulation study to assess the performance of the method, using results from the application in Section 4 as a starting point to simulate data. In Section 6 we describe an R package to apply the method. In Section 7 we discuss the advantages and limitations of our proposed approach, with potential extensions.

2. Likelihood model

Here we describe the data generating (likelihood) model for a piecewise growth mixture model with linear segments and latent number of changepoints. Hierarchical prior

distributions for unknown model parameters and other details that are specific to a Bayesian framework are discussed later in Section 3. In what follows, Greek characters denote unknown parameters to be estimated.

We first describe the model for a single class and fixed number of changepoints. For a single subject i , assume data are available for M_i measurement occasions. For each occasion $j = 1, \dots, M_i$, let y_{ij} denote the measured outcome for subject i and let x_{ij} denote the time of measurement. The longitudinal trajectory of the outcome is characterized by $K + 1$ segments of linear growth or decline, where K is the number of changepoints. That is, for changepoint locations $\lambda_{i1}, \dots, \lambda_{iK}$,

$$y_{ij} = \beta_{i,0} + \beta_{i,1}x_{ij} + \sum_{k=1}^K \beta_{i,k+1}(x_{ij} - \lambda_{i,k})^+ + \varepsilon_{ij}, \quad (1)$$

where $(\cdot)^+$ is the positive part function,

$$(x_{ij} - \lambda_{i,k})^+ = \begin{cases} x_{ij} - \lambda_{i,k} & \text{if } x_{ij} - \lambda_{i,k} > 0 \\ 0 & \text{otherwise,} \end{cases}$$

and the error terms ε_{ij} are independent and normally distributed with variance σ_ε^2 :

$$\varepsilon_{ij} \stackrel{iid}{\sim} N(0, \sigma_\varepsilon^2). \quad (2)$$

Because $\beta_{i,k+1}$ denotes the change in slope at the k 'th changepoint, the slope of the k 'th linear segment is given by $\beta_{i1} + \beta_{i2} + \dots + \beta_{ik}$.

If subject i belongs to a single-class cohort of size N , $i = 1, \dots, N$, it is natural to assume a Gaussian random effects model for the subject-specific intercepts, slope changes, and changepoints:

$$\begin{aligned} \beta_{i,k} &\stackrel{iid}{\sim} N(\bar{\beta}_k, \sigma_{\beta_k}^2) \text{ for } k = 0, \dots, K+1, \text{ and} \\ \lambda_{i,k} &\stackrel{iid}{\sim} N(\bar{\lambda}_k, \sigma_{\lambda_k}^2) \text{ for } k = 1, \dots, K, \end{aligned}$$

where $\{\bar{\beta}_k, \bar{\lambda}_k\}$ give the mean effects over all subjects and $\{\sigma_{\beta_k}^2, \sigma_{\lambda_k}^2\}$ give the variances of the subject-level effects. For identifiability of the index k we presume $\bar{\lambda}_1 < \bar{\lambda}_2 < \dots < \bar{\lambda}_K$.

If the number of changepoints is unknown, we extend model (1) by introducing the latent parameter $\mathcal{K}$:

$$y_{ij} = \beta_{i,0} + \beta_{i,1}x_{ij} + \sum_{k=1}^K \beta_{i,k+1}(x_{ij} - \lambda_{i,k})^+ \mathbb{1}_{\{k \leq \mathcal{K}\}} + \varepsilon_{ij}, \quad (3)$$

Where $\mathbb{1}_{\{\cdot\}}$ is the indicator

$$\mathbb{1}_{\{k \leq \mathcal{K}\}} = \begin{cases} 1 & \text{if } k \leq \mathcal{K} \\ 0 & \text{otherwise.} \end{cases}$$

Note that $\mathcal{K} \in \{0, 1, \dots, K\}$ defines the number of changepoints, with a maximum of K possible changepoints.

If the N subjects belong to C classes, we let $\psi(i) \in \{1, \dots, C\}$ represent the class to which subject i belongs for $i = 1, \dots, N$. In our context the class memberships $\psi(i)$ are unknown, and $\psi(i) = c$ with probability ν_c for $i = 1, \dots, N$ and $c = 1, \dots, C$. Here ν_c defines the marginal probabilities in each class (i.e., the mixing proportions). The parameters $\{\beta_{\psi,k}, \lambda_{\psi,k}\}$ arise from a C -component Gaussian mixture with class specific means and variances:

$$\begin{aligned} \beta_{i,k} &\stackrel{iid}{\sim} N(\bar{\beta}_{\psi(i),k}, \sigma_{\psi(i),k}^2) \text{ for } k = 0, \dots, K+1, \text{ and} \\ \lambda_{i,k} &\stackrel{iid}{\sim} N(\bar{\lambda}_{\psi(i),k}, \sigma_{\psi(i),k}^2) \text{ for } k = 1, \dots, K. \end{aligned} \quad (4)$$

The number of changepoints may also depend on the classes:

$$y_{ij} = \beta_{i,0} + \beta_{i,1}x_{ij} + \sum_{k=1}^K \beta_{i,k+1}(x_{ij} - \lambda_{i,k})^+ \mathbb{1}_{\{k \leq \mathcal{K}_{\psi(i)}\}} + \varepsilon_{ij}, \quad (5)$$

where $\mathcal{K}_1, \dots, \mathcal{K}_C$ give the latent number of changepoints in each class. The class membership may either be known or latent.

3. Parameter inference

In this section we discuss inference for the unknown parameters of the model defined by Equations (2), (4) and (5). We use a Bayesian framework for inference. Our reasons for adopting a Bayesian framework, rather than, for instance, a frequentist framework with maximum likelihood estimation, are several fold. Bayesian inference gives a flexible philosophical framework for assessing uncertainty in all parameters, incorporating subjective prior beliefs, and borrowing information across multiple studies. Furthermore, Bayesian inference has generally been shown to outperform maximum likelihood inference for less complex piecewise random effects models, in terms of both the accuracy of the estimated model parameters and computational feasibility (Wang and McArdle, 2008; Kohli et al., 2015b). Moreover, maximum likelihood is generally prone to overfitting richly parametrized

models (Myung, 2000). For example, a direct application of maximum likelihood to estimate the unknown parameters of the model given in Section 2 will always select the maximum number of changepoints ($\mathcal{K}_1 = \dots = \mathcal{K}_C = K$), because this case subsumes models where $\mathcal{K}_c < K$ for some class c . Alternatively, one could specify and fix $\{\mathcal{K}_c\}_{c=1}^C$ and estimate the remaining parameters via maximum likelihood. Under this approach post-hoc penalized likelihood heuristics such as the Akaike information criterion (AIC) (Akaike, 1974) or Bayesian information criterion (BIC) (Schwarz et al., 1978) may be used to compare models with different candidate values for $\{\mathcal{K}_c\}_{c=1}^C$. While AIC/BIC values can be used for model selection, they provide little intuition on the relative strength of evidence and uncertainty in the number of changepoints. Moreover, AIC and BIC are justified asymptotically and their direct application is often inappropriate for models with multilevel or clustered effects (see, e.g., Vaida and Blanchard (2005) and Delattre et al. (2014)), and estimating all permutations of the number of possible changepoints for multiple classes can be computationally intensive. Rather, we treat $\{\mathcal{K}_c\}_{c=1}^C$ as random variables to be inferred in a unified Bayesian framework.

