

Article

# Bayesian Nonlinear Models for Repeated Measurement Data: An Overview, Implementation, and Applications

Se Yoon Lee<sup>1\*</sup>

<sup>1</sup> Department of Statistics, Texas A&M University, College Station, Texas, USA; stat\_math@tamu.edu  
\* Correspondence

**Abstract:** Bayesian nonlinear mixed effects models for data in the form of continuous, repeated measurements from a population, also known as Bayesian hierarchical nonlinear models, are a popular platform for analysis when interest focuses on individual specific characteristics and relevant uncertainty quantification. Due to the limitation of computational power, this framework was relatively dormant until the late 1980s, but in recent years, the statistical research community saw vigorous development of new methodological and computational techniques for these models, the emergence of software, and wide application of the models in numerous industrial and academic fields. This article presents an overview of the formulation, interpretation, and implementation of Bayesian nonlinear mixed effects models and surveys recent advances and applications.

**Keywords:** Bayesian nonlinear hierarchical model; Bayesian nonlinear mixed effects models; Inter-individual variation; Intra-individual variation; Markov chain Monte Carlo technique

## 1. Introduction

One of the common challenges in biological, agricultural, environmental, epidemiological, financial, and medical applications is to make inferences on characteristics underlying profiles of continuous, repeated measures data from multiple individuals within a population of interest [1–4]. By ‘repeated measures data’ we mean the data type generated by observing a number of individuals repeatedly under differing experimental conditions where the individuals are assumed to constitute a random sample from a population of interest. A common type of repeated measures data is longitudinal data such that the observations are ordered by time [5,6].

Linear mixed effects models for repeated measures data have become popular due to their straightforward interpretations, flexibility allowing correlation structure among the observations, and utility accommodating unbalanced and multi-level data structure (i.e., clustered designs that vary among individuals) [7,8]. The modeling framework is also intuitively appealing: the central idea that individuals’ responses are governed by a linear model with slope or intercept parameters that vary among individuals seems to be appropriate in many scientific problems (for e.g., see [9,10]). It also allows practitioners to test and evaluate multivariate causal relationships by conducting regression analysis at the population level. By preserving the multi-level structure in a single model, estimation or prediction for the analyses can take advantage of information borrowing [11].

For many applications, researchers often want to theorize that time courses of individual response commonly follow a certain nonlinear function dictated by a finite number of parameters [12]. These nonlinear functions are based on reasonable scientific hypotheses, typically represented as a differential equation system. By tuning the parameters, the shape of the function in terms of curvature, steepness, scale, height, etc, may change, which is used as the rationale behind describing heterogeneity between subjects. Nonlinear mixed effects models, also referred to as hierarchical nonlinear models, have gained broad acceptance as a suitable framework for these purposes [13–15]. Analyses based on this model are now routinely reported in various industrial problems, which is, in part, enabled by the

breakthrough development of software [16–20]. The excellent books and review papers were published by [14,15,21]. Although their works were published more than 20 years ago, they still provide statisticians, programmers, and researchers with many pedagogical insights about the modeling framework, implementations, and practical applications of using the nonlinear mixed effects models.

While frequentist analysis of nonlinear mixed effects models has a long history, Bayesian analysis for the models was a relatively dormant field until the late 1980s. This is due primarily to the time-consuming nature of the calculations required for Bayesian computation to implement a Bayesian model [22]. Since the early 1990s Bayesian approaches began to re-emerge, motivated both by exploitation of rapid developments in computing power and by the growing desire to quantify the uncertainty associated with parameter estimation and prediction [23–25]. Since then, Bayesian nonlinear mixed effects models, also called Bayesian hierarchical nonlinear models, have been extensively used in diverse industrial and academic researches, endowed with new computational tools providing a far more flexible framework for statistical inference matching exactly the increasing complexity of scientific research [26–31].

The objective of this article is to present an updated look at the Bayesian nonlinear mixed effects models. Although the works of [14,15] discuss some of the Bayesian approaches for the nonlinear mixed effects models, the main perspective adopted in the works is much more oriented to the frequentist framework, and prior distributions and Bayesian computing strategy explained in the works are quite outdated. In the literature, it is striking that very few research works provide an updated overview of the Bayesian methodologies on the nonlinear mixed effects models. Motivated this, in this article, we endeavor to survey a recent perspective of formulation, interpretation, and implementation of Bayesian nonlinear mixed effects models, with recently developed prior distributions and Bayesian computation techniques. Illustrated methodologies can be directly exploited in a wide variety of applications, ranging from small data to big data problems in both industrial and academic research.


This article is organized as follows. In Section 2, we motivate readers to understand why it is necessary to use the Bayesian nonlinear mixed effects model by illustrating four real problems, which will be conceptualized as a statistical problem. To solve the statistical problem, we suggest a basic version of the Bayesian nonlinear mixed effects models in Section 3, and its likelihood is analyzed in Section 4 wherein frequentist computations are briefly discussed. Section 5 describes modern Bayesian computation strategies to implement the basic model. Popularly used prior distributions are presented in Section 6. Section 7 discusses model selection, and Section 8 reviews recent advances and extensions that build on the basic model. Finally, Section 9 concludes the article.

2. Bayesian Nonlinear mixed effects model

2.1. Pharmaceutical industry

Studies of the pharmacokinetics of drugs help us learn about the variability in drug disposition in a population [32]. Figure 1 shows theophylline concentration in the plasma as a function of time after oral administration of the same amount of anti-asthmatic theophylline for 12 subjects. (The data considered here are courtesy of Dr. Robert A. Upton of the University of California, San Francisco.) As seen in the panel, concentration trajectories have a similar functional shape for all individuals. However,  $C_{max}$  and  $t_{max}$  (peak concentration and time when it is achieved), absorption, and elimination phases are substantially different across subjects. Clinical pharmacologists believe that these differences are attributable to between-subject variation in the underlying pharmacokinetic processes, explained by Absorption, Distribution, Metabolism, and Excretion (ADME), understanding of which is crucial in a new drug development in the pharmaceutical industry.

In pharmacokinetics analysis, often abbreviated by ‘PK analysis’, it is routine to use compartmental modeling to describe the amount of drug in the body by dividing the whole body into one or more compartments [33]. For theophylline, a one-compartment model is



**Figure 1.** Theophylline concentrations for 12 subjects following an oral dose

normally used, which assumes that the entire body acts like a single, uniform compartment; see page 30 from [34] for a detailed explanation about the model:

$$C(t) = \frac{DFk_a}{V(k_a - Cl/V)} \left\{ \exp\left(-\frac{Cl}{V}t\right) - \exp(-k_at) \right\}, \quad (1)$$

where  $C(t)$  is drug concentration at time  $t$  for a single subject following oral dose  $D$  at  $t = 0$ . Here,  $F$  is the bioavailability which expresses the proportion of a drug that gains access to the systemic circulation.  $k_a$  is the absorption rate constant describing how quickly drug is absorbed from the gut into the systemic circulation.  $V$  is the volume of the central compartment.  $Cl$  is the clearance rate representing the volume of plasma from which drug is eliminated per unit time. Eventually, the pharmacokinetic processes for a given subject is summarized by the 4-dimensional vector with 'PK parameters' ( $F, k_a, V, Cl$ ). Obviously, it is the modeler's discretion to proceed with a more complex PK model such as a three compartment models with nonlinear clearance to fit the data, but in this case, over-parameterization should be carefully examined [35].

Typically, the dataset collected in a drug development program includes demographic and clinical covariates obtained from each subject, for e.g., body weight, height, age, sex, creatinine clearance, albumin, etc; and furthermore, one can also involve genetic information in an individual's response to drugs. Most covariates are measured at baseline, before assigning the drug, while some covariates can be measured at every sampling time. One of the crucial goals of PK analysis is to illustrate the effect of such covariates on the PK parameters [36]. The causal relationship inferred by the covariate analysis can be used to support physicians in making the necessary judgments about the medicines that they prescribe, tailored to individual patients [37].


In PK report for a new drug application to government authorities like U.S. Food and Drug Administration (FDA) or European Medicines Agency (EMA), PK parameters are summarized by mean or median, and very importantly, estimates of parameter precision. Estimates of parameter precision can provide valuable information regarding the adequacy of the data to support those parameters [38]. Parameter uncertainty can be estimated through several methods, including bootstrap procedures [39], log-likelihood profiling [40], or using the asymptotic standard errors of parameter estimates, and recently, Bayesian approaches draw a lot of attention from the pharmaceutical industry [41]. Particularly,

Bayesian approaches for the population PK analysis can be very useful when there is prior knowledge about PK parameters learned from preclinical studies, published works, etc, and one wants to incorporate them into the prior specification for PK parameters [28].

## 2.2. Oil and gas industry

The US shale boom—a product of technological advances in horizontal drilling and hydraulic fracturing that unlocked new stores of energy—has greatly benefited the growth in the US economy. Horizontal drilling is a directional drilling technology such that a well is drilled parallel to the reservoir bedding plane [42]. Well productivity of a horizontal well is known to be often 3 to 5 times greater than that of a vertical well [43,44], but also costs 1.5 to 2.5 times more than a vertical well [45]. Therefore, the eventual success of the drilling project of unconventional shale wells relies on a large degree of well construction costs [46]. Because of very low permeability, and a flow mechanism very different from that of conventional reservoirs, estimates for the shale well construction cost often contain high levels of uncertainty. For this reason, one of the crucial tasks of petroleum engineers is to quantify the uncertainty associated with the process of oil or gas production to reduce the extra initial risk for the projects.

Figure 2 shows monthly production rate trajectories of 360 shale oil wells completed in the Eagle Ford Shale of South Texas, studied by [31]. The declining pattern manifested in the trajectories is commonly observed in almost all oil production rate time series data following well completion. (Here, the completion is terminology in petroleum engineering, meaning the process of transforming a well ready for the initial production [47].) Decline curve analysis (DCA), introduced by [48] around 100 years ago, is one of the most popularly utilized methods for petroleum engineers. Its purpose is to (i) theorize a curve describing the declining pattern, (ii) analyze the declining production rates, (iii) characterize the well-productivity, and (iv) forecast the future performance of oil and gas wells. Particularly, estimation and uncertainty quantification of estimated ultimate recovery (EUR) (here, EUR is a special jargon defined as an approximated quantity of oil from a well which is potentially recoverable by the end of its producing life [49]) is the utmost important task and a starting point in the decision-making process for future drilling projects. Also, the oil and gas companies comply with financial regulations about EUR outlined by the U.S. Securities and Exchange Commission: see [www.sec.gov/](http://www.sec.gov/) - for the regulations.



**Figure 2.** Production rates for 360 shale oil wells after completion.


Most curves used in DCA are derived from solving certain differential equations that describe a hidden dynamic from production rate trajectory [50–55]. See [56–59] for an overview of such curves. [31] studied Arps' hyperbolic, stretched exponentiated decline, Duong, and Weibull curves to fit the trajectories shown in the Figure 2. Particularly, the Duong model was developed for unconventional reservoirs with very low permeability:

$$P(t) = q_1 t^{-m} \exp \left\{ \frac{a(t^{1-m} - 1)}{1 - m} \right\},$$

where  $P(t)$  is the production rate at time  $t$  for a single well following completion.  $q_1$  is the initial rate coefficient, and  $m$  and  $a$  are additional model parameters. We note that the parameters,  $q_1$ ,  $m$  and  $a$ , have their own meanings in terms of well-productivity: see [55] for the interpretation. That being said, the well-productivity for a given well is summarized by the 3-dimensional parameter vector,  $(q_1, m, a)$ . In modeling perspective, the variation of the well-productivity across different wells is attributable to the different values for  $(q_1, m, a)$ . To explain this variability, one can regress the values  $(q_1, m, a)$  on the well-design parameters such as true vertical depth, measure depth, etc. The causal relationship inferred by the covariate analysis will be used in a future drilling project. Geological information of wells can be also incorporated to make a spatial prediction for the EUR at a new location, as researched by [31].

### 2.3. Financial industry

Macroeconomists, financial economists, and market participants all attempt to build good models of the 'yield curve' [60]. The yield curve on a given day is a curve showing the interest rates across different maturity spans (1 month, one year, five years, etc.) for a similar debt contract at a particular date. It determines the interest rate pattern (i.e., cost of borrowing), which can be used to calculate a bond's price [61]. Figure 3 shows daily treasury par yield curve rates spanning from January 3rd to January 13th, 2022, with maturities up to 30 years. The data source is from the U.S. Department of the treasury (<https://www.treasury.gov/>). As seen from the panel, the shape of the yield curve displays a slightly delayed humped shape. Economists believe that such a shape of the yield curve has an important implication on the economic growth [62].



**Figure 3.** Daily Treasury par yield curve rates from January 3rd to January 13th, 2022.



The Nelson-Siegel model [63] is a very popular model in the literature to fit the term structure:

$$Y(\tau) = \beta_0 + \beta_1 \left\{ \frac{1 - \exp(-\lambda\tau)}{\lambda\tau} \right\} + \beta_2 \left\{ \frac{1 - \exp(-\lambda\tau)}{\lambda\tau} - \exp(-\lambda\tau) \right\},$$

where  $Y(\tau)$  denotes the (zero-coupon) yield at evaluated at  $\tau$ , and  $\tau$  denotes the time to maturity. The model parameters have a specific financial meaning:  $\beta_0$ ,  $\beta_1$ , and  $\beta_2$  are related long-term, short-term, and midterm effects on the interest rate, respectively, and  $\lambda$  is referred to a decay factor [64]. Each of the yield curves is summarized by the 4-dimensional parameter  $(\beta_0, \beta_1, \beta_2, \lambda)$ , and it is known that model can capture a wide range of possible shapes of the yield curve [63–66]. Therefore, the Nelson-Siegel model is extensively used by central banks and monetary policymakers [67]. For example, The Federal Reserve updates estimates of  $(\beta_0, \beta_1, \beta_2, \lambda)$  once per week: visit the website (<https://www.federalreserve.gov/>). In recent years, there has been a great deal of interest in the uncertainty quantification of the Nelson-Siegel parameters over time, and their relationship with macroeconomic variables such as inflation and real activity, etc, in financial applications: refer to [66,68–70] for some of those works.

2.4. Early stage of epidemic


Novel coronavirus disease 2019 (COVID-19) is a big threat to global health. The rapid spread of the virus has created a pandemic, and countries all over the world are struggling with a surge in COVID-19 infected cases. Figure 4 displays the daily infection trajectories describing the cumulative numbers of infected cases for 40 countries, spanning from January 22nd to May 14th, 2020, studied by [30]. The data source is from COVID-19 Data Repository by the Center for Systems Science and Engineering at Johns Hopkins University (<https://coronavirus.jhu.edu/>). Refer to Table S.1 in [30] for the list of 40 countries. The time frame of the authors’ research was the early stage of the pandemic when there was no drug or other therapeutics approved by the US FDA.

In general, during an early phase of a pandemic, information regarding the disease is very limited and scattered even if it exists. In spite of that, it is crucial to predict future cases of infection or death. In such a situation, one consideration is to use data integration (also called ‘borrowing information’), combining data from diverse sources and eliciting useful information with a unified view of them. Additionally, it is very important to find risk factors relevant to the disease. Reliable and early risk assessment of a developing infectious disease outbreak allow policymakers to make swift and well-informed decisions that would be needed to ensure epidemic control. Quantifying uncertainty about the final epidemic size is also very important.

Richards growth curve [71], so-called the generalized logistic curve [72], is a popularly used growth curve for population studies in situations where growth is not symmetrical about the point of inflection [73,74]. There are variant reparameterized forms of the Richards curve in the literature [75–78], and one of the frequently used form is

$$I(t) = \frac{a}{[1 + \xi \exp\{-b(t - c)\}]^{1/\xi}},$$

where  $I(t)$  is the cumulative number of infected cases at time  $t$ . Here, epidemiological meanings of the parameters,  $a$ ,  $b$ , and  $c$ , are the final epidemic size, infection rate, and lag phase of the trajectory, respectively. The parameter  $\xi$  is the shape parameter, and there seems no clear epidemiological meaning [79]. Each infection trajectory in Figure 4 can be characterized by the 4-dimensional parameters  $(a, b, c, \xi)$  if the Richards curve is used. Due to its flexibility originating from the shape parameter  $\xi$ , Richards curve has been widely used in epidemiology for real-time prediction of outbreak of diseases, possibly at an early phase of the pandemic when there is no second wave. Examples include SARS [80,81], dengue fever [82,83], pandemic influenza H1N1 [84], and COVID-19 outbreak [30,85].



**Figure 4.** Daily trajectories for cumulative numbers of COVID-19 infections for 40 countries from January 22nd to May 14th, 2020.

2.5. Statistical problem

In the previous subsections, we presented a range of examples in which nonlinear mixed effects models can be exploited. They have their own challenges to solve the problems that are representative of issues many researchers have to deal with in other areas: for example, (1) how to describe a possible nonlinear clearance with a limited number of patients; (2) how to handle an enormously large number of shale oil wells and make a spatial prediction of EUR at a new location; (3) how to describe the dynamic of the financial parameters over time; and (4) how to integrate data from different sources to produce more accurate forecast on the epidemic size.

An emerging issue accompanied by these problems, requested from researchers, government entities, domain experts, etc, is how to quantify the uncertainty associated with parameter estimation and prediction. Although the traditional nonlinear mixed effects models, based on the maximum likelihood method, can provide confidence intervals and statistical tests, calculations of those generally involve approximations that are most accurate for large sample sizes. On the other hand, in the Bayesian approach - in which the prior automatically imposes the parameter constraints - inferences about parameter values based on the posterior distribution usually require integration rather than maximization, and no further approximation is involved. For that reason, the Bayesian approach is often suggested as a viable alternative to the frequentist approach to solving the problems.

We now formulate these problems as a statistical problem. First, we summarize a common feature of the dataset for the analysis.

- (1) There exists repeated measures of a continuous response over time for each subject;
- (2) There exists a variation of individual observations over time;
- (3) There exists a variation from subject-to-subject in trajectories;
- (4) There exist covariates measured at baseline for each subject.

The subject of sampling units considered in the statistical analysis is quite comprehensive. We have seen that it can be a patient, a shale oil well, a particular date, and a country. As the unique identifier, we assign the index  $i$  to each individual. By denoting  $N$  as the number of individuals (i.e., the sample size), the index  $i$  will take an integer from 1 to  $N$ . The sample size  $N$  available for the data analysis substantially varies across different industrial problems as well as subfields within the same industry. For example,

the number of shale oil wells on Eagle Ford Shale Play can be as large as 6,000 [31]. As for the pharmaceutical industry, in phase I cancer clinical trials, the number of cancer patients  $N$  may be strictly confined to 25 [86], but for phase III trials for non-oncology drug studies,  $N$  can be as large as 2,000 [87].

Here the term time is meant in the broadest sense. It can be a calendar time, a nominal time, a time after some event (for e.g., the time after dose from Figure 1 and the time after well completion from Figure 2), or a time to some event (for e.g., the time to maturity from Figure 3). Essentially, time can be defined as a physical quantity that can be indexed with consecutive integers to produce a temporal record. Another important characteristic of the time is that each subject may have different time points where observations are measured. In this article, we use  $t_{ij}$  to represent the time point, where the integer  $j = 1, 2, \dots, M_i$  indexes the time point from the earliest to the last observations. Thus,  $M_i$  represents the number of repeated observations for the  $i$ -th individual. When  $M_i$  is relatively small (or large), we say the repeated measures are sparsely observed (or densely observed). For example, the theophylline and yield curve data shown in Figure 1 and 3 are sparse data, while the oil production and COVID-19 data shown in Figure 2 and 4 are dense data.

As for the repeated measures,  $y_{ij}$  denotes the continuous response of the  $i$ -th subject at the time point  $t_{ij}$ . We assume that  $y_{ij}$  has been already pre-processed so that it is ready to be used for statistical modeling. For most applications, it may be necessary first to transform the data into some new representation before training the model. For example, as seen from Figure 2, oil productions vary substantially across different wells. For that reason, the authors [31] take a logarithm on the productions to derive the response  $y_{ij}$ , followed by appropriate statistical modeling on the log-scale. To some extent, data pre-processing may enhance the performance of the model.


Suppose that researchers collected  $P$  number of covariates at the baseline from each subject  $i$  ( $i = 1, \dots, N$ ). Here, the baseline refers to the time point  $t_{i1}$  (or possibly right before the time point  $t_{i1}$ ), where at the first response  $y_{i1}$  has not been observed yet. Let  $x_{ib}$  denote the  $b$ -th covariate of the  $i$ -th subject ( $b = 1, \dots, P$ ). In general, there are two types of covariates: time-invariant and time-varying covariates. This article mainly concerns the former type. As similar to  $N$ , the number of covariates  $P$  substantially varies across industries and specific problems. For instance, in pharmacogenetics analysis, the number of protein-coding genes  $P$  would around 20,000 [88]. In the oil and gas industry, if we consider most of the covariates obtained from the well completion procedure,  $P$  could be at least 100 [31].

In conclusion, the dataset for the statistical analysis can be represented by the collection of the  $N$  triplets  $\{(\mathbf{y}_i, \mathbf{t}_i, \mathbf{x}_i)\}_{i=1}^N$ . Here, for each subject  $i$  ( $i = 1, \dots, N$ ), we formulated two  $M_i$ -dimensional vectors  $\mathbf{y}_i = (y_{i1}, \dots, y_{ij}, \dots, y_{iM_i})^\top$  and  $\mathbf{t}_i = (t_{i1}, \dots, t_{ij}, \dots, t_{iM_i})^\top$ , and a  $P$ -dimensional vector  $\mathbf{x}_i = (x_{i1}, \dots, x_{ib}, \dots, x_{iP})^\top$ .

The data structures described up to this point are commonly encountered in longitudinal data studies [7]. Essentially, the feature of dataset motivating the use of nonlinear mixed effects models is that, for each subject  $i$ , the response vector  $\mathbf{y}_i$  displays some nonlinear tendency over time  $\mathbf{t}_i$ , as seen in Figure 1, 2, 3, and 4. To explain this nonlinearity, a researcher needs to theorize some nonlinear function, denoted as  $f$ , such as one compartment, Duong, Nelson-Siegel, and Richards models, depending on the contexts. The construction of such functions relies on human modelers' abstraction of data into a suitable dynamical system, which is often represented by a differential equation. Such a differential equation has a finite number of parameters that control the dynamic of the solution of the system, understanding of which is vital for causal inference for the nature of the system by associating with covariates  $\mathbf{x}_i$ .

Figure 5 displays a pictorial description about how a PK modeler would see the theophylline concentration trajectory from the modeling perspective, where she theorized that the one compartment model (1) would be suitable to describe the trajectories  $\mathbf{y}_i$  over time  $\mathbf{t}_i$  for each subject  $i$  ( $i = 1, \dots, N$ ). Then the 10-dimensional vector  $\mathbf{y}_i$  is summarized by a 4-dimensional PK parameter vector  $(F_i, k_{ai}, V_i, Cl_i)$ ; the dimension reduction is intrinsic





**Figure 5.** Pictorial illustration of PK modeling for the theophylline data.

sically embedded in this process. As each of the parameters  $F_i$ ,  $k_{ai}$ ,  $V_i$ , and  $Cl_i$  has an important clinical meaning, it is very natural to ask how they are related with  $P$  covariates  $\mathbf{x}_i$  to induce a causal relationship. For the purpose of modeling, it may be necessary to transform the original PK parameters  $(F_i, k_{ai}, V_i, Cl_i) \in [0, 1] \times (0, \infty)^3$  to model parameters  $(\theta_{1i}, \theta_{2i}, \theta_{3i}, \theta_{4i}) \in \mathbb{R}^K$  ( $K = 4$ ) so that elements  $\theta_{li}$  ( $l = 1, \dots, K$ ) are supported on the real number, by taking transformations  $\theta_{1i} = \log\{F_i/(1 - F_i)\}$ ,  $\theta_{2i} = \log k_{ai}$ ,  $\theta_{3i} = \log Cl_i$ , and  $\theta_{4i} = \log Cl_i$ . As this transformation was taken only for the modeling purpose, interpretations on the PK parameter for the PK report should be carried out after transforming back to the original scale.

### 3. The model

#### 3.1. Basic model

Assume that we have dataset for a statistical analysis  $\{(\mathbf{y}_i, \mathbf{t}_i, \mathbf{x}_i)\}_{i=1}^N$  from  $N$  subjects, as explained in Subsection 2.5. The usual Bayesian nonlinear hierarchical model may then be written as a three-stage hierarchical model as follows:

##### • Stage 1: Individual-Level Model

$$y_{ij} = f(t_{ij}; \boldsymbol{\theta}^i) + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma^2), \quad (i = 1, \dots, N; j = 1, \dots, M_i). \quad (2)$$

In (2), the conditional mean  $\mathbb{E}[y_{ij} | \boldsymbol{\theta}^i, \sigma^2] = f(t_{ij}; \boldsymbol{\theta}^i)$  is a known function governing within-individual temporal behavior dictated by a  $K$ -dimensional parameter  $\boldsymbol{\theta}^i = (\theta_{1i}, \theta_{2i}, \dots, \theta_{li}, \dots, \theta_{Ki})^\top \in \mathbb{R}^K$  specific to the subject  $i$ .  $\epsilon_{ij}$  is an additive Gaussian noise with variance  $\sigma^2$  across individuals and times.

##### • Stage 2: Population Model

$$\theta_{li} = \alpha_l + \mathbf{x}_i^\top \boldsymbol{\beta}_l + \eta_{li}, \quad \eta_{li} \sim \mathcal{N}(0, \omega_l^2), \quad (i = 1, \dots, N; l = 1, \dots, K). \quad (3)$$

In (3), the  $l$ -th model parameter  $\theta_{li}$  is used as the response of an ordinary linear regression with predictor  $\mathbf{x}_i$ , with intercept  $\alpha_l \in \mathbb{R}$  and coefficient vector  $\boldsymbol{\beta}_l = (\beta_{l1}, \beta_{l2}, \dots, \beta_{lP}) \in \mathbb{R}^P$ . By letting  $\boldsymbol{\eta}_i = (\eta_{1i}, \eta_{2i}, \dots, \eta_{li}, \dots, \eta_{Ki}) \in \mathbb{R}^K$ , we assume that the  $\boldsymbol{\eta}_i$  is distributed according a  $K$ -dimensional Gaussian distribution  $\mathcal{N}(\mathbf{0}, \boldsymbol{\Omega})$


with covariance matrix  $\Omega = \text{diag}(\omega_1^2, \omega_2^2, \dots, \omega_1^2, \dots, \omega_K^2) \in \mathbb{R}^{K \times K}$ . The diagonality in  $\Omega$  implies that each model parameter are uncorrelated across  $l$ .

• **Stage 3: Prior**

$$\sigma^2 \sim \pi(\sigma^2), \quad \alpha_l \sim \pi(\alpha_l), \quad \beta_l \sim \pi(\beta_l), \quad \omega_l^2 \sim \pi(\omega_l^2), \quad (l = 1, \dots, K). \quad (4)$$

Distributions in (4) are chosen to encapsulate any information or belief which have been formulated about the parameters. We suggest some popularly used prior options in Section 6.