Let D denote the set of all observed data

$$D = \{(y_{ij}, x_{ij}) : i = 1, \dots, N \text{ and } j = 1, \dots, M_i\}$$

and let Θ denote the set of all unknown model parameters

$$\Theta = \left\{ \sigma_\epsilon^2, \{\bar{\beta}_{c,k}\}, \{\bar{\lambda}_{c,k}\}, \{\sigma_{c,\beta_k}^2\}, \{\sigma_{c,\lambda_k}^2\}, \{\mathcal{K}_c\}, \nu \right\}. \quad (6)$$

The maximum number of possible changepoints K and the maximum number of classes C are fixed hyper-parameters. Let $L(D|\Theta)$ be the multivariate Gaussian likelihood defined by Equations (2), (4) and (5). Under a Bayesian framework, one puts a *prior* probability distribution on the parameter space, $p(\Theta)$. An application of Bayes' rule gives the probability distribution of Θ conditional on the observed data:

$$p(\Theta|D) = \frac{p(\Theta)L(D|\Theta)}{\int p(\Theta)L(D|\Theta)d\Theta}, \quad (7)$$

this is called the *posterior* probability distribution, and serves as the basis for parametric inference.

3.1. Prior selection

The choice of a prior distribution $p(\Theta)$ reflects the *a priori* probability distribution of the parameters, which is refined by the observed data to give the posterior. In practice the choice of a prior can be informed by previously observed data or expert opinion. Alternatively, the

prior can be non-informative, chosen to minimize subjective bias and maximize the influence of the observed data D on the posterior. We describe in detail a prior that is intended to be non-informative for a given maximum of changepoints (K) and number of possible classes (C). This is the prior we use for our implementation, and it is designed to be applicable in a wide variety of situations by default when there is no clear rationale for a more informative prior. Alternative priors may be used, depending on the application, especially when there is rationale to use a more subjective prior. For example, here we construct a uniform discrete prior for the number of changepoints and a uniform prior over the inner range of X for the location of the changepoints; if subject matter knowledge for a given application suggests more information on the number and location of the changepoints, this information can be expressed in an alternative prior.

We use the notation

$$X = \{x_{ij} : i = 1, \dots, N \text{ and } j = 1, \dots, M_i\}$$

and

$$Y = \{y_{ij} : i = 1, \dots, N \text{ and } j = 1, \dots, M_i\}.$$

Note that the measured time-points x_{ij} need not be common across subjects, and the number of measurement occasions M_i may depend on the subject i . For example, if the measurement times x_{ij} are generally common across subjects but some outcome measurements y_{ij} are missing, time points with missing measurements for subject i can be removed for subject i only.

The prior distribution for each model parameter is given with appropriate justification below. For general applicability these priors should be robust to the scale of measurement for X and Y . Thus for some parameters the prior involves *empirical hyper-parameters*, i.e., parameters for a prior that depend on the data. However, these hyper-parameters are given only in terms of the baseline mean and overall scale of the observed measurements, to ensure that they are invariant to the scale of measurement for a given application. This is equivalent to using non-empirical priors after linearly scaling the data before analysis, in a way that will not affect the shape and structure of the trajectories (just scale the x- and y- axes).

- $\sigma_e^2 \sim \text{Inverse-gamma}(0.001, 0.001)$. The inverse-gamma distribution is a *conjugate* prior for the variance of a normal distribution, meaning both the prior and the posterior belong to the same distributional family (here, the inverse-gamma family). The Inverse-gamma(0.001, 0.001) prior is a commonly used non-informative prior for the residual variance (see Spiegelhalter et al. (1996)), as it accommodates both very small ($\sigma_e^2 \rightarrow 0$) and very large ($\sigma_e^2 \rightarrow \infty$) values.
- $\bar{\beta}_{c,0} \stackrel{iid}{\sim} N(\bar{y}_{\cdot,1}, \text{var}(y_{\cdot,1}))$ for $c = \{1, \dots, C\}$, where $\bar{y}_{\cdot,1}$ and $\text{var}(y_{\cdot,1})$ are the sample mean and variance of outcomes at the initial time point. The normal

distribution is conjugate prior for the mean parameters of normally distributed random effects. Here $\bar{y}_{.,1}$ and $\text{var}(y_{.,1})$ are empirical hyper-parameters assuming that the initial time point corresponds to 0 ($x_{p1} = 0$) for all subjects. If not, the time scale can be shifted so that the first measured time point corresponds to $x = 0$.

- $\bar{\beta}_{c,k} \stackrel{iid}{\sim} N(0, (\text{sd}(Y)/\text{sd}(X))^2)$ for $c = \{1, \dots, C\}$ and $k = \{1, \dots, K+1\}$, where $\text{sd}(X)$ is the sample standard deviation of the time points and $\text{sd}(Y)$ is the sample standard deviation of the outcome over all time points. The empirical hyper-parameter $\text{sd}(Y)/\text{sd}(X)$ is motivated by the observation that for the simple model

$$y_{ij} = \beta_{0i} + \beta_i x_{ij} + \varepsilon_{ij} \text{ with } \beta_i \sim N(\bar{\beta}, \sigma_\beta^2), \quad (8)$$

$|\bar{\beta}| < \text{sd}(Y)/\text{sd}(X)$; thus, we generally expect slope parameters to have absolute value smaller than $\text{sd}(Y)/\text{sd}(X)$.