Directed asymmetric graphical (DAG) model representation of the basic model (2) – (4) is depicted in Figure 6. Following the grammar of the graphical model (Chapter 8 of [89]), the circled variables indicate stochastic variables, while observed ones are additionally colored in grey. Non-stochastic quantities are uncircled. The arrows indicate the conditional dependency between the variables.



**Figure 6.** The basic model (2) – (4) as a graphical model.

### 3.2. Vectorized form of the basic model

We will often wish to write the hierarchy (2) – (4) for the  $i$ -th individual's entire response vector and represent it with an equivalent vector-form. This turns out to be useful to develop relevant computational techniques. We first introduce a  $K \times N$  dimensional matrix frequently used throughout this article:

$$\Theta = \begin{bmatrix} \theta_{11} & \cdots & \theta_{1i} & \cdots & \theta_{1N} \\ \vdots & & \vdots & & \vdots \\ \theta_{l1} & \cdots & \theta_{li} & \cdots & \theta_{lN} \\ \vdots & & \vdots & & \vdots \\ \theta_{K1} & \cdots & \theta_{Ki} & \cdots & \theta_{KN} \end{bmatrix} \in \mathbb{R}^{K \times N}. \quad (5)$$

The matrix (5) is referred to as *model matrix* because it comprises of model parameters  $\{\theta_{li}\}_{l=1, i=1}^{K, N}$  aggregating all subjects. Indeed, most of computational techniques either via frequentist or Bayesian setting in the literature have been developed to overcome an obstacle of a nonlinear association of the model matrix  $\Theta$  into the mean function  $f$ .

In (5), the subject index  $i$  is stacked column-wisely, while model parameter index  $l$  is stacked row-wisely, different from the usual way adopted in most statistics. The column indexing for the subjects (i.e., stacking individual-based vector horizontally) shown in (5) is often adopted in modern computation theory of deep learning [90], and one of the main advantages of using this indexing is that it may give some pedagogical insights on the use

of vectorization toward the entries  $\{\theta_{li}\}_{l=1,i=1}^{K,N}$  to exploit parallel computations, stochastic updating, etc, in optimization or sampling techniques.

We can simplify the model matrix  $\Theta$  (5) by denoting  $\Theta = [\theta^1 \dots \theta^i \dots \theta^N] \in \mathbb{R}^{K \times N}$  by stacking the individual model parameter vector in Stage 1 (2). In another aspect, we can represent the matrix with  $\Theta = [\theta_1 \dots \theta_l \dots \theta_K]^\top \in \mathbb{R}^{K \times N}$  by defining a  $N$ -dimensional vector corresponding the  $l$ -th model parameter across all subjects  $\theta_l = (\theta_{l1}, \theta_{l2}, \dots, \theta_{lN})^\top \in \mathbb{R}^N$  ( $l = 1, \dots, K$ ). Former and latter indexing method are referred to as  $i$ -indexing and  $l$ -indexing, respectively.

We are now in a position to re-write the hierarchy (2) – (4) using the vector notations:

• **Stage 1: Individual-Level Model**

$$\mathbf{y}_i = f_i(\mathbf{t}_i, \theta^i) + \epsilon_i, \quad \epsilon_i \sim \mathcal{N}_{M_i}(\mathbf{0}, \sigma^2 \mathbf{I}), \quad (i = 1, \dots, N). \quad (6)$$

In (6),  $f_i(\mathbf{t}_i, \theta^i)$  is a  $M_i$ -dimensional vector whose elements are temporally stacked:  $f_i(\mathbf{t}_i, \theta^i) = (f(t_{i1}; \theta^i), f(t_{i2}; \theta^i), \dots, f(t_{iM_i}; \theta^i))^\top$  for the subject  $i$ . The vector  $\epsilon_i$  is distributed according to the  $M_i$ -dimensional Gaussian distribution with mean  $\mathbf{0}$  and covariance matrix  $\sigma^2 \mathbf{I}$ .

• **Stage 2: Population Model ( $l$ -indexing)**

$$\theta_l = \mathbf{1}\alpha_l + \mathbf{X}\beta_l + \eta_l, \quad \eta_l \sim \mathcal{N}_N(0, \omega_l^2 \mathbf{I}), \quad (l = 1, \dots, K). \quad (7)$$

In (7), for each  $l$ , the  $N$ -dimensional model parameter vector  $\theta_l$  is used as the response vector of an ordinary linear regression: (i)  $N$ -by- $P$  design matrix  $\mathbf{X} = [\mathbf{x}_1 \mathbf{x}_2 \dots \mathbf{x}_N]^\top$ ; (ii) intercept  $\alpha_l$ ; (iii) coefficient vector  $\beta_l$ , and (iv) isotropic Gaussian error vector  $\eta_l = (\eta_{l1}, \eta_{l2}, \dots, \eta_{lN})^\top$  with variance  $\omega_l^2$ . (Notation  $\mathbf{1}$  in (7) represents an all-ones vector.).

• **Stage 2': Population Model ( $i$ -indexing)**

$$\theta^i = \alpha + \mathbf{B}\mathbf{x}_i + \eta^i, \quad \eta^i \sim \mathcal{N}_K(\mathbf{0}, \Omega), \quad (i = 1, \dots, N). \quad (8)$$

Equation in (8) is derived by incorporating each of the  $N$  columns of the model matrix (5). Here,  $\alpha$  represents a  $K$ -dimensional vector  $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_K)^\top$ , and  $\mathbf{B}$  represents a  $K$ -by- $P$  matrix with rows  $\beta_l$  ( $l = 1, \dots, K$ ). Here, the  $K$ -dimensional vector  $\mathbf{B}\mathbf{x}_i$  in the right-hand side of (8) is the mathematically identical to  $\mathbf{X}_i\beta$ , where  $\mathbf{X}_i = \mathbf{I}_K \otimes \mathbf{x}_i^\top \in \mathbb{R}^{K \times KP}$  and  $\beta = (\beta_1, \beta_2, \dots, \beta_K) \in \mathbb{R}^{KP}$  ( $\mathbf{I}_K$  is the  $K$ -by- $K$  identity matrix and  $\otimes$  represents the Kronecker matrix product.). The error vector  $\eta^i = (\eta_{1i}, \eta_{2i}, \dots, \eta_{Ki})^\top$  is distributed according a  $K$ -dimensional Gaussian distribution with mean  $\mathbf{0}$  and covariance matrix  $\Omega = \text{diag}(\omega_1^2, \omega_2^2, \dots, \omega_K^2)$ .

• **Stage 3: Prior**

$$\sigma^2 \sim \pi(\sigma^2), \quad \alpha \sim \pi(\alpha), \quad \mathbf{B} \sim \pi(\mathbf{B}), \quad \Omega \sim \pi(\Omega). \quad (9)$$

Each of the parameter blocks in  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  is assumed to be independent a *priori*.


To summarize, we have two equivalent vector forms representing the basic model (2) – (4) according to how the model matrix  $\Theta$  (5) is vectorized:

• **Vector-form (a):** Stage 1-(6), Stage 2-(7) ( $l$ -indexing), and Stage 3-(9);

• **Vector-form (b):** Stage 1-(6), Stage 2'-(8) ( $i$ -indexing), and Stage 3-(9).

Figure 7 displays DAG representations of the two vector forms of the basic model. In vector-form (a),  $K$  latent nodes  $\{\theta_l\}_{l=1}^K$  are fully connected toward the  $N$  response vectors  $\{\mathbf{y}_i\}_{i=1}^N$ .

On the other hand, in vector-form (b),  $N$  latent nodes  $\{\theta^i\}_{i=1}^N$  are bijectively connected to the  $N$  response vectors  $\{y_i\}_{i=1}^N$  for each subject  $i$ . These two ways of looking at the framework of the Bayesian nonlinear hierarchical models complement a more proper understanding of the framework and provide modelers with a statistical insight. For example, vector-form (a) is useful to understand the mathematics for  $P$  regressors in the population regression, while vector-form (b) makes us easy to comprehend the role of the population covariance matrix  $\Omega$ .



**Figure 7.** DAG representations of the basic model (2) – (4) in vector-form (a) [Stage 1-(6), Stage 2-(7) ( $l$ -indexing), and Stage 3-(9)] (left) and vector-form (b) [Stage 1-(6), Stage 2'-(8) ( $i$ -indexing), and Stage 3-(9)] (right). Two vector forms are equivalent except for the way how the model matrix  $\Theta$  (5) is vectorized.

#### 4. Likelihood

##### 4.1. Likelihood based on Stage 1

As in most of the statistical models, a natural starting point for inference is maximum likelihood estimation. Likelihood theory is also fundamental of Bayesian inference, which is the main topic in this article. We start with considering only Stage 1 from the basic model in Section 3.1 and ignore Stage 2 and 3 for now, to get some pedagogical insights. Then the likelihood function for the  $i$ -th subject is

$$\mathcal{L}(\theta^i, \sigma^2 | y_i) = p(y_i | \theta^i, \sigma^2) = \mathcal{N}_{M_i}(y_i | f_i(t_i, \theta^i), \sigma^2 \mathbf{I}), \quad (i = 1, \dots, N).$$

Therefore, the likelihood function based on the  $N$  subjects  $y_{1:N} = \{y_i\}_{i=1}^N$  is

$$\mathcal{L}(\Theta, \sigma^2 | y_{1:N}) = \prod_{i=1}^N \mathcal{N}_{M_i}(y_i | f_i(t_i, \theta^i), \sigma^2 \mathbf{I}). \quad (10)$$

Now, we maximize the likelihood (10) with respect to the model matrix  $\Theta$  (5) given  $\sigma^2$  fixed:

$$\begin{aligned} \hat{\Theta} &= \operatorname{argmax}_{\Theta \in \mathbb{R}^{K \times N}} \log \mathcal{L}(\Theta, \sigma^2 | y_{1:N}) \\ &= \operatorname{argmin}_{\Theta \in \mathbb{R}^{K \times N}} \sum_{i=1}^N \|y_i - f_i(t_i, \theta^i)\|_2^2 \\ &= \operatorname{argmin}_{\Theta \in \mathbb{R}^{K \times N}} \sum_{i=1}^N \sum_{j=1}^{M_i} (y_{it} - f(t_{ij}; \theta^i))^2, \end{aligned} \quad (11)$$

where  $\|\cdot\|_2$  is the Euclidean norm.

The estimator  $\hat{\Theta} = [\hat{\theta}^1 \dots \hat{\theta}^i \dots \hat{\theta}^N] \in \mathbb{R}^{K \times N}$  can be obtained by various optimization techniques such as Newton-Raphson method or Gradient descent method [91]. Noting from the summation across  $i$  in (11),  $N$  estimators  $\{\hat{\theta}^i\}_{i=1}^N$  are independent. We can obtain an estimator of the variance  $\sigma^2$  by plugging  $\hat{\Theta}$  into the likelihood (10), and then maximize

with respect to the  $\sigma^2$ . To investigate a denoised temporal tendency for the trajectory  $\mathbf{y}_i$ , we can simply plug  $\hat{\theta}^i$  into the function  $f(t_{ij}; \theta^i)$  ( $j = 1, \dots, M_i$ ). To see a future pattern, we can extrapolate the function by extending the time index beyond the last time point  $t_{iM_i}$ . Eventually, the illustrated approach is based on traditional least squares estimation.

Unfortunately, there are three major drawbacks in this approach. First, it forfeits the opportunity to use ‘information borrowing’ [30] to improve a predictive accuracy due to the ignorance of Stage 2. What happens in Stage 2 (3) is to borrow strength across  $N$  individuals to produce a better estimator for  $\Theta$  than an estimator simply based on individual data. A similar issue can be found in the Clemente problem from [11] where the James-Stein estimator [92] better predicts than an individual hitter-based estimator. Another example applied to epidemic data can be found in [30]. Second, it does not well-aligned with the generic motivation to use the mixed effects models whose primary purpose is to understand “typical” values for the model parameters in  $f$ , representing whole subjects, which should be addressed by making an inference about the parameters  $\alpha$ ,  $B$ , and  $\Omega$ . Third, it only produces point estimates for the parameters, failing to describe the underlying uncertainty.

A remedy of the first two drawbacks is the consideration of Stage 1 (2) and 2 (3) hierarchically in a single model, leading to a frequentist version of nonlinear mixed effects models, which will be discussed in Subsection 4.2 and 4.3. To describe relevant uncertainty within the frequentist framework one may resort to bootstrap methods [39]. Another solution resolving all the three drawbacks at once is to incorporate Stage 1 (2), 2 (3), and 3 (4) in a fully Bayesian way, resulting in a Bayesian version of nonlinear mixed effects models, which is the main topic in this paper.

#### 4.2. Likelihood based on Stage 1 and 2 from vector-form (a)

A likelihood function based on vector-form (a) is derived. More specifically, we consider the frequentist setting where the assumptions of Stage 1-(6) and Stage 2-(7) are considered, while the parameters introduced in Stage 3-(9) are regarded as fixed quantities (i.e., no prior assumptions).