- $\bar{\lambda}_{c,k} \stackrel{iid}{\sim} \text{Uniform}(a, b)$ for $c = \{1, \dots, C\}$ and $k = \{1, \dots, K\}$, where a is the second-to-earliest measured time point in X (over all subjects), and b is the second-to-latest time point. We choose not to let the mean changepoints have uniform support over the full range of X ($\text{Uniform}(\min(X), \max(X))$) to avoid model identifiability issues at the boundaries.
- $\sigma_{c,\beta_0} \stackrel{iid}{\sim} \text{Uniform}(0, \text{sd}(y_{.,1}))$ for $c = \{1, \dots, C\}$, where $\text{sd}(y_{.,1})$ is the sample standard deviation of outcomes at the initial time point. Here and elsewhere we use a uniform distribution over a range of plausible values for the standard deviation of latent subject-specific parameters (see Gelman et al. (2006)). The upper bound $\text{sd}(y_{.,1})$ is motivated by the observation that the variance of the subject-specific intercepts will always be less than the variance of the observed values at time 0.
- $\sigma_{c,\beta_k} \stackrel{iid}{\sim} \text{Uniform}(0, \text{sd}(Y)/\text{sd}(X))$ for $c = \{1, \dots, C\}$ and $k = \{1, \dots, K+1\}$. The upper bound $\text{sd}(Y)/\text{sd}(X)$ is motivated by the observation that in model (8), $\sigma_\beta^2 < \text{sd}(Y)/\text{sd}(X)$.
- $\sigma_{c,\lambda_k} \stackrel{iid}{\sim} \text{Uniform}(0, b)$ where

$$b = \frac{\max(X) - \min(X)}{4},$$

for $c = \{1, \dots, C\}$ and $k = \{1, 2, \dots, K\}$. The upper bound b is motivated by the observation that anything larger than b necessarily implies plausible values of the

random changepoint occur outside of the range of X . That is, $\sigma_{c,\lambda_k} = b$ implies that

$$P(\lambda_{i,k} \in \bar{\lambda}_{c,k} \pm \text{range}(X)/2) \approx 0.95,$$

and so if the changepoint mean is at the midpoint, $\bar{\lambda}_{c,k} = (\max(X) + \min(X))/2$, plausible values of the changepoint cover the entire range of X .

- $\mathcal{K}_c \stackrel{iid}{\sim} \text{Uniform}\{0, 1, \dots, K\}$ for $c \in \{1, \dots, C\}$. That is, we give equal prior probability to each possible number of changepoints in each class. In practice, this is accomplished by introducing auxiliary Bernoulli variables $I_{c,k}$ for $k = 1, \dots, K$, with

$$P(I_{c,k}) = \frac{K - k + 1}{K - k + 2},$$

then

$$\mathcal{K}_c = \sum_{k=1}^K \prod_{k'=1}^k I_{c,k'}.$$

As an alternative prior, we also consider $\mathcal{K}_c \stackrel{iid}{\sim} \text{Binomial}(K, p)$, where the binomial probability p controls the expected number of changepoints and can be used to specify a more conservative prior. The binomial prior is constructed with Bernoulli variables $I_{c,k}$ with $P(I_{c,k}) = p$, where

$$\mathcal{K}_c = \sum_{k=1}^K I_{c,k}.$$

- $\psi(i) = c$ with probability ν_c for $i = 1, \dots, N$ and $c = 1, \dots, C$. By default, we use a Dirichlet(1, ..., 1) prior for $(\nu_1, \dots, \nu_C)$, which is uniformly distributed over the standard $C - 1$ simplex ($\nu_1 + \nu_2 + \dots + \nu_C = 1$). When $C = 2$, the Dirichlet(α_1, α_2) distribution is equivalent to a Beta(α_1, α_2) distribution for ν_1 . Therefore the default prior corresponds to a Uniform(0, 1) prior on ν_1 , because Beta(1, 1) is Uniform(0, 1). More generally, a Dirichlet($\alpha_1, \dots, \alpha_C$) prior can be used. When $\alpha_1 = \dots = \alpha_C = \alpha^*$, smaller values of α^* suggest less parity in the class sizes; this can lead to under-identified classes in the posterior, because the prior allows some classes to have negligible probabilities. Larger values of α^* suggest more parity in the class sizes.

3.2. Posterior computation

The posterior distribution (7) cannot be derived analytically because the integral in the denominator is not tractable. Rather, we simulate draws from this posterior distribution via

Markov chain Monte Carlo (MCMC). In particular, we use the package *rjags* for R (Plummer, 2016), which performs Gibbs sampling in combination with other techniques to draw from the posterior distribution. Gibbs sampling was first described in Geman and Geman (1984) and is widely used to simulate draws from a Bayesian posterior with several parameters. We refer to Casella and George (1992) for an accessible and thorough introduction to Gibbs sampling. Sampling results in a Markov chain $\{\boldsymbol{\theta}^{(t)} = (\theta_1^{(t)}, \dots, \theta_d^{(t)})\}_{t=1}^T$ containing T dependent samples from the joint posterior distribution defined by $p(\boldsymbol{\Theta} \mid D)$.

The posterior sampling algorithm requires specifying initial parameter values $\boldsymbol{\theta}^{(0)} = \{\theta_1^{(0)}, \dots, \theta_d^{(0)}\}$. The initial samples $\boldsymbol{\Theta}^{(1)}, \boldsymbol{\Theta}^{(2)}, \dots$ are dependent on $\boldsymbol{\Theta}^{(0)}$, and so may not be representative of the posterior distribution. Typically, samples before a certain number of iterations T_0 are ignored, and sometimes called the *burn-in*. Those samples after iterations T_0 , $\{\boldsymbol{\theta}^{(T_0+1)}, \dots, \boldsymbol{\theta}^{(T)}\}$, are used for posterior inference. We obtain a point estimate for each parameter θ_i by taking the mean of posterior samples

$$\hat{\theta}_i = \frac{1}{T - T_0} \sum_{t=T_0+1}^T \theta_i^{(t)}.$$

We obtain a 95% *credible interval* (CI) $[a_i, b_i]$ where a_i is the 2.5th percentile and b_i is the 97.5th percentile of $\{\theta_i^{(T_0+1)}, \dots, \theta_i^{(T)}\}$; because the sample percentiles approximate quantiles of the posterior distribution, this gives

$$P(\theta_i \in [a_i, b_i] \mid D) \approx 0.95.$$

for discrete parameters, such as the number of changepoints $\mathcal{K}_c$ for each class, the relative frequency of each value over posterior samples approximates the posterior probability for that value.