The individual model on Stage 1 (6) yields a conditional density  $p(\mathbf{y}_i | \theta^i, \sigma^2) = \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I})$  for each subject  $i = 1, \dots, N$ . Under the population assumption on Stage 2 (7), we have the density  $p(\theta_l | \alpha_l, \beta_l, \omega_l^2) = \mathcal{N}_N(\theta_l | \mathbf{1}\alpha_l + \mathbf{X}\beta_l, \omega_l^2 \mathbf{I})$  for each model parameter index  $l = 1, \dots, K$ . The joint density of  $(\mathbf{y}_{1:N}, \theta_{1:K})$  given parameters  $\sigma^2, \alpha, B$  and  $\Omega$  is a product-form distribution:

$$p(\mathbf{y}_{1:N}, \theta_{1:K} | \sigma^2, \alpha, B, \Omega) = \left\{ \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \right\} \cdot \left\{ \prod_{l=1}^K \mathcal{N}_N(\theta_l | \mathbf{1}\alpha_l + \mathbf{X}\beta_l, \omega_l^2 \mathbf{I}) \right\},$$

where  $\mathbf{y}_{1:N} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$  and  $\theta_{1:K} = \{\theta_1, \theta_2, \dots, \theta_K\}$ . Now, the next step is to integrate out the latent model parameters  $\theta_{1:K}$  from the density above to get a likelihood for the  $(\sigma^2, \alpha, B, \Omega)$ :

$$\mathcal{L}(\sigma^2, \alpha, B, \Omega | \mathbf{y}_{1:N}) = \int p(\mathbf{y}_{1:N}, \theta_{1:K} | \sigma^2, \alpha, B, \Omega) d\theta_{1:K}. \quad (12)$$

In most cases, the integral (12) is not tractable due to the non-linearity of the function  $f_i(\mathbf{t}_i, \theta^i)$  with respect to the  $\theta^i$ . Although it may be possible to use numerical techniques for the evaluation of the integral (12), this might require enormous computational effort, which is not really appreciated in the literature due to the high-dimensionality of the integral involving the  $KN$  dimensional model parameters  $\theta_{1:K}$ .

#### 4.3. Likelihood based on Stage 1 and 2 from vector-form (b)

A likelihood function based on vector-form (b) adopting  $i$ -indexing is derived here. As similar to Subsection 4.2, we preserve the assumption of Stage 1-(6) and Stage 2'-(8), but work with fixing the parameters in Stage 3-(9). In these specification, for each index  $i = 1, \dots, N$ , the individual model on Stage 1 (6) and population model on Stage 2 (8) lead



to densities  $p(\mathbf{y}_i|\boldsymbol{\theta}^i, \sigma^2) = \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, \boldsymbol{\theta}^i), \sigma^2\mathbf{I})$  and  $p(\boldsymbol{\theta}^i|\boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}) = \mathcal{N}_K(\boldsymbol{\theta}^i|\boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i, \boldsymbol{\Omega})$ , respectively. Thus, the joint density of  $(\mathbf{y}_i, \boldsymbol{\theta}^i)$  given parameters  $\sigma^2, \boldsymbol{\alpha}, \mathbf{B}$  and  $\boldsymbol{\Omega}$  is

$$\begin{aligned} p(\mathbf{y}_i, \boldsymbol{\theta}^i|\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}) &= p(\mathbf{y}_i|\boldsymbol{\theta}^i, \sigma^2) \cdot p(\boldsymbol{\theta}^i|\boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}) \\ &= \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, \boldsymbol{\theta}^i), \sigma^2\mathbf{I}) \cdot \mathcal{N}_K(\boldsymbol{\theta}^i|\boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i, \boldsymbol{\Omega}). \end{aligned}$$

Given the parameters  $(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega})$ , the ordered pairs in the collection  $\{(\mathbf{y}_i, \boldsymbol{\theta}^i)\}_{i=1}^N$  are conditionally independent across individuals. Therefore, a likelihood for the  $(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega})$  is based on the marginal density of  $\mathbf{y}_{1:N} = \{\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N\}$ :

$$\mathcal{L}(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}|\mathbf{y}_{1:N}) = \prod_{i=1}^N \int \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, \boldsymbol{\theta}^i), \sigma^2\mathbf{I}) \cdot \mathcal{N}_K(\boldsymbol{\theta}^i|\boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i, \boldsymbol{\Omega}) d\boldsymbol{\theta}^i \quad (13)$$

$$= \prod_{i=1}^N \int \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, \boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i + \boldsymbol{\eta}^i), \sigma^2\mathbf{I}) \cdot \mathcal{N}_K(\boldsymbol{\eta}^i|\mathbf{0}, \boldsymbol{\Omega}) d\boldsymbol{\eta}^i, \quad (14)$$

where the last equality is derived by using the change of variable (8). The last expression (14) is a standard mathematical formulation many frequentist approaches are based on: see Equation (3.2) from [14].

As the model parameter  $\boldsymbol{\theta}^i$  in (13) (or similarly,  $\boldsymbol{\eta}^i$  in (14) which is often called *random effect* in the frequentist framework) participates to the function  $f$  in a non-linear fashion, the integral generally cannot be obtained in a closed-form. Benefiting from a conditional independence [93], dimensionality of the  $N$  integrals (13) is much lower than that of the integral (12) based on vector-form (a). Analytically, the likelihood functions of the basic model (2) – (4) based on vector-form (a) (12) and vector-form (b) (13) may be equivalent. That being said, minimization of the two functions with respect to the parameters  $(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega})$  yields the same solution,  $(\hat{\sigma}^2, \hat{\boldsymbol{\alpha}}, \hat{\mathbf{B}}, \hat{\boldsymbol{\Omega}})$ , so-called maximum likelihood estimators (MLE).

We shall briefly discuss on MLE computations. One approach would be to perform a multivariate numerical integration (for e.g., Gauss-Hermite quadrature [94]) to each of the  $N$  integrals (13), and then obtain the MLE by maximizing the product of the  $N$  numerical integrals with respect to the parameters  $(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega})$  [95]. This approach turns out to be computationally so expensive and may have poor converge properties due to the following two reasons [96]. First, the numerical integration necessitates increasingly expensive iterative procedures within a MLE algorithm as the correlation of the model parameters (or equivalently, random effects) increases. Second, convergence property may be highly deteriorated when the number of model parameters  $K$  is large (i.e., high-dimensional integral) and the number of sampling times  $M_i$  is small (i.e., sparse data) due to the ‘curse of dimensionality’ [97].

A class of common approaches for the MLE computations is based on analytical approximation to each of the  $N$  integrals (14) [13,98–101], and some of them have been successfully adopted to industrial software like NONMEM [41,102] and SAS [103]. Here, we illustrate a key idea of the first-order method attributed to [98]. Let us define a mapping  $\mathbf{g}_i(\boldsymbol{\eta}^i) = f_i(\mathbf{t}_i, \boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i + \boldsymbol{\eta}^i) : A \subset \mathbb{R}^K \rightarrow \mathbb{R}^{M_i}$  for each subject  $i$ , where  $A$  is an open set with  $\mathbf{0} \in A$ . Suppose that  $\mathbf{g}_i(\boldsymbol{\eta}^i)$  is smooth on the set  $A$ : then, by Taylor’s theorem (page 375 of [104]), we have the best linear approximation of the mapping  $\mathbf{g}_i(\boldsymbol{\eta}^i)$  at the origin  $\mathbf{0}$  given by  $\mathbf{g}_i(\boldsymbol{\eta}^i) \approx \mathbf{g}_i(\mathbf{0}) + \mathbf{D}\mathbf{g}_i(\mathbf{0})\boldsymbol{\eta}^i$ , where  $\mathbf{D}\mathbf{g}_i(\mathbf{0}) \in \mathbb{R}^{M_i \times K}$  is the Jacobian matrix of  $\mathbf{g}_i(\boldsymbol{\eta}^i)$  at  $\mathbf{0}$ . Now, we shall replace the function  $\mathbf{g}_i(\boldsymbol{\eta}^i) = f_i(\mathbf{t}_i, \boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i + \boldsymbol{\eta}^i)$  in integral (14) with the resulting approximation  $\mathbf{g}_i(\mathbf{0}) + \mathbf{D}\mathbf{g}_i(\mathbf{0})\boldsymbol{\eta}^i$  for each  $i$  ( $i = 1, \dots, N$ ), leading to a closed-form expression

$$\begin{aligned}\tilde{\mathcal{L}}(\sigma^2, \alpha, \mathbf{B}, \Omega | \mathbf{y}_{1:N}) &= \prod_{i=1}^N \int \mathcal{N}_{M_i}(\mathbf{y}_i | \mathbf{g}_i(\mathbf{0}) + \mathbf{D}\mathbf{g}_i(\mathbf{0})\boldsymbol{\eta}^i, \sigma^2 \mathbf{I}) \cdot \mathcal{N}_K(\boldsymbol{\eta}^i | \mathbf{0}, \Omega) d\boldsymbol{\eta}^i \\ &= \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | \mathbf{g}_i(\mathbf{0}), \mathbf{D}\mathbf{g}_i(\mathbf{0})\Omega\mathbf{D}\mathbf{g}_i(\mathbf{0})^\top + \sigma^2 \mathbf{I}).\end{aligned}\quad (15)$$

To summarize, a linearization was used to convert the nonlinear mixed effects model to a linear mixed effects model, in some sense, equivalent to Lindley-Smith form [105]. This enables us to integrate out the random vector  $\boldsymbol{\eta}^i$  from the  $N$  integrals (15), deriving a marginal likelihood (15) to approximate the exact marginal likelihood (13). MLE  $(\hat{\sigma}^2, \hat{\alpha}, \hat{\mathbf{B}}, \hat{\Omega})$  can be obtained by jointly maximizing  $\tilde{\mathcal{L}}(\sigma^2, \alpha, \mathbf{B}, \Omega | \mathbf{y}_{1:N})$  (15) assuming the approximation is exact.

Another way to compute the MLE is through the use of expectation-maximization (EM) algorithm [106]. Borrowing terms from EM updating process [107],  $\mathbf{y}_i, \boldsymbol{\theta}^i, (\mathbf{y}_i, \boldsymbol{\theta}^i), (\sigma^2, \alpha, \mathbf{B}, \Omega)$  and  $p(\mathbf{y}_i, \boldsymbol{\theta}^i | \sigma^2, \alpha, \mathbf{B}, \Omega)$  (i.e., the integrand in (13)) can be viewed as observable incomplete data, missing data, complete data, unknown parameters, and density of complete data, respectively, for the  $i$ -th subject. The goal is to maximize the exact marginal likelihood  $\mathcal{L}(\sigma^2, \alpha, \mathbf{B}, \Omega | \mathbf{y}_{1:N})$  (13) by iterating E-step and M-step, leading to the MLE  $(\hat{\sigma}^2, \hat{\alpha}, \hat{\mathbf{B}}, \hat{\Omega})$ . The E-step computes a conditional expected log-likelihood of  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  based on the hierarchy (2) – (3), followed by the M-step that maximizes the function with respect to  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$ . The nonlinearity associated with the model matrix  $\boldsymbol{\Theta}$  (5) makes the E-step intractable. As a remedy, variant versions of the EM algorithm are proposed; see [108–111] for a technical detail applied to a hierarchy similar to the basic model. Among them, the scheme of stochastic approximation EM algorithm proposed by [112], splitting the E-step into two steps, namely a simulation step and a stochastic approximation step, is widely used in many applications for its numerical stability, fast computation, and theoretical soundness [113,114], which has been successfully deployed as industrial software including MONOLIX [115] as well as open source software such as R package NLMIXR [18].

## 5. Bayesian inference and implementation

### 5.1. Bayesian inference

We briefly overview two contrasting workflows of Bayesian and frequentist approaches for nonlinear mixed effects models before moving to a technical detail. Both settings allow the randomness in the model matrix  $\boldsymbol{\Theta}$  (5), but then, they diverge when it comes to how parameters  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  are treated. Bayesians treat  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  as random, while frequentists regard it as fixed. To conceptualize a subtlety arising from this difference, let us recap maximum likelihood approaches discussed in Subsection 4.2 and 4.3. There, the model matrix  $\boldsymbol{\Theta}$  was eventually integrated out from the joint density of  $(\mathbf{y}_{1:N}, \boldsymbol{\Theta})$ , either approximately or exactly, to derive a marginal likelihood of  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  from which the MLE  $(\hat{\sigma}^2, \hat{\alpha}, \hat{\mathbf{B}}, \hat{\Omega})$  is computed via various optimization methods. After that, frequentists apply standard Bayesian formulas, such as posterior density, posterior mean, and so on, to estimate  $\boldsymbol{\Theta}$  [7].

In contrast, to drive the Bayesian engine, one would need an appropriate prior  $\pi(\sigma^2, \alpha, \mathbf{B}, \Omega)$ . After that, the entire collection of parameters  $(\boldsymbol{\Theta}, \sigma^2, \alpha, \mathbf{B}, \Omega)$  will be updated through the Bayes' theorem post observing the data  $\mathbf{y}_{1:N}$ , leading to the posterior density  $\pi(\boldsymbol{\Theta}, \sigma^2, \alpha, \mathbf{B}, \Omega | \mathbf{y}_{1:N})$  [116]. The essence of the Bayesian viewpoint is that there is no logical distinction between  $\boldsymbol{\Theta}$  and  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$ , which are associated with the random and fixed effects, respectively, from the frequentist perspective. In Bayesian framework, both  $\boldsymbol{\Theta}$  and  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  are random quantities. It is important to point out that the likelihood principle is naturally incorporated in the Bayes' theorem [117]. Clearly, modern data complications such as enormous volume, large dimensionality, and multi-level structures may necessitate a sophistication on the prior specifications.

We are now in a position to describe the Bayesian analysis for the basic model (2)–(4), that assumed independence, *a priori*, for each parameter blocks  $\sigma^2$ ,  $\alpha$ ,  $B$ , and  $\Omega$ ,  $\pi(\sigma^2, \alpha, B, \Omega) = \pi(\sigma^2) \cdot \pi(\alpha) \cdot \pi(B) \cdot \pi(\Omega)$ . (Our logic below can be generalized to a more complex prior setting.) As discussed in Section 3, it is the discretion of modeler how she would treat the model matrix  $\Theta$  (5) with  $l$ -indexing  $\theta_{1:K} = \{\theta_l\}_{l=1}^K$  or  $i$ -indexing  $\theta^{1:N} = \{\theta^i\}_{i=1}^N$ , leading to vector forms (a) and (b), respectively. For the sake of readability, we illustrate the Bayesian inference by using the vector-form (a), but we will sometimes use the vector-form (b) when this seems more understandable.

A central task in the application of the Bayesian nonlinear mixed effects models is to evaluate the posterior density, or indeed to compute expectation with respect to the density:

$$\begin{aligned} \pi(\Theta, \sigma^2, \alpha, B, \Omega | \mathbf{y}_{1:N}) &= \frac{\pi(\mathbf{y}_{1:N}, \Theta, \sigma^2, \alpha, B, \Omega)}{p(\mathbf{y}_{1:N})} \propto \pi(\mathbf{y}_{1:N}, \Theta, \sigma^2, \alpha, B, \Omega) \\ &= \underbrace{p(\mathbf{y}_{1:N} | \Theta, \sigma^2)}_{\text{Stage 1}} \cdot \underbrace{\pi(\Theta | \alpha, B, \Omega)}_{\text{Stage 2}} \cdot \underbrace{\pi(\sigma^2) \cdot \pi(\alpha) \cdot \pi(B) \cdot \pi(\Omega)}_{\text{Stage 3}}, \end{aligned}$$

where the last equation can be detailed as follows

$$\left\{ \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \right\} \cdot \left\{ \prod_{l=1}^K \mathcal{N}_N(\theta_l | \mathbf{1}\alpha_l + \mathbf{X}\beta_l, \omega_l^2 \mathbf{I}) \pi(\alpha_l) \pi(\omega_l^2) \pi(\beta_l) \right\} \cdot \pi(\sigma^2). \quad (16)$$

From a Bayesian perspective, all inferential problems regarding the parameter  $(\Theta, \sigma^2, \alpha, B, \Omega)$  may be addressed in terms of the posterior distribution  $\pi(\Theta, \sigma^2, \alpha, B, \Omega | \mathbf{y}_{1:N})$  (16). Unfortunately, for almost all problems, the distribution is intractable. In such situations, we need to resort to approximation techniques, and these fall broadly into two classes, according to whether they rely on stochastic [118–121] or deterministic [122–125] approximations. See [126,127] for review papers for these techniques. In this article, we are mainly focused on the stochastic approximation. The basic idea behind the methodology is to construct a Markov chain whose stationary distribution is the posterior distribution  $\pi(\Theta, \sigma^2, \alpha, B, \Omega | \mathbf{y}_{1:N})$  (16).

## 5.2. Gibbs sampling algorithm

We resort to Markov chain Monte Carlo (MCMC) technique [126] to sample from the full joint density  $\pi(\Theta, \sigma^2, \alpha, B, \Omega | \mathbf{y}_{1:N})$  (16). Among many MCMC techniques, we use the Gibbs sampling algorithm [118,128] to exploit the conditional independence [93] induced by the hierarchical formulation. A generic Gibbs sampler would cycle in turn through each of the conditional distributions for the parameter blocks  $\Theta, \sigma^2, \alpha, B$ , and  $\Omega$  as follows:

**Step 1.** Sample  $\Theta$  from its full conditional distribution

$$\pi(\Theta | \sigma^2, \alpha, B, \Omega, \mathbf{y}_{1:N}) \propto \left\{ \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \right\} \cdot \left\{ \prod_{l=1}^K \mathcal{N}_N(\theta_l | \mathbf{1}\alpha_l + \mathbf{X}\beta_l, \omega_l^2 \mathbf{I}) \right\}; \quad (17)$$

**Step 2.** Sample  $\sigma^2$  from its full conditional distribution

$$\pi(\sigma^2 | \Theta, \alpha, B, \Omega, \mathbf{y}_{1:N}) \propto \left\{ \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \right\} \cdot \pi(\sigma^2); \quad (18)$$

**Step 3.** Sample  $\alpha$  from its full conditional distribution

$$\pi(\alpha | \sigma^2, \Theta, B, \Omega, \mathbf{y}_{1:N}) \propto \prod_{l=1}^K \mathcal{N}_N(\theta_l | \mathbf{1}\alpha_l + \mathbf{X}\beta_l, \omega_l^2 \mathbf{I}) \cdot \pi(\alpha_l); \quad (19)$$

**Step 4.** Sample  $\mathbf{B}$  from its full conditional distribution

$$\pi(\mathbf{B}|\boldsymbol{\Theta}, \sigma^2, \boldsymbol{\alpha}, \boldsymbol{\Omega}, \mathbf{y}_{1:N}) \propto \prod_{l=1}^K \mathcal{N}(\boldsymbol{\theta}_l | \mathbf{1}\alpha_l + \mathbf{X}\boldsymbol{\beta}_l, \omega_l^2 \mathbf{I}) \cdot \pi(\boldsymbol{\beta}_l); \quad (20)$$

**Step 5.** Sample  $\boldsymbol{\Omega}$  from its full conditional distribution

$$\pi(\boldsymbol{\Omega}|\boldsymbol{\Theta}, \sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \mathbf{y}_{1:N}) \propto \prod_{l=1}^K \mathcal{N}(\boldsymbol{\theta}_l | \mathbf{1}\alpha_l + \mathbf{X}\boldsymbol{\beta}_l, \omega_l^2 \mathbf{I}) \cdot \pi(\omega_l^2). \quad (21)$$

Sampling the model matrix  $\boldsymbol{\Theta}$  (5) at **Step 1** is independent of the choice of the priors, which we discuss shortly. On the other hand, sampling the parameters  $(\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega})$  at **Step 2, 3, 4,** and **5** depends on the prior choices for the parameters in Stage 3 (9); we discuss this topic in Section 6.

### 5.3. Parallel computation for model matrix

One of the most computer-intensive steps to implement the Gibbs sampler in Subsection 5.2 is **Step 1** to sample the model matrix  $\boldsymbol{\Theta} \in \mathbb{R}^{K \times N}$  (5), or equivalently its entries  $\{\theta_{li}\}_{l=1, i=1}^{K, N}$ , from the full conditional distributions  $\pi(\boldsymbol{\Theta}|\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}, \mathbf{y}_{1:N})$  (17). Clearly, the nonlinear participation of the model parameters to the function  $f$  makes the conditional distribution intractable, hence, non-conjugate sampling is unavoidable, which may suffer from a slow convergence. At the same times, due to the Markovian nature of the Gibbs algorithm, it is difficult to parallelize the whole steps of the Gibbs sampler, which creates difficulties in slower languages like R [129]. Nevertheless, the increasing number of parallel cores that are available at a very low price drives more and more interest in ‘parallel sampling algorithms’ that can benefit from the available parallel processing units on computers [130,131].

We suggest a framework of parallel computations to efficiently update the model matrix  $\boldsymbol{\Theta} \in \mathbb{R}^{K \times N}$ . This framework can be particularly appreciated under the setting of Bayesian nonlinear mixed effects models when the number of subjects  $N$  is a lot larger than the number of model parameters  $K$  ( $N \gg K$ ).


The first version of parallel sampling algorithms is based on scalar updating. For the derivation, we start with analyzing the full conditional posterior distribution of a single element  $\theta_{li}$  ( $l = 1, \dots, K; i = 1, \dots, N$ ):

$$\begin{aligned} \pi(\theta_{li}|-) &= \pi(\theta_{li}|\theta_{-li}, \sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}, \mathbf{y}_{1:N}) \propto \pi(\boldsymbol{\Theta}|\sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}, \mathbf{y}_{1:N}) \\ &\propto \left\{ \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki}), \sigma^2 \mathbf{I}) \right\} \cdot \left\{ \prod_{l=1}^K \prod_{i=1}^N \mathcal{N}(\theta_{li} | \alpha_l + \mathbf{x}_i^\top \boldsymbol{\beta}_l, \omega_l^2) \right\} \\ &\propto \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki}), \sigma^2 \mathbf{I}) \cdot \mathcal{N}(\theta_{li} | \alpha_l + \mathbf{x}_i^\top \boldsymbol{\beta}_l, \omega_l^2), \end{aligned} \quad (22)$$

where the notation  $\theta_{-li}$  represents the all the entries of  $\boldsymbol{\Theta}$  except for  $\theta_{li}$ , that is,  $\theta_{-li} = \{\theta_{li}\}_{l=1, i=1}^{K, N} - \{\theta_{li}\}$ . Here, we used a conventional notation in Bayesian computation: ‘ $\pi(\theta_{li}|-)$ ’ indicates the density  $\pi(\theta_{li}|\theta_{-li}, \sigma^2, \boldsymbol{\alpha}, \mathbf{B}, \boldsymbol{\Omega}, \mathbf{y}_{1:N})$ , where the notation ‘-’ in  $\pi(\theta_{li}|-)$  implies all the parameters except for the  $\theta_{li}$  in the basic model (2) – (4) along with  $N$  observations.

Note that the proportional part of the full conditional  $\pi(\theta_{li}|-)$  (22) only involves the  $i$ -th column vector of the model matrix  $\boldsymbol{\Theta}$  (5), that is,  $\boldsymbol{\theta}^i = (\theta_{1i}, \dots, \theta_{Ki})^\top \in \mathbb{R}^K$  in its analytic expression. This implies that we can update the  $K$  entries of the vector  $\boldsymbol{\theta}^i$  ( $i = 1, \dots, N$ ) independently across subjects. Parallel sampling algorithm can be completed by assigning a single CPU process to each of the subjects  $i$ . Within the step to sample the  $K$  entries of  $\boldsymbol{\theta}^i$ , Gibbs iterative procedure is required to update the scalar components. Authors [30,31] applied this technique to update the model matrix to train Bayesian nonlinear mixed effects

model to fit the dataset explained in 2.2 and 2.4. Figure 8 displays the schematic idea of the parallel sampling algorithm.



**Figure 8.** A pictorial description on the use of parallel computation to the basic model (2) – (4) to update the model matrix  $\Theta = [\theta^1 \dots \theta^i \dots \theta^N] \in \mathbb{R}^{K \times N}$  (5). In the parallel computation, a single CPU is assigned to an individual subject  $i = 1, \dots, N$ .

The second version of parallel sampling algorithms is based on vector updating. We analyze the full conditional posterior distribution of the vector  $\theta^i$  ( $i = 1, \dots, N$ ):

$$\begin{aligned} \pi(\theta^i | -) &= \pi(\theta^i | \theta^{-i}, \sigma^2, \alpha, B, \Omega, \mathbf{y}_{1:N}) \propto \pi(\Theta | \sigma^2, \alpha, B, \Omega, \mathbf{y}_{1:N}) \\ &\propto \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \cdot \mathcal{N}_K(\theta^i | \alpha + B\mathbf{x}_i, \Omega) \\ &\propto \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \cdot \mathcal{N}_K(\theta^i | \alpha + B\mathbf{x}_i, \Omega), \end{aligned} \quad (23)$$

where the notation  $\theta^{-i}$  represents the all the column vectors of  $\Theta$  except for  $\theta^i$ . As similar to the first version, we can use the parallel computation to update the model matrix  $\Theta$  by simultaneously sampling from the full-conditional density  $\pi(\theta^i | -)$  (23) across subjects by assigning one CPU to each individual.

#### 5.4. Elliptical slice sampler

Due the issue of non-conjugacy to sample from the univariate density  $\pi(\theta_{li} | -)$  (22) or  $K$ -dimensional density  $\pi(\theta^i | -)$  (23), the choice of a suitable MCMC method and further the choice of a proposal distribution is crucial for the convergence of the Markov chain simulated from the **Step 1** within the Gibbs sampler in Section 5.2. The Metropolis-Hastings (MH) algorithm [132,133] is the first solution to consider in such intractable situations: see the Algorithm 1 from [134]. In practice, the performances of the MH algorithm are highly dependent on the choice of the proposal density [135]. In the past decades, numerous MH-type algorithms to improve computational efficiency have been developed, and these fall broadly into two classes, according to whether the proposal density reflects a gradient information [119,136–138] or not [120,139]. In specific, the gradient information, here, refers to the first-order derivative of the minus of the log of the target density (i.e.,  $\nabla U(\theta_{li}) = -\nabla \log \pi(\theta_{li} | -) \in \mathbb{R}$  or  $\nabla U(\theta^i) = -\nabla \log \pi(\theta^i | -) \in \mathbb{R}^K$ , where the notation  $\nabla$  represents the gradient operator). Typically, gradient-based samplers are attractive in terms of rapid exploration of the state space, but the cost of the gradient computation can be prohibitive when the sample size  $N$  or model dimension  $K$  is extremely large [140,141]. Fortunately, this requirement can be made less burdensome by using automatic differentiation [142].

In the present subsection, we introduce an efficient gradient-free sampling technique, the elliptical slice sampler (ESS) proposed by [120], to simulate a Markov chain from the density  $\pi(\theta_{li} | -)$  (22). (The sampling logic can be directly applied to the situation to sample from the density  $\pi(\theta^i | -)$  (23) by simply replacing  $\theta_{li}$  by  $\theta^i$  and changing the dimension of



relevant distributions, stochastic processes, etc, from 1 to  $K$ .) Conceptually, MH and ESS algorithms are similar in that both comprise two steps: a proposal step and a criterion step. A difference between the two algorithms arises in the criterion step. If a new candidate does not pass the criterion, then MH algorithm takes the current state as the next state: whereas, ESS re-proposes a new candidate until rejection does not take place, rendering the algorithm rejection-free. The utility of ESS can be appreciated when an analytic expression of a target density can be factored to have a Gaussian prior distribution. Unlike the traditional MH algorithm that requires the proposal variance or density, ESS is fully automated, no tuning is required.