In practice we initialize the model in two stages. First we perform Gibbs sampling for a small number of iterations for the simplified model with no within-class variability,

$$\sigma_{c, \beta_k}^2 = 0 \text{ and } \sigma_{c, \lambda_k}^2 = 0$$

for all c, k . In this first stage initial values are generated from the prior (Section 3.1), and Gibbs sampling is much faster than for the full model with random effects. The mean over iterations for parameters $\{\sigma_{\epsilon}^2, \{\bar{\beta}_{c,k}\}, \{\bar{\lambda}_{c,k}\}\}$, and the mode for parameters $\{\{\mathcal{K}_c\}, \boldsymbol{\psi}\}$, are then used as initial values for the full model.

For accuracy of posterior inference it is important that simulated draws after the burn-in have negligible dependence on the initial values, and that enough samples are taken to adequately approximate the full posterior distribution. If both of these conditions are satisfied, we say the algorithm has *converged*. We assess convergence by running multiple MCMC chains, from different initial values, in parallel. If the algorithm has converged we expect results obtained from each chain, from samples after the burn-in, to be similar. In particular we use the *potential scale reduction factor* (PSRF) (Gelman and Rubin, 1992), which is the square root of the ratio of the total variation across chains over the average variation within chains, as a robust measure of convergence. The multivariate PSRF (Brooks and Gelman, 1998) is a generalization of the PSRF for a single parameter, to assess convergence of a parameter vector. A PSRF or multivariate PSRF less than 1.1 or 1.2 are common thresholds (Sinharay, 2004).

When Gibbs sampling for a mixture model there is a risk of *label switching*. For example, what is labeled class 1 in the beginning of the chain may be labeled class 2 at the end of the chain, and vice versa. To address label switching we apply the Equivalence Classes Representatives (ECR) algorithm (Papastamoulis, 2014) available from the `label.switching` package for R (Papastamoulis, 2015), as a post-hoc step after posterior samples are collected. This algorithm separately permutes the class labels at each iteration to best match a specified *pivot* labeling. The pivot is initially given by the most common label for each sample across the iterations, and it is updated for repeated applications of the ECR algorithm until convergence. In addition to label-switching of the classes, we must consider identifiability of the changepoint labels. For $\mathcal{K}_c^{(t)} > 1$, we permute labels of the piecewise components so that $\bar{\lambda}_{c,i}^{(t)} < \bar{\lambda}_{c,j}^{(t)}$ for all $i < j < \mathcal{K}_c^{(t)}$, which is necessary for the components to be consistently labeled across the MCMC iterations.

Note that correlations between the random effects $\{\{\beta_{i,k}\}, \{\lambda_{i,k}\}\}$ are not explicitly modeled by the prior given in Section 3.1. This choice was motivated by simplicity and practicality. Conjugate priors for a correlation/covariance matrix that facilitate Gibbs sampling such as the inverse-Wishart prior are inflexible and biased, whereas more flexible priors inhibit computational speed and convergence (see, e.g., Daniels and Kass (1999)). Alternatively, we compute the empirical sample correlation between each pair of random effects, averaged over each sampling iteration, as a post-hoc estimate. For example, a point estimate for the correlation between the first changepoint and the initial slope for Class 1 is given by

$$\frac{1}{T - T_0} \sum_{t=T_0+1}^T \text{cor}\left\{\left(\beta_{i,1}^{(t)}, \lambda_{i,1}^{(t)}\right); \psi(i) = 1\right\}.$$

4. Application

We apply the estimation approach detailed in Section 3 to mouse-tracking data for a language recognition task, originally described in Incera and McLennan (2016). For this task, subjects were shown text of a given color, and instructed to move a computer mouse and click on the word corresponding to that color as quickly as possible. Thus the conditions

were analogous to a Stroop task, with certain experimental modifications. For this illustration we focus on the control task, in which participants were shown meaningless text of the given color (e.g., “XXXXXX” in red); the study also included a congruent task in which participants were shown the word of the same color (e.g., “red” in red) and an incongruent task in which participants were shown the word of a different color (e.g., “blue” in red). The x-coordinate of the mouse was tracked for each trial, on the range of -100 (away from the correct color) to $+100$ (toward the correct color). The x-location was measured at 20ms intervals for a period of 1000ms (1 second). The task was repeated for 16 trials for each of 60 subjects. We consider the average x-location at each time point over the trials, yielding a sample size of $N = 60$ with $M_j = 50$ time points for each subject.

The position of the mouse is considered a proxy for the cognitive process of determining the word corresponding to the correct color. Therefore, the data for each subject represents a longitudinal learning trajectory for a simple task on a very short time scale (1 second). We are interested in modeling the pattern of this trajectory, and interpreting this pattern as it relates to the cognitive processes of recognition and action. Intuitively, we expect these trajectories to have 2 changepoints, that vary from subject to subject. All trajectories begin at the origin, and are initially flat as the subjects process the image. Thus the first changepoint has a critical interpretation as the “decision point”, when the subject recognizes the correct word and begins to move their mouse in that direction. For the second changepoint we hope to capture the “conclusion point”, after the subject has reached the correct word and settles on their decision. We estimate the two change-point model for one, two, or three possible classes. The two-class model had the lowest deviance information criterion (DIC) (Spiegelhalter et al., 2002), with a DIC of 15863 (the DIC for the single class model was 15978, and for the three class model was 15898). Thus in what follows we focus on the two-class results, which also had the clearest interpretation. The authors of the original study further hypothesized a possible difference in the trajectory between monolingual speakers and bilingual speakers. Of the subjects, 20 were fluent only in English, and 40 were fluent in English and one other language. Thus, we also consider whether the identified classes are representative of monolingual and bilingual speakers.

We run Gibbs sampling for the full model with $T = 50,000$ total iterations, including $T_0 = 20,000$ iterations for burn-in, for 3 MCMC chains. The sampling algorithm converges satisfactorily, with a PSRF across chains of less than 1.10 for each of the parameters (6), and a multivariate PSRF of 1.08 over all parameters. Table 1 gives parameter estimates and 95% credible intervals for the overall and class-level parameters. Figure 1 shows a spaghetti plot showing the trajectory for each individual, and the class trajectories defined by the mean parameters for each class. The class membership of each trajectory shown in Figure 1 is determined by whichever class has the highest posterior probability for that individual. As shown in Table 1, there is considerable individual variability about the mean parameters for each class. Figure 2 shows the individual parameter fits for the model for three subject curves in each class; these curves are chosen to be representative (with posterior probability > 0.975 for their class) and to illustrate variation in the subject-level trajectories within each class.

The algorithm identifies one large class (Class 1; $\nu \approx 0.77$, or about 77% of subjects) and one smaller class (Class 2). There was strong evidence of two changepoints in each class, with posterior probability $P(\mathcal{K}_c = 2 | D)$ above 0.99 for $c = 1, 2$. The mean trajectory of the larger class agrees with the intuition above, wherein the mouse cursor is stationary at first, then increases sharply after the first changepoint when the subject makes a decision and moves in the correct direction, then is mostly stationary after the second changepoint after the subject settles on their decision. However, for the smaller class there is a more gradual increase after the first changepoint, and a sharper increase after the second changepoint. Our interpretation is that the subjects in Class 1 react quickly and confidently after making an initial decision, whereas subjects in Class 2 proceed more tentatively at first and gradually move toward their decision with increasing confidence. Interestingly, the decision point, given by the first changepoint, is similar for both classes; it is the swiftness with which they react after their initial decision, and the pattern of their reaction, that distinguishes the two classes.

The 95% credible intervals in Table 1 are relatively wide for the class proportion ν_2 , and for certain parameters in Class 2. This is partly due to uncertainty in the class memberships: 8 individuals had posterior class probabilities between 0.1 and 0.9. Uncertainty in the class memberships can drive uncertainty in the parameter estimates, particularly for the smaller class (Class 2). Moreover, even for given class memberships, the flat uniform prior will in general yield more posterior variance for the underlying class proportion than, e.g., a more informative Dirichlet prior with larger hyperparameters.

Of the 11 subjects that were allocated to the larger class, 5 were monolingual speakers. The association between the identified class for each subject and monolingual/bilingual status was not significant (p-value = 0.48, Fisher's exact test), suggesting that monolingual/bilingual status was not the dominant distinguishing characteristic in these data.

5. Simulation

Here we describe a comprehensive simulation study to evaluate different aspects of the proposed method, using the results from the application in Section 4 as a starting point. We simulate data under the piecewise linear model of Section 2 to assess recovery of the true parameters, using the estimates in Table 1 as parameter values, with certain manipulated conditions. We consider different values of the sample size N , and number of time points M_j , to assess how either affect the overall accuracy of the fitted model. We consider different class proportions ν_c , to assess how the relative sizes of each cluster affect class allocation. We also manipulate the number of changepoints $\mathcal{K}_c$, as detecting the number of changepoints is a key innovation of the present model. Specifically, we simulate under the following assumptions:

- The overall sample size, N , is 60 or 120.
- The number of time-points, M_j , is 25 or 50. For each simulation scenario the number of measurement occasions M_j is the same for each subject $i = 1, \dots, N$. If $M_j = 25$, the measurement occasions occur after intervals of length 40, $x_j = (40,$

80, 120, ..., 1000), rather than intervals of length 20 for $M_i = 50$, $x_i = (20, 40, 60, \dots, 1000)$.

- The proportion of samples in Class 1, ν_1 , is 0.8 or 0.5. If $\nu_1 = 0.8$ the classes are of size 48 and 12 for $N = 60$ or 96 and 24 for $N = 120$, if $\nu_1 = 0.5$ the classes are of size 30 and 30 for $N = 60$ or 60 and 60 for $N = 120$.
- The number of changepoints in Class 2, $\mathcal{K}_2$, is 0, 1, or 2. For $\mathcal{K}_2 = 1$ the effect at changepoint 2 is removed, and for $\mathcal{K}_2 = 0$ the effects at both changepoints are removed, but other parameters for Class 2 remain the same. We include both changepoints for Class 1 ($\mathcal{K}_1 = 2$) for all scenarios.

We implement a fully crossed factorial design with the above conditions, resulting in $2 \times 2 \times 2 \times 3 = 24$ simulation scenarios. We generate 100 replicated datasets under each simulation scenario, and apply the estimation algorithm outlined in Section 3 to each replication.

For each of the 2,400 total replications, we ran the algorithm for 50,000 sampling iterations per chain for each of the 2,400 total replications, with a 20,000 iteration burn-in; we ran three chains for each replication, to assess convergence. The percent of replications that converged, with a mean PSRF over all parameters less than 1.2, ranged from 26% to 82% across the 24 simulation scenarios. To assess the effect of the number of sampling iterations on convergence we repeated all 100 replications in the application-motivated scenario with $N = 60$, $M_i = 50$, $\nu_1 = 0.80$, and $\mathcal{K}_2 = 2$ for a larger number of iterations. For this scenario 44% of replications converged within the first 50,000 iterations (with 20,000 burn-in), 76% converged when the algorithm was run for 50,000 more iterations (100,000 total; 50,000 burn-in), and 91% converged when the algorithm was run for 100,000 more iterations (200,000 total; 100,000 burn-in). These results suggest that poor convergence can be overcome by increasing the number of sampling iterations for certain situations.

Among converged replications across all simulation conditions, the classes were generally well identified; when class membership was determined based on whichever class had higher posterior probability, the mean misclassification rate across all replications was 3.9%. The number of changepoints in each class was also generally well-recovered; the average estimated posterior probability for the true number of changepoints, across all classes and simulation conditions, was 0.958.

We used multi-factor ANOVA to investigate the effects of the manipulated simulation conditions on three key performance metrics representing convergence, misclassification, and recovery of the number of changepoints. For each metric we performed ANOVA with main effects for each of the four manipulated conditions. Table 2 gives the significance of each effect, as well as the overall mean within each level. Interestingly, while results generally improved with a larger sample size ($N = 120$ vs. $N = 60$), the number of measurement occasions ($M_i = 50$ vs. $M_i = 25$) did not significantly affect any metric.

For Class 1, which always had 2 changepoints, the correct number of changepoints had posterior probability greater than 0.5 for all simulation replications, and had posterior probability greater than 0.95 for 99.8% of replications. For Class 2, the mean probability for

each number of changepoints is shown against the true number of changepoints in Table 3. Here, too, the number of changepoints is generally well recovered. However, the scenario with 1 true changepoint gave non-trivial probability to the 2 changepoint model; this is likely in part due to misclassification, in which some of the subjects with 2 changepoints (from Class 1) are misallocated to Class 2.

Table 4 shows the bias and variability, as well as coverage rates for the 95% credible interval, for the converged replications across all simulation conditions. The parameters are generally well-recovered, but some are estimated with bias. This bias is apparently due in part to misclassification; parameter estimates tend to shrink toward the midpoint of both classes, as a small proportion of subjects in Class 1 are misallocated to Class 2 and vice-versa. Note also that the estimates for Class 2 are generally less precise, due to its smaller sample size in several simulation scenarios. For the $N = 601$ converged replications when $\nu_1 = 0.5$, the estimates for ν_1 had mean 0.492 and standard deviation 0.049, yielding a z-standardized bias of -0.16 . For the $N = 525$ converged replications when $\nu_1 = 0.8$, the estimates for ν_1 had mean 0.77 and standard deviation 0.079, yielding a z-standardized bias of 0.295.