To adapt the ESS to our example, we re-write the density  $\pi(\theta_{li}|-)$  (22) as the following form:

$$\pi(\theta_{li}|-) = \frac{1}{Z} \mathcal{L}(\theta_{li}) \cdot \mathcal{N}(\theta_{li}|\alpha_l + \mathbf{x}_l^\top \boldsymbol{\beta}_l, \omega_l^2), \quad (24)$$

where  $\mathcal{L}(\theta_{li}) = \exp \{-\|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i, \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 / (2\sigma^2)\}$ , and  $Z$  is the normalization constant. Introducing the notation  $\mathcal{L}(\theta_{li})$  is our intention that we shall treat the function as a likelihood part temporarily only when sampling from the density (24). Alternatively, one can proceed with the choice  $\mathcal{L}'(\theta_{li}) = \mathcal{N}_{M_i}(\mathbf{y}_i | \mathbf{f}_i(\mathbf{t}_i, \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki}), \sigma^2 \mathbf{I})$  as a likelihood part, which then change the normalization constant accordingly. We recommend to use the simplest functional form for the likelihood part, if possible, to reduce computation cost.

Algorithm 1 summarizes the ESS in an algorithmic form where the situation is at the  $(s+1)$ -th iteration of the Gibbs sampler. Therefore, the goal is to draw  $\theta_{li}^{(s+1)}$  from the target density  $\pi(\theta_{li}|-)$  ( $l = 1, \dots, K; i = 1, \dots, N$ ) (24), where we already have  $\theta_{li}^{(s)}$  as the current state for the target variable  $\theta_{li}$  realized from the  $s$ -th iteration:

### 5.5. Metropolis adjusted Langevin algorithm

We introduce Metropolis adjusted Langevin algorithm (MALA) [136,137] which is popular for its use of problem-specific proposal distribution based on the gradient information of the target density. The main idea of MALA is to use Langevin dynamics to construct the Markov chain. To adapt the sampling technique to our example, we re-write the density  $\pi(\theta_{li}|-)$  (22) as the following form:

$$\pi(\theta_{li}|-) = \frac{\exp(-U(\theta_{li}))}{\int_{-\infty}^{\infty} \exp(-U(\theta_{li})) d\theta_{li}} \quad \text{for all } \theta_{li} \in \mathbb{R}, \quad (25)$$

where  $U(\theta_{li}) = (1/\{2\sigma^2\}) \cdot \|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i; \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 + (1/\{2\omega_l^2\}) \cdot (\theta_{li} - \alpha_l - \mathbf{x}_l^\top \boldsymbol{\beta}_l)^2$ . Now, we consider a stochastic differential equation [143] that characterizes the evolution of the Langevin diffusion with the drift term set by the gradient of the log of the density (25):

$$d\theta_{li}(t) = \nabla \log \pi(\theta_{li}(t)|-) dt + \sqrt{2} dW(t) = -\nabla U(\theta_{li}(t)) dt + \sqrt{2} dW(t), \quad (26)$$

where  $\{W(t) | t \geq 0\}$  is a standard 1-dimensional Wiener process, or Brownian motion [144]. In (26),  $t$  indexes a fictitious continuous time.  $\nabla$  represents the gradient operator with respect to  $\theta_{li}$ . Under fairly mild conditions on the function  $U(\theta_{li})$ , the equation (26) has a strong solution  $\{\theta_{li}(t) | t \geq 0\}$  that is a Markov process [145]. Furthermore, it can be shown that the distribution of  $\{\theta_{li}(t) | t \geq 0\}$  converges to the invariant distribution  $\pi(\theta_{li}|-)$  (25) as  $t \rightarrow \infty$ .

Since solving the equation (26) is very difficult, a first-order Euler-Maruyama discretization [146] is used to approximate the solution to the equation:

$$\theta_{li}^{[s+1]} \leftarrow \theta_{li}^{[s]} - \delta \cdot \nabla U(\theta_{li}^{[s]}) + \sqrt{2\delta} Z, \quad Z \sim \mathcal{N}(0, 1), \quad (27)$$

**Algorithm 1:** ESS to sample from  $\pi(\theta_{li}|-)$  (22)**Goal :** Sampling from the full conditional posterior distribution

$$\pi(\theta_{li}|-) \propto \mathcal{L}(\theta_{li}) \cdot \mathcal{N}(\theta_{li}|\mu_{li}, \omega_l^2),$$

where  $\mathcal{L}(\theta_{li}) = \exp \{-\|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i, \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 / (2\sigma^2)\}$  and  $\mu_{li} = \alpha_l + \mathbf{x}_i^\top \boldsymbol{\beta}_l$ .

**Input :** Current state  $\theta_{li}^{(s)}$ .**Output :** A new state  $\theta_{li}^{(s+1)}$ .

- Choose an ellipse centered at  $\mu_{li}$ :  $v \sim \mathcal{N}(\mu_{li}, \omega_l^2)$ .
- Define a criterion function:

$$\alpha(\theta_{li}, \theta_{li}^{(s)}) = \min\{\mathcal{L}(\theta_{li}) / \mathcal{L}(\theta_{li}^{(s)}), 1\} : \mathbb{R} \rightarrow [0, 1].$$

- Choose a threshold and fix:  $u \sim \text{Unif}[0, 1]$ .
- Draw an initial proposal  $\theta_{li}^*$ :

$$\phi \sim \text{Unif}(-\pi, \pi];$$

$$\theta_{li}^* = (\theta_{li}^{(s)} - \mu_{li}) \cos \phi + (v - \mu_{li}) \sin \phi + \mu_{li}.$$

- if** ( $u < \alpha(\theta_{li}^*, \theta_{li}^{(s)})$ ) **{**  $\theta_{li}^{(s+1)} = \theta_{li}^*$  **}** **else** {  
 Define a bracket :  $(\phi_{\min}, \phi_{\max}] = (-\pi, \pi]$ .  
**while** ( $u \geq \alpha(\theta_{li}^*, \theta_{li}^{(s)})$ ) {  
 Shrink the bracket and try a new point :  
**if** ( $\phi > 0$ )  $\phi_{\max} = \phi$  **else**  $\phi_{\min} = \phi$   
 $\phi \sim \text{Unif}(\phi_{\min}, \phi_{\max}]$   
 $\theta_{li}^* = (\theta_{li}^{(s)} - \mu_{li}) \cos \phi + (v - \mu_{li}) \sin \phi + \mu_{li}$ .  
**}**  
 $\theta_{li}^{(s+1)} = \theta_{li}^*$   
**}**

where  $\delta$  is the step size of discretization, and  $[s]$  indexes the discrete time steps. This recursive update defines the Langevin Monte Carlo algorithm. Typically, to handle the discretization error and satisfy the detailed balance [147] to make Markov chain converge to the target distribution  $\pi(\theta_{li}|-)$  (25), the MH correction is needed. Algorithm 2 details MALA to sample from the  $\pi(\theta_{li}|-)$  ( $l = 1, \dots, K; i = 1, \dots, N$ ) (25):

## 5.6. Hamiltonian Monte Carlo

We introduce Hamiltonian Monte Carlo (HMC) algorithm that employs Hamiltonian dynamics to efficiently explore the parameter space [119,138]. Among many MH-type sampling algorithms, HMC has been recognized as one of the most effective algorithms due to its rapid mixing rate and small discretization error. By that reason, HMC has been deployed as the default sampler in many open packages such as STAN [148] and TENSORFLOW [149]. A key idea of HMC distinctive from ESS and MALA is the introduction of an auxiliary momentum variable, which is typically assumed to follow as a Gaussian distribution and independent of the target variable. By doing so, the HMC can produce distant proposals for the target variable, thereby avoiding the slow exploration of the state space that results from the diffusive behavior of simple random-walk proposals.

**Algorithm 2:** MALA to sample from  $\pi(\theta_{li}|-)$  (22)**Goal :** Sampling from the full conditional posterior distribution

$$\pi(\theta_{li}|-) \propto \exp(-U(\theta_{li})),$$

where

$$U(\theta_{li}) = \|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i; \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 / (2\sigma^2) + (\theta_{li} - \alpha_l - \mathbf{x}_i^\top \boldsymbol{\beta}_l)^2 / (2\omega_l^2).$$

**Input :** Current state  $\theta_{li}^{(s)}$  and step size  $\delta$ .**Output :** A new state  $\theta_{li}^{(s+1)}$ .

a. Define a criterion function:

$$\alpha(\theta_{li}, \theta_{li}^{(s)}) = \min \left\{ \frac{\exp(-U(\theta_{li}))}{\exp(-U(\theta_{li}^{(s)}))} \cdot \frac{\mathcal{J}(\theta_{li}^{(s)}|\theta_{li})}{\mathcal{J}(\theta_{li}|\theta_{li}^{(s)})}, 1 \right\} : \mathbb{R} \rightarrow [0, 1].$$

b. Choose a threshold  $u$ :  $u \sim \text{Unif}[0, 1]$ .c. Draw a proposal  $\theta_{li}^*$ :

$$\theta_{li}^* \sim \mathcal{J}(\theta_{li}|\theta_{li}^{(s)}) = \mathcal{N}(\theta_{li}|\theta_{li}^{(s)} - \delta \cdot \nabla U(\theta_{li}^{(s)}), 2\delta).$$

e. **if** ( $u < \alpha(\theta_{li}^*, \theta_{li}^{(s)})$ ) **{**  $\theta_{li}^{(s+1)} = \theta_{li}^*$  **}** **else** **{**  $\theta_{li}^{(s+1)} = \theta_{li}^{(s)}$  **}**

We adapt the HMC to our example. We shall first take a look at a joint density:

$$\begin{aligned} \pi(\theta_{li}, \phi_{li}|-) &= \pi(\theta_{li}|-) \cdot \pi(\phi_{li}) \\ &= \frac{\exp(-U(\theta_{li}))}{\int_{-\infty}^{\infty} \exp(-U(\theta_{li})) d\theta_{li}} \cdot \frac{1}{\sqrt{2\pi m_{li}}} \exp\left(-\frac{\phi_{li}^2}{2m_{li}}\right) \\ &\propto \exp(-U(\theta_{li}) - K(\phi_{li})) \quad \text{for all } (\theta_{li}, \phi_{li}) \in \mathbb{R} \times \mathbb{R}, \end{aligned} \quad (28)$$

where  $U(\theta_{li}) = (1/\{2\sigma^2\}) \cdot \|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i; \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 + (1/\{2\omega_l^2\}) \cdot (\theta_{li} - \alpha_l - \mathbf{x}_i^\top \boldsymbol{\beta}_l)^2$  and  $K(\phi_{li}) = \phi_{li}^2 / (2m_{li})$ . The auxiliary variable  $\phi_{li}$  is distributed according to the univariate Gaussian distribution  $\mathcal{N}(\phi_{li}|0, m_{li})$  with variance  $m_{li}$ . Note that it holds  $\int \pi(\theta_{li}, \phi_{li}|-) d\phi_{li} = \pi(\theta_{li}|-)$  due to the independence between  $\theta_{li}$  and  $\phi_{li}$ . Therefore, our ultimate goal is to sample from the joint density  $\pi(\theta_{li}, \phi_{li}|-)$  (28), and take only  $\theta_{li}$  by marginalization.

Noting from (28), the negative of joint log-posterior is

$$H(\theta_{li}, \phi_{li}) = U(\theta_{li}) + K(\phi_{li}), \quad \text{for all } (\theta_{li}, \phi_{li}) \in \mathbb{R} \times \mathbb{R}. \quad (29)$$

The physical analogy of the bivariate function  $H(\theta_{li}, \phi_{li}) : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$  (29) is a Hamiltonian [138,150], which describe the sum of a potential energy  $U(\theta_{li})$ , defined at the position  $\theta_{li}$ , and a kinetic energy  $K(\phi_{li}) = \phi_{li}^2 / (2m_{li})$ , where the auxiliary variable  $\phi_{li}$  can be interpreted as a momentum variable and the variance  $m_{li}$  denotes a mass.

Now, we construct a Hamiltonian system by taking a derivative of  $H$  (29) with respect to  $\theta_{li}$  and  $\phi_{li}$ , and introducing a continuous fictitious time  $t$ :

$$\frac{d\theta_{li}(t)}{dt} = \frac{\partial H(\theta_{li}, \phi_{li})}{\partial \phi_{li}} = \frac{\partial K(\phi_{li})}{\partial \phi_{li}} = \frac{\phi_{li}(t)}{m_{li}}, \quad (30)$$

$$\frac{d\phi_{li}(t)}{dt} = -\frac{\partial H(\theta_{li}, \phi_{li})}{\partial \theta_{li}} = -\frac{\partial U(\theta_{li})}{\partial \theta_{li}} = -\nabla U(\theta_{li}(t)), \quad (31)$$

where  $\nabla$  represents the gradient operator with respect to  $\theta_{li}$ .

The Hamiltonian system (30) - (31) has three nice properties. Assume that  $(\theta_{li}(t), \phi_{li}(t)) : [a, b] \rightarrow \mathbb{R} \times \mathbb{R}$  is a solution curve of the system, where  $a, b \in \mathbb{R} \cup \{\pm\infty\}$ . Then, the following relationships hold:

- (a) *preservation of total energy*:  $H(\theta_{li}(t), \phi_{li}(t)) = H(\theta_{li}(0), \phi_{li}(0))$  for all  $t \in [a, b]$ ;
- (b) *preservation of volume*:  $d\theta_{li}(t)d\phi_{li}(t) = d\theta_{li}(0)d\phi_{li}(0)$  for all  $t \in [a, b]$ ;
- (c) and *time reversibility*: The mapping  $T_s$  from state at  $t$ ,  $(\theta_{li}(t), \phi_{li}(t))$ , to the state at time  $t + s$ ,  $(\theta_{li}(t + s), \phi_{li}(t + s))$ , is one-to-one, and hence has an inverse  $T_{-s}$ .