A spreadsheet with summary results for each of the 24 simulation scenarios is available as a supplementary file online. Additional simulation studies and illustrations are described in a supplementary document available online. A simulation involving up to five changepoints ($K = 1, 2, 3, 4$, or 5) demonstrates the feasibility to accurately detect the number and location of a larger number of changepoints. A simulation with up to 4 classes ($C = 1, 2, 3$, 4) demonstrates the feasibility to correctly identify the class memberships even when the maximum number of classes is over-specified. Additional simulation studies illustrate the effect of prior specification on the posterior for key model parameters, which is important to consider for any Bayesian approach. In a study with different prior choices for the number of changepoints, the correct number of changepoints had the highest posterior probability for all priors considered, but certain prior choices favored over- or under-detection of the changepoints. Thus, the choice of prior can be used to imply a more conservative model that avoids over-estimating changepoints. A study on the effect of the cluster concentration parameter α illustrates how a higher α favors more equality among the latent class proportions in the posterior; however, the accuracy of the class memberships and other model parameters are relatively robust to changes in α . A study using alternative priors for the variances of the random effects reveals that an appropriately scaled half-Cauchy prior performs similarly to the uniform priors used here; generic (unscaled) half-Cauchy priors and uninformative inverse-gamma priors perform poorly and are not recommended.

6. Implementation

We have created a well-documented and user friendly package for R, called *BayesianPGMM*, to estimate the model under the default priors given in Section 3.1. This package provides estimates and 95% credible intervals for all model parameters, assesses convergence of the algorithm, and includes functions to automatically visualize and summarize results. The package is included as a supplementary file online [and will be uploaded to the CRAN

package repository upon acceptance of the manuscript]. A brief tutorial on the use of the package is given in Section 2 of the supplemental material available online.

7. Discussion

In this article we have described a mathematical framework for piecewise growth mixture models with unknown number of changepoints, and we have proposed a Bayesian estimation scheme that is appropriate for a wide variety of application scenarios. The advantages of our proposed method are its *flexibility* and its *simple interpretation*. The model can accommodate many diverse patterns of growth, with subject-level variation and minimal a priori assumptions on important parameters such as the number and location of the changepoints. This facilitates interpretation, as the locations at which the functional form of a growth trajectory changes, and the nature of those changes, is invaluable information for many applications in psychometrics and in other fields.

We have described an application to mouse-tracking data for a language recognition task, which may be extended in several ways. Here we have focused on modeling the mean mouse trajectories over 16 trials of the same task, across a population of individuals. Alternatively, one could directly model intra-subject variability over the 16 trials, which would involve a more complex model with multiple hierarchical levels for subjects and trials within a subject. One could also allow potential covariates of interest, such as mono/bilingual status, to enter the model explicitly in a number of ways.

By necessity our framework involves many unknown parameters that are influential and highly interdependent, and this can pose a challenge to their estimation. In particular, the bias in estimation of class-specific parameters can be non-negligible, especially if there is considerable overlap between the two classes. And the computational time required to fit the model is not trivial for large datasets. The application described in Section 4 took approximately 1 hour to run on a laptop with a 2.5 GHz i7 processor, and computing time scales approximately linearly with the sample size, the number of measurement occasions, and the number of MCMC iterations. In practice we recommend monitoring convergence closely, and only interpreting parameter estimates after the MCMC algorithm has reached the desired level of convergence. Fortunately, in certain situations poor convergence may be addressed by increasing the number of MCMC iterations for a given application, if necessary.

In the application and simulation we have focused on the model with a maximum of 2 changepoints ($K = 2$), and 2 classes ($C = 2$). The framework readily allows for more changepoints and classes, and a supplemental document describing additional simulations and illustrations when $K > 2$ or $C > 2$ is available online. However, we caution against over-parameterization. At minimum, K must be restricted by the number of time-points M_i that are commonly observed across subjects. In particular, if the number of active changepoints is greater than $M_i - 1$, this requires inferring multiple changepoints between two measurement occasions for subject i , which presents identifiability issues. DIC or other model selection metrics may be used to select these hyper-parameters. But the asymptotically motivated assumptions for DIC, such as multivariate normality of the

posterior, are likely not well satisfied for smaller sample sets. Diagnostics such as the widely applicable information criterion (WAIC) or leave-one-out cross-validation (LOO) (Vehtari et al., 2017) may be useful alternatives. WAIC and LOO consider the accuracy of the posterior predictive distribution for held-out data points, and in this context, a “new data point” may be the entire trajectory for an individual, or a single measurement occasion.

Our framework may be extended in several ways. In this article we focus on a model with no discontinuity (jumps) in the outcome over time. We think continuity is a reasonable assumption for most applied situations, but the model may be extended to accommodate discontinuity at the changepoints if necessary. In this article we have also only considered piecewise segments that have a linear form. In theory, these segments may take an exponential, quadratic, or other polynomial form. Allowing these forms to be unknown and vary across different segments would improve the flexibility of the model, but would increase the complexity of estimation and interpretation.

More advanced and targeted estimation techniques may improve the efficiency and accuracy of posterior computation. For example, reversible jump MCMC (Green, 1995) inherently infers the presence of fundamentally different parameters, and may be used in this context to jump between models with different number of changepoints or different functional forms for each phase during MCMC. Moreover, there is a considerable literature on Bayesian estimation for changepoint problems involving a single process with multiple phases realized from distinct stationary distributions (see, e.g., Carlin et al. (1992); Chib (1998); Fearnhead (2006)). That is, $y_j \sim p(\eta_j)$ where parameters η_j are given by phases over time:

$$\eta_j = \begin{cases} \theta_1 & \text{if } x_j < \lambda_1 \\ \theta_2 & \text{if } \lambda_1 \leq x_j < \lambda_2 \\ \vdots & \\ \theta_{K+1} & \text{if } \lambda_{K-1} \leq x_j < \lambda_K. \end{cases}$$

In particular, Chib (1998) suggests a hidden Markov approach, that models the probability of transitioning from one phase to the next for discrete time points. Such an approach is not readily extended to our context, chiefly because the assumption of stationarity within each phase is not satisfied, and modeling segment-wise constant probabilities of change at discrete time points is not readily incorporated into a random effects framework where the location of the changepoint(s) are of substantial interest. However, generally this literature provides many tools that may improve Bayesian estimation of segmented growth models for a population. Improving the estimation and flexibility of segmented growth models remains an exciting area of active and future research.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We would like to thank Sara Incera and Conor T. McLennan of Cleveland State University for graciously providing the mouse-tracking data described in Section 4.