Three properties are eventually related with the following nice properties of the HMC: (a) a high probability of acceptance of proposals; (b) a simple analytic form of acceptance ratio (no need to consider a hard-to-compute Jacobian factor); and (c) a detailed balance with respect to the target density  $\pi(\theta_{li}|-)$ . For a detailed description and extensive review see [119].

For practical applications, the differential equation system (30) - (31) cannot be solved analytically and numerical methods are required. As Hamiltonian  $H$  in the system is separable (or equivalently, the joint density  $\pi(\theta_{li}, \phi_{li}|-)$  is factorizable), to traverse the state space more efficiently, the leapfrog integrator method is typically used, which involves a discretized step of the dynamics. As similar to the construction of MALA, the discretization errors arising from the leapfrog integration are addressed by MH correction step. Algorithm 3 details the HMC to sample from the target density  $\pi(\theta_{li}|-) = \int \pi(\theta_{li}, \phi_{li}|-)d\phi_{li}$  (22). In the algorithm,  $(s)$  indexes the sampling iteration within the Gibbs sampler, while  $[d]$  represents the index introduced due to the discretization.

One caveat in the HMC is that no matter what we accept or reject the proposal, we draw a new momentum from the kinetic energy at every iteration. To check this, see the Step a in Algorithm 3, where  $\phi_{li}^{[0]}$  is drawn from the kinetic density  $K(\phi_{li}) \propto \mathcal{N}(0, m_{li})$ . The momentum  $\phi_{li}^{[0]}$  is only used to formulate the initial pair  $(\theta_{li}^{[0]}, \phi_{li}^{[0]})$ , where  $\theta_{li}^{[0]}$  is the current target state  $\theta_{li}^{(s)}$ , that will be guided by Hamiltonian dynamics (30) - (31) via leapfrog integrator, eventually reaching to the last pair  $(\theta_{li}^{[L]}, \phi_{li}^{[L]})$  which is used as the proposal  $(\theta_{li}^*, \phi_{li}^*)$ . The momentum  $\phi_{li}^{[0]}$  is deleted and we will draw a new momentum in the next iteration. This independent drawing of the momentum is the engine that enables HMC to produce distant proposals, but nevertheless maintains a high probability of acceptance.

The naive HMC (Algorithm 3) requires the users to specify at least three parameters: a step size  $\delta$ , a number of steps  $L$ , and a mass  $m_{li}$ , for which to run a simulated Hamiltonian system. A poor choice of either of these parameters will result in a dramatic drop in the efficiency HMC. No-U-Turn Sampler (NUTS) developed by [151] is an extension of HMC, which is designed to automatically turn the parameters  $(\delta, L)$  while fixing  $m_{li} = 1$ , making it possible to run NUTS with no hand-tuning at all. HMC and NUTS are general-purpose inference engines deployed in STAN.

We would like to highlight a difference between MALA (Algorithm 2) and HMC (Algorithm 3). Although both algorithms utilizes the gradient information (that is,  $\nabla U(\theta_{li}) = -\nabla \log \pi(\theta_{li}|-)$ ), the former is based on stochastic differential equation (26) and the latter is based on ordinary differential equation (30) - (31). From the algorithmic perspective, MALA exhibits a single loop structure and is designed to directly employ the discretization of the underlying Langevin dynamics: see that the index  $[s]$  resulting from the discretization in (27) is directly used as the sampling index  $(s)$  in Algorithm 2. On the other hand, HMC has a double loop structure: the inner loop (i.e., Step c in Algorithm 3) solves the Hamiltonian dynamics (30) - (31) to make a proposal, while the outer loop judges the proposals. The index  $[d]$  in the inner loop, resulting from the leapfrog integrator, and the sampling index  $(s)$  of the outer loop in Algorithm 3 are not related [152]. Therefore, one can set the number of steps  $L$  for the leapfrog integrator by an arbitrary integer, which does not exist in the MALA.

---

**Algorithm 3:** HMC to sample from  $\pi(\theta_{li}|-) = \int \pi(\theta_{li}, \phi_{li}|-)d\phi_{li}$  (22)

---

**Goal :** Sampling from the full conditional posterior distribution

$$\pi(\theta_{li}|-) = \int \pi(\theta_{li}, \phi_{li}|-)d\phi_{li} \propto \exp(-U(\theta_{li})),$$

where  $\pi(\theta_{li}, \phi_{li}|-) \propto \exp(-H(\theta_{li}, \phi_{li})) = \exp(-U(\theta_{li}) - K(\phi_{li}))$  with  $U(\theta_{li}) = \|\mathbf{y}_i - \mathbf{f}_i(\mathbf{t}_i; \theta_{1i}, \dots, \theta_{li}, \dots, \theta_{Ki})\|_2^2 / (2\sigma^2) + (\theta_{li} - \alpha_l - \mathbf{x}_i^\top \boldsymbol{\beta}_l)^2 / (2\omega_l^2)$  and  $K(\phi_{li}) = \phi_{li}^2 / (2m_{li})$ .

**Input :** Current state  $\theta_{li}^{(s)}$ , step size  $\delta$ , number of steps  $L$ , and mass  $m_{li}$ .

**Output :** A new state  $\theta_{li}^{(s+1)}$ .

- Generate the initial momentum with mass  $m_{li}$ :  $\phi_{li}^{[0]} \sim \mathcal{N}(0, m_{li})$ .
- Define a criterion function:

$$\alpha((\theta_{li}, \phi_{li}), (\theta_{li}^{(s)}, \phi_{li}^{[0]})) = \min \left\{ \frac{\exp(-H(\theta_{li}, \phi_{li}))}{\exp(-H(\theta_{li}^{(s)}, \phi_{li}^{[0]}))}, 1 \right\} : \mathbb{R} \times \mathbb{R} \rightarrow [0, 1].$$

- Simulate discretization of Hamiltonian dynamics (30) – (31):
    - Set the initial pair of the solution curve:  $(\theta_{li}^{[0]}, \phi_{li}^{[0]}) = (\theta_{li}^{(s)}, \phi_{li}^{[0]})$ .
    - Make a half step for the momentum:  $\phi_{li}^{[0]} \leftarrow \phi_{li}^{[0]} - (\delta/2)\nabla U(\theta_{li}^{[0]})$ .
    - Alternate full steps for position and momentum:
 

**for** ( $d = 1, \dots, L$ ) {

Update position:  $\theta_{li}^{[d]} \leftarrow \theta_{li}^{[d-1]} + (\delta/m_{li})\phi_{li}^{[d-1]}$ .

Update momentum: **if** ( $d = L$ ) {  $\phi_{li}^{[d]} \leftarrow \phi_{li}^{[d-1]} - \delta\nabla U(\theta_{li}^{[d]})$ . }

}
    - Make a half step for momentum:  $\phi_{li}^{[L]} \leftarrow \phi_{li}^{[L-1]} - (\delta/2)\nabla U(\theta_{li}^{[L]})$ .
    - Negate the last momentum:  $\phi_{li}^{[L]} \leftarrow -\phi_{li}^{[L]}$ .
    - Set the last pair of the solution curve as the proposal:  $(\theta_{li}^*, \phi_{li}^*) = (\theta_{li}^{[L]}, \phi_{li}^{[L]})$ .
  - Choose a threshold  $u$ :  $u \sim \text{Unif}[0, 1]$ .
  - if** ( $u < \alpha((\theta_{li}^*, \phi_{li}^*), (\theta_{li}^{(s)}, \phi_{li}^{[0]}))$ ) {  $\theta_{li}^{(s+1)} = \theta_{li}^*$  } **else** {  $\theta_{li}^{(s+1)} = \theta_{li}^{(s)}$  }
- 

## 6. Prior options

### 6.1. Priors for variance

Provided the assumption of the basic form (2) – (4), the random error terms  $\{\epsilon_{ij}\}_{i=1, j=1}^{N, M_i}$  in Stage 1 and  $\{\eta_{li}\}_{l=1, i=1}^{K, N}$  in Stage 2 are the stochastic sources of (remaining) intra-individual and inter-individual variabilities, respectively [8]. Both terms are assumed to follow univariate Gaussian distributions in the basic model. This assumption can be generalized to multivariate Gaussian distribution,  $t$ -distribution, mixture of Gaussian distributions, etc, depending on the exhibition of the data or prior guess of perturbation associated with model matrix  $\Theta$  (5) [153].

Recall that the basic model assumes that data-level errors  $\epsilon_{ij}$  are distributed according to  $\mathcal{N}(0, \sigma^2)$  with variance  $\sigma^2$ , independently across times  $j = 1, \dots, M_i$  and subjects  $i = 1, \dots, N$ . (We discuss about the  $\eta$ -term in Subsection 6.3.) Therefore, the standard deviation  $\sigma$  describes a vertical difference (i.e., measurement error) between the observation  $y_{ij}$  and theory  $f(t_{ij}; \boldsymbol{\theta}^i)$  across time and individuals. We can generalize the basic setting by replacing  $\sigma^2$  with (a)  $\sigma_i^2$  ( $i = 1, \dots, N$ ) or (b)  $\sigma_{ij}^2$  ( $i = 1, \dots, N; j = 1, \dots, M_i$ ) to accommodate the heterogeneity (a) across subjects and (b) across subjects and time, respectively, provided sufficiently large sampling times  $M_i$  [28].



For any prior  $\pi(\sigma^2)$ , the full conditional posterior distribution of  $\sigma^2$  (18) is given as

$$\pi(\sigma^2 | \Theta, \alpha, B, \Omega, \mathbf{y}_{1:N}) \propto (\sigma^2)^{-\sum_{i=1}^N M_i/2} \exp\left(-\frac{1}{2\sigma^2} \sum_{i=1}^N \|\mathbf{y}_i - f_i(\mathbf{t}_i, \boldsymbol{\theta}^i)\|_2^2\right) \cdot \pi(\sigma^2), \quad (32)$$

where  $\|\mathbf{a}\|_2^2$  represents the Euclidean norm of the vector  $\mathbf{a}$ .

Popularly used priors  $\pi(\sigma^2)$  (or  $\pi(\sigma)$ ) are (i) the Jeffreys prior  $\pi(\sigma^2) \propto 1/\sigma^2$  [154]; (ii) inverse-gamma prior  $\pi(\sigma^2) = \mathcal{IG}(a_{\sigma^2}, b_{\sigma^2})$  with shape  $a_{\sigma^2} > 0$  and scale  $b_{\sigma^2} > 0$ ; and (iii) half-Cauchy prior  $\pi(\sigma) = \mathcal{C}^+(0, b_\sigma) = \{2/(\pi b_\sigma)\} \cdot 1/\{1 + (\sigma/b_\sigma)^2\}$  with scale  $b_\sigma > 0$ . Note that half-Cauchy distribution should be given to the standard deviation  $\sigma$ , not variance  $\sigma^2$ . The first two prior options lead to the conjugate update to sample from the density  $\pi(\sigma^2 | -)$  (32). Although the third one induces non-conjugate update to sample from the density  $\pi(\sigma | -)$ , computationally efficient sampling can be constructed by using parameter expansion technique [155] or slice sampler [121].

## 6.2. Priors for intercept and coefficient vector

One of the central goals of using nonlinear mixed effects models is to identify significant covariates among the  $P$  covariates  $\mathbf{x} = (x_1, \dots, x_P)^\top$ , explaining each of the model parameter  $\theta_l$  ( $l = 1, \dots, K$ ). This is because the function  $f$  in Stage 1 (2) is typically derived from a differential equation system. Such a differential equation has model parameters  $\{\theta_l\}_{l=1}^K$  which controls the dynamic of the solution of the system, and how the parameters are related with covariates is vital to understand causality. For example, in PK analysis, understanding whether and to what extent weight, renal status, disease status, etc, are associated with drug clearance may dictate how these factors can be considered in a dosing schedule.

We explain popularly used priors for the intercept and coefficient vector by taking the vector-form (a) [Stage 1-(6), Stage 2-(7), and Stage 3-(9)] because it directly embeds the framework of linear regression. For each model parameter index  $l = 1, \dots, K$ , we re-write the equation (7) for the purpose of illustration:

$$\begin{bmatrix} \theta_{l1} \\ \vdots \\ \theta_{li} \\ \vdots \\ \theta_{lN} \end{bmatrix} = \begin{bmatrix} \alpha_l \\ \vdots \\ \alpha_l \\ \vdots \\ \alpha_l \end{bmatrix} + \begin{bmatrix} x_{11} & \cdots & x_{1b} & \cdots & x_{1P} \\ \vdots & & \vdots & & \vdots \\ x_{i1} & \cdots & x_{ib} & \cdots & x_{iP} \\ \vdots & & \vdots & & \vdots \\ x_{N1} & \cdots & x_{Nb} & \cdots & x_{NP} \end{bmatrix} \begin{bmatrix} \beta_{l1} \\ \vdots \\ \beta_{lb} \\ \vdots \\ \beta_{lP} \end{bmatrix} + \begin{bmatrix} \eta_{l1} \\ \vdots \\ \eta_{li} \\ \vdots \\ \eta_{lN} \end{bmatrix} \quad (33)$$

where  $\boldsymbol{\eta}_l = (\eta_{l1}, \dots, \eta_{li}, \dots, \eta_{lN})^\top \sim \mathcal{N}_N(\mathbf{0}, \omega_l^2 \mathbf{I})$ . By the assumption (9), we have priors  $\alpha_l \sim \pi(\alpha_l)$  and  $\beta_{lb} \sim \pi(\beta_{lb})$  ( $b = 1, \dots, P$ ). Note that the equation (33) is a Bayesian multivariate linear regression (page 149 of [156]), and the only difference from the usual context is that the response vector in (33) is latent. Therefore, almost all Bayesian regression techniques [157,158] can be used to the latent regression (33) provided that the model matrix  $\Theta$  (5) is efficiently realized in **Step 1** within the Gibbs sampler.