References

- Akaike H. A new look at the statistical model identification. *IEEE transactions on automatic control*. 1974; 19(6):716–723.
- Bradway KP, Thompson CW. Intelligence at adulthood: A twenty-five year follow-up. *Journal of Educational Psychology*. 1962; 53(1):1.
- Brooks SP, Gelman A. General methods for monitoring convergence of iterative simulations. *Journal of computational and graphical statistics*. 1998; 7(4):434–455.
- Carlin BP, Gelfand AE, Smith AF. Hierarchical Bayesian analysis of changepoint problems. *Applied statistics*. 1992:389–405.
- Casella G, George EI. Explaining the Gibbs sampler. *The American Statistician*. 1992; 46(3):167–174.
- Chib S. Estimation and comparison of multiple change-point models. *Journal of econometrics*. 1998; 86(2):221–241.
- Cudeck R, Klebe KJ. Multiphase mixed-effects models for repeated measures data. *Psychological methods*. 2002; 7(1):41. [PubMed: 11928890]
- Daniels MJ, Kass RE. Nonconjugate Bayesian estimation of covariance matrices and its use in hierarchical models. *Journal of the American Statistical Association*. 1999; 94(448):1254–1263.
- Delattre M, Lavielle M, Poursat MA, et al. A note on BIC in mixed-effects models. *Electronic journal of statistics*. 2014; 8(1):456–475.
- Dominicus A, Ripatti S, Pedersen NL, Palmgren J. A random change point model for assessing variability in repeated measures of cognitive function. *Statistics in medicine*. 2008; 27(27):5786–5798. [PubMed: 18680123]
- Fearnhead P. Exact and efficient Bayesian inference for multiple changepoint problems. *Statistics and computing*. 2006; 16(2):203–213.
- Fearnhead P, Liu Z. Efficient Bayesian analysis of multiple changepoint models with dependence across segments. *Statistics and Computing*. 2011; 21(2):217–229.
- Gallistel CR, Fairhurst S, Balsam P. The learning curve: implications of a quantitative analysis. *Proceedings of the national academy of Sciences of the United States of America*. 2004; 101(36):13124–13131. [PubMed: 15331782]
- Gelman A, et al. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian analysis*. 2006; 1(3):515–534.
- Gelman A, Rubin DB. Inference from iterative simulation using multiple sequences. *Statistical science*. 1992:457–472.
- Geman S, Geman D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence*. 1984; (6):721–741. [PubMed: 22499653]
- Green PJ. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*. 1995; 82(4):711–732.
- Hall CB, Lipton RB, Sliwinski M, Stewart WF. A change point model for estimating the onset of cognitive decline in preclinical Alzheimer's disease. *Statistics in medicine*. 2000; 19(11–12):1555–1566. [PubMed: 10844718]
- Hall CB, Ying J, Kuo L, Sliwinski M, Buschke H, Katz M, Lipton RB. Estimation of bivariate measurements having different change points, with application to cognitive ageing. *Statistics in medicine*. 2001; 20(24):3695–3714. [PubMed: 11782027]
- Incera S, McLennan CT. Mouse tracking reveals that bilinguals behave like experts. *Bilingualism: Language and Cognition*. 2016; 19(03):610–620.
- Karlberg J. On the modelling of human growth. *Statistics in medicine*. 1987; 6(2):185–192. [PubMed: 3589247]
- Karlberg J, Jalil F, Lam B, Low L, Yeung C. Linear growth retardation in relation to the three phases of growth. *European journal of clinical nutrition*. 1994; 48:S25–43. [PubMed: 8005089]
- Kieffer MJ. Before and after third grade: Longitudinal evidence for the shifting role of socioeconomic status in reading growth. *Reading and Writing*. 2012; 25(7):1725–1746.

- Kohli N, Harring J, Zopluoglu C. A finite mixture of nonlinear random coefficient models for continuous repeated measures data. *Psychometrika*. 2015a; 81(3):851–880. [PubMed: 25925010]
- Kohli N, Harring JR, Hancock GR. Piecewise linear–linear latent growth mixture models with unknown knots. *Educational and Psychological Measurement*. 2013; 73(6):935–955.
- Kohli N, Hughes J, Wang C, Zopluoglu C, Davison M. Fitting a linear–linear piecewise growth mixture model with unknown knots: A comparison of two common approaches to inference. *Psychological methods*. 2015b; 20(2):259–275. [PubMed: 25867487]
- Kohli N, Sullivan AL, Sadeh S, Zopluoglu C. Longitudinal mathematics development of students with learning disabilities and students without disabilities: A comparison of linear, quadratic, and piecewise linear mixed effects models. *Journal of school psychology*. 2015c; 53(2):105–120. [PubMed: 25746821]
- Kreisman MB. Evaluating academic outcomes of head start: An application of general growth mixture modeling. *Early Childhood Research Quarterly*. 2003; 18(2):238–254.
- Lai Y, Albert PS. Identifying multiple change points in a linear mixed effects model. *Statistics in medicine*. 2014; 33(6):1015–1028. [PubMed: 24114935]
- Morrell CH, Pearson JD, Carter HB, Brant LJ. Estimating unknown transition times using a piecewise nonlinear mixed-effects model in men with prostate cancer. *Journal of the American Statistical Association*. 1995; 90(429):45–53.
- Muniz Terrera G, Van den Hout A, Matthews F. Random change point models: investigating cognitive decline in the presence of missing data. *Journal of Applied Statistics*. 2011; 38(4):705–716.
- Myung IJ. The importance of complexity in model selection. *Journal of Mathematical Psychology*. 2000; 44(1):190–204. [PubMed: 10733864]
- Naumova EN, Must A, Laird NM. Tutorial in biostatistics: evaluating the impact of critical periods in longitudinal studies of growth using piecewise mixed effects models. *International Journal of Epidemiology*. 2001; 30(6):1332–1341. [PubMed: 11821342]
- Papastamoulis P. Handling the label switching problem in latent class models via the ECR algorithm. *Communications in Statistics-Simulation and Computation*. 2014; 43(4):913–927.
- Papastamoulis P. label. switching: An R package for dealing with the label switching problem in MCMC outputs. arXiv preprint arXiv:1503.02271. 2015
- Plummer M. rjags: Bayesian Graphical Models using MCMC. R package version. 2016:4–6.
- Schwarz G, et al. Estimating the dimension of a model. *The annals of statistics*. 1978; 6(2):461–464.
- Sinharay S. Experiences with Markov chain Monte Carlo convergence assessment in two psychometric examples. *Journal of Educational and Behavioral Statistics*. 2004; 29(4):461–488.
- Spiegelhalter D, Thomas A, Best N, Gilks W. BUGS 0.5: Bayesian inference using Gibbs sampling manual. MRC Biostatistics Unit, Institute of Public Health; Cambridge, UK: 1996. 1–59.
- Spiegelhalter DJ, Best NG, Carlin BP, Van Der Linde A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2002; 64(4):583–639.
- Sullivan AL, Kohli N, Farnsworth EM, Sadeh S, Jones L. Longitudinal models of reading achievement of students with learning disabilities and without disabilities. *School Psychology Quarterly*. 2016 in press.
- Vaida F, Blanchard S. Conditional Akaike information for mixed-effects models. *Biometrika*. 2005; 92(2):351–370.
- Vehtari A, Gelman A, Gabry J. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*. 2017; 27(5):1413–1432.
- Wang L, McArdle JJ. A simulation study comparison of Bayesian estimation with conventional methods for estimating unknown change points. *Structural Equation Modeling*. 2008; 15(1):52–74.
- Wang Y, Fang Y. Least square and empirical Bayes approaches for estimating random change points. *Journal of Data Science*. 2009; 7(1):111–127.
- Zhao L, Banerjee M. Bayesian piecewise mixture model for racial disparity in prostate cancer progression. *Computational statistics & data analysis*. 2012; 56(2):362–369. [PubMed: 25540470]

Figure 1.

Spaghetti plot of the mouse-tracking data for the color recognition task (a positive position corresponds to the correct color). The trajectory defined by the mean parameters for each class are shown in **bold**.

Figure 2.
Individual model fits for three representative subjects in Class 1 and three subjects in Class 2. Raw data are shown in gray, piecewise model fits are shown in black.

Table 1

Estimates, with 95% credible interval, for model parameters in each class.

Parameter	Class 1 estimate	Class 1 95% CI	Class 2 estimate	Class 2 95% CI
ν	0.77	(0.53, 0.88)	0.23	(0.13, 0.47)
σ	3.16	(3.08, 3.24)	3.16	(3.08, 3.24)
β_0	0.003	(−0.054, 0.054)	0.000	(−0.052, 0.053)
β_1	−0.002	(−0.006, 0.001)	−0.005	(−0.012, 0.002)
β_2	0.194	(0.172, 0.232)	0.060	(0.036, 0.108)
β_3	−0.171	(−0.202, −0.143)	0.081	(−0.052, 0.135)
σ_{β_0}	0.020	(0.001, 0.038)	0.019	(0.001, 0.038)
σ_{β_1}	0.010	(0.008, 0.013)	0.008	(0.004, 0.014)
σ_{β_2}	0.064	(0.049, 0.082)	0.027	(0.013, 0.054)
σ_{β_3}	0.079	(0.056, 0.103)	0.068	(0.035, 0.122)
λ_1	362	(332, 391)	321	(222, 409)
σ_{λ_1}	93.6	(70.6, 119)	132	(72.2, 221)
λ_2	643	(567, 695)	726	(626, 827)
σ_{λ_2}	149	(102, 191)	128	(68.3, 222)

Table 2

ANOVA p-values and mean effects for four different simulations conditions (sample size N , measurement occasions M_i , class proportion ν_1 , and number of changepoints in Class 2 $\mathcal{K}_2$) on three different performance metrics.

	$\text{Log}_2(\text{PRS})$	Misclassification	P(Correct $\mathcal{K}_2$)
$N = 60$	0.290	0.043	0.954
$N = 120$	0.366	0.035	0.962
p-value	< 0.001	0.001	0.256
$M_i = 25$	0.334	0.041	0.957
$M_i = 50$	0.321	0.038	0.958
p-value	0.880	0.307	0.896
$\nu_1 = 0.8$	0.320	0.056	0.917
$\nu_1 = 0.5$	0.336	0.025	0.993
p-value	0.001	< 0.001	< 0.001
$\mathcal{K}_2 = 2$	0.464	0.071	0.985
$\mathcal{K}_2 = 1$	0.334	0.061	0.902
$\mathcal{K}_2 = 0$	0.185	0.005	0.979
p-value	< 0.001	< 0.001	< 0.001

Table 3

True number of changepoints and mean posterior probability, for Class 2.

	True $\mathcal{K}_2 = 2$	True $\mathcal{K}_2 = 1$	True $\mathcal{K}_2 = 0$
Posterior $\mathcal{K}_2 = 2$	0.97	0.17	0.00
Posterior $\mathcal{K}_2 = 1$	0.02	0.81	0.04
Posterior $\mathcal{K}_2 = 0$	0.01	0.02	0.96

Table 4

Summary of parameter estimates in Class 1 and 2 over all converged replications (1126 total replications), including true values (in **bold**), mean of the estimates, standard deviation of the estimates across simulation replications, standardized bias ($Z\text{-bias}=(\text{Mean}-\text{Truth})/(\text{SD})$) and coverage rate of the 95% credible interval.

	Class 1	Mean	SD	Z-bias	Coverage	Class 2	Mean	SD	Z-bias	Coverage
σ	3.16	3.17	0.054	0.28	0.929	3.16	3.17	0.054	0.28	0.929
β_1	-0.002	-0.002	0.002	-0.05	0.947	-0.005	-0.004	0.009	0.19	0.903
β_2	0.194	0.191	0.014	-0.25	0.885	0.060	0.060	0.022	0.04	0.887
β_3	-0.171	-0.164	0.018	0.38	0.867	0.081	0.059	0.048	-0.48	0.838
σ_{β_1}	0.010	0.010	0.003	-0.04	0.926	0.008	0.012	0.017	0.22	0.863
σ_{β_2}	0.064	0.055	0.012	-0.76	0.634	0.027	0.036	0.025	0.35	0.833
σ_{β_3}	0.079	0.077	0.014	-0.13	0.877	0.068	0.077	0.027	0.33	0.818
λ_1	362	363	16.6	0.08	0.956	321	328	60.7	0.11	0.918
$\lambda_1\sigma$	93.6	97.6	12.8	0.31	0.931	132	142	27.9	0.36	0.929
λ_2	643	643	33.4	0.002	0.926	726	743	69.8	0.26	0.904
σ_{λ_2}	149	149	19.4	-0.03	0.970	128	144	30.0	0.53	0.970