Default choice for the prior  $\pi(\alpha_l)$  is the flat prior  $\pi(\alpha_l) \propto 1$ , also called a uniform prior. Alternatively, a diffuse Gaussian prior  $\pi(\alpha_l) = \mathcal{N}(a_l, b_l^2)$  is also often used by fixing  $b_l$  to be sufficiently large (saying  $b_l = 10$  or  $100$ ) and  $a_l = 0$ . In either case, the full conditional density (19) enjoys the conjugate update, hence, **Step 3** in the Gibbs sampler in Subsection 5.2 seldom imposes computational burden. The idea behind the use of (nearly) non-informative priors for the intercept is that such priors induce almost minimal degree of Bayesian shrinkage, hence, allow the data to have (nearly) maximum effect on the posterior estimate for the intercept [159].

Now, we discuss on the Bayesian analysis for the coefficients. There are numerous choices for the prior of the coefficient vector  $\boldsymbol{\beta}_l = (\beta_{l1}, \dots, \beta_{lb}, \dots, \beta_{lP}) \sim \pi(\boldsymbol{\beta}_l)$ , which is not surprising because the linear regression is arguably one of the most researched topics

in statistics. Here, we suggest some popular priors whose main utility fall broadly into two classes, according to whether the design matrix  $\mathbf{X} \in \mathbb{R}^{N \times P}$  in (33) is tall ( $N \geq P$ ) or fat ( $N < P$ ). In regression theory, the former setting is referred to as low-dimensional regression, and latter one is called high-dimensional regression [160–162].

Under the *tall design* ( $N \geq P$ ), particularly when the number of subjects  $N$  is much larger than the number of covariates  $P$  ( $N \gg P$ ), one important theoretical consideration is that, it is expected to see Bernstein–von Mises type results [163–165] on the posterior inference for coefficients  $\beta_l$ . Roughly speaking, the theorem, sometimes called the “Bayesian Central Limit Theorem”, states that the posterior distribution of  $\beta_l$  is approximately a normal distribution following a likelihood theory as sample size  $N$  goes to infinity for any prior choice  $\pi(\beta_l)$  under certain regularity conditions. In reality, it is possible that due to an ill-conditioned design matrix  $\mathbf{X}$ , a misspecification of error distribution for  $\eta_{li}$ , a small sample size  $N$ , an inappropriate choice of prior  $\pi(\beta_l)$ , etc, the Bernstein-von Mises Phenomenon (page 151 from [166]) may not be empirically observed. But even such cases, it is known the influence of the prior distribution diminishes as  $N$  grows, which means that using different priors for  $\beta_l$  may not sensitively change the resulting Bayesian inferences about  $\beta_l$ , and furthermore, the inference outcomes obtained by Bayesian and frequentist methods agree in most instances under the tall design. For example, see results of [167] and [28] for pharmacokinetic applications. Some possible options for prior  $\pi(\beta_l)$  are (i) flat prior  $\pi(\beta_l) \propto 1$ ; (ii) Gaussian diffuse prior  $\pi(\beta_l) = \mathcal{N}_P(\mathbf{0}, \sigma_{\beta_l}^2 \mathbf{I})$  with a large variance  $\sigma_{\beta_l}^2$  [168], and (iii)  $g$ -prior  $\pi(\beta_l) = \mathcal{N}_P(\mathbf{0}, g \cdot \omega_l^2 [\mathbf{X}^\top \mathbf{X}]^{-1})$  for some positive value  $g$  [169]. The suggested priors yield the conjugate update for **Step 4** within the Gibbs sampler.

Now, we discuss some priors for  $\pi(\beta_l)$  in the linear regression (33) under the *fat design setting* ( $N \ll P$ ). This setting can be applied to pharmacogenetics where one of the main interests is to find important genes that may influence pharmacokinetics or pharmacodynamics [170–172], where the number of genes  $P$  is allowed to be a few thousand, while number of patients  $N$  is confined to a few hundred. A fundamental assumption in this setting is *sparsity assumption* on the coefficients  $\beta_l$ . This means that many of the coefficients of  $\beta_l = (\beta_{l1}, \dots, \beta_{lP})^\top$  are (close to) zero. The true non-zero coefficients in the  $\beta_l$  are referred to as signal coefficients, while the remaining are called noise coefficients.

Statisticians have devised a number of penalized regression techniques for estimating  $\beta_l$  under the sparsity assumption [173]. From a Bayesian point of view, sparsity favoring mixture priors with separate control on the signal and noise coefficients have been proposed [174–177], which is called the ‘spike-and-slab priors’. Although these priors often lead to attractive theoretical properties [157,178], computational issues and considerations that many of the  $\beta_{lb}$ ’s ( $b = 1, \dots, P$ ) may be small but not exactly zero has led to a wide variety of ‘continuous shrinkage priors’ [179–183], which can be unified through a global-local scale mixture representation [184]. The following hierarchies describe the sparse favoring priors:

- **Spike-and-slab priors.** Each component of the coefficients  $\beta_l$  is assumed to be drawn from

$$\beta_{lb} | \tau_l \sim (1 - \tau_l) \cdot \delta_0(\beta_{lb}) + \tau \cdot f(\beta_{lb}), \quad (l = 1, \dots, K; b = 1, \dots, P),$$

where  $\tau_l = \mathbf{Pr}[\beta_{lb} \neq 0]$ . The function  $\delta_0(\beta_{lb})$  is the Dirac-delta function and  $f(\beta_{lb})$  is a density supported on  $\mathbb{R}$ , called the spike and slab densities, respectively. The spike density shrinks noise coefficients to the exact zero, while the slab density captures signal coefficients by allowing a positive mass on the tail region [157,185,186].

• **Continuous shrinkage priors.** Each component of the coefficients  $\beta_l$  is assumed to be drawn from

$$\beta_{lb} | \lambda_{lb}, \tau_l, \omega_l^2 \sim \mathcal{N}(0, \lambda_{lb}^2 \tau_l^2 \omega_l^2), \quad (l = 1, \dots, K; b = 1, \dots, P),$$

$$\lambda_{lb} \sim f(\lambda_{lb}), \tau_l \sim g(\tau_l), \omega_l^2 \sim h(\omega_l^2), \quad (l = 1, \dots, K; b = 1, \dots, P),$$

where  $f$ ,  $g$ , and  $h$  are priors for  $\lambda_{lb}$ ,  $\tau_l$ , and  $\omega_l^2$ , respectively, supported on  $(0, \infty)$ . Here,  $\lambda_{lb}$  and  $\tau_l$  are referred to as local-scale and global-scale parameters, respectively. The choices of  $f$  and  $g$  play a key role in controlling the effective sparsity and concentration of the prior and posterior distributions [184,187–191].

Roughly speaking, the roles of the  $\tau_l$  in both prior frameworks are similar in the sense that they control the degree of the sparsity [184]. Slab density and local-scale prior density are expected to put a sufficient mass on the tail regions of the densities to detect signal coefficients and produce a robust Bayes estimator for  $\beta_l$  [158,182]; heavy-tailed densities (for e.g., double generalized Pareto distribution, Cauchy distributions [192,193]) are preferably used by that reason. Refer to [194,195] for comprehensive surveys on Bayesian variable selection.

### 6.3. Priors for covariance matrix

Consider a nonlinear function  $f(t; \theta)$  indexed by a  $K$ -dimensional model parameter  $\theta = (\theta_1, \dots, \theta_l, \dots, \theta_K)^\top$  to describe an individual trajectory. A basic assumption is that all the components  $\theta_l$  are unrelated across  $l$  ( $l = 1, \dots, K$ ), which is referred to as *uncorrelated design setting*. In many practical problems, this setting is reasonably accepted since one of the fundamental assumptions on  $f$  is that each component  $\theta_l$  has its own role in modifying a functional shape of  $f$ . A central goal of researcher in using nonlinear mixed models is to examine these roles mathematically, endowed with interpretations by domain experts in terms of physiology, epidemiology, or pharmacology, etc. The basic model (2) – (4) that we illustrated so far is designed under this assumption; recall that the covariance matrix  $\Omega \in \mathbb{R}^{K \times K}$  on Stage 2 was assumed to be diagonal,  $\Omega = \text{diag}(\omega_1^2, \dots, \omega_l^2, \dots, \omega_K^2)$ . Under this uncorrelated design setting, possible options for priors for the scale components  $\omega_l^2$  (or  $\omega_l$ ) are (i) Jeffreys prior  $\pi(\omega_l^2) \propto 1/\omega_l^2$  [154]; (ii) inverse-gamma prior  $\pi(\omega_l^2) = \mathcal{IG}(a_{\omega_l^2}, b_{\omega_l^2})$  with shape  $a_{\omega_l^2} > 0$  and scale  $b_{\omega_l^2} > 0$ ; and (iii) half-Cauchy prior  $\pi(\omega_l) = \mathcal{C}^+(0, b_{\omega_l}) = \{2/(\pi b_{\omega_l})\} \cdot 1/\{1 + (\omega_l/b_{\omega_l})^2\}$  with scale  $b_{\omega_l} > 0$ . See discussion by [196] for the prior options implemented on *8-schools example*.

We discuss Bayesian inference about a population covariance matrix  $\Omega$  under a *correlated design setting*, where the off-diagonal entries of  $\Omega$  (34) are allowed to be nonzeros:

$$\Omega = \begin{bmatrix} \omega_{11} & \cdots & \omega_{1l} & \cdots & \omega_{1K} \\ \vdots & & \vdots & & \vdots \\ \omega_{l1} & \cdots & \omega_{ll} & \cdots & \omega_{lK} \\ \vdots & & \vdots & & \vdots \\ \omega_{K1} & \cdots & \omega_{Kl} & \cdots & \omega_{KK} \end{bmatrix} \in \mathbb{R}^{K \times K}. \quad (34)$$

In (34), the  $(l_1, l_2)$ -th entry is denoted by  $\Omega[l_1, l_2] = \omega_{l_1 l_2} = \text{cov}(\eta_{l_1 i}, \eta_{l_2 i})$  ( $l_1, l_2 = 1, \dots, K; i = 1, \dots, N$ ). With  $l_1 = l_2$ , we have  $\omega_{l_1 l_2} = \omega_{l_1}^2$  ( $l_1 = 1, \dots, K$ ).

Researchers often wish to work with the correlated design setting to examine whether any pair of model parameters,  $\theta_{l_1}$  and  $\theta_{l_2}$ , are physiologically (or epidemiology, pharmacology, financially, etc) associated or not. Taking the term structure modeling discussed in Subsection 2.3 as an example, it is a valid question whether the Nelson-Siegel parameters are correlated or not as they are all associated with the interest rate [64]. Statistically, having a well-designed covariance structure can also improve the model fitting and produce reliable estimators for the model parameters compared to uncorrelated designs. To make a

fully Bayesian inference about the  $\Omega$  (34), we need to specify an appropriate prior  $\pi(\Omega)$  which we will discuss shortly. After that, we operate the Gibbs sampler in Subsection 5.2, with some modifications, if necessary. For instance, to implement the parallel computation in *Step 1*, we recommend to sample from the joint density  $\pi(\theta^i | -)$  across  $i$ , rather than sampling from the individual  $\pi(\theta_{li} | -)$ . Especially, among the five steps of the Gibbs sampler, implementation of the *Step 5* needs special care in sampling from the full-conditional posterior density  $\Omega$ . This step can be highly complicated depending on the chosen prior.

A challenge in choosing a workable prior  $\pi(\Omega)$  is briefly mentioned. Researches regarding this subject are vast, growing, and deep. As similar to the obstacles encountered in classical covariance estimation [197–201], there are three major aspects, among many others, in the consideration of a thoughtful prior  $\pi(\Omega)$  to produce a reliable Bayes estimator of  $\Omega$ : (1) sample size  $N$ ; (2) the number of model parameters  $K$ ; and (3) positive-definiteness of  $\Omega$  [202]. The first two aspects are related to theoretical constraints. In general, it is known that the estimation of the covariance  $\Omega$  can be distorted unless the ratio  $K/N$  is sufficiently small (see, e.g., [203–205]). The third one is germane to the modeling consideration and computation strategies to estimate  $\Omega$ , typically resolved via principal component analysis, Cholesky decomposition, and Gaussian graphical models, etc [206,207].

For any prior  $\pi(\Omega)$ , the full conditional density of  $\Omega$  is analytically expressed as follows (see [208] for a similar derivation):

$$\begin{aligned} \pi(\Omega | \Theta, \sigma^2, \alpha, B, \mathbf{y}_{1:N}) &\propto \pi(\Theta | \alpha, B, \Omega) \cdot \pi(\Omega) \propto \left\{ \prod_{i=1}^N \mathcal{N}_K(\theta^i | \alpha + B\mathbf{x}_i, \Omega) \right\} \cdot \pi(\Omega) \quad (35) \\ &\propto (\det \Omega)^{-N/2} \exp \left( -\frac{1}{2} \text{tr} \left[ \Omega^{-1} \left\{ (N-1)\mathbf{G} + N(\bar{\phi} - \alpha)(\bar{\phi} - \alpha)^\top \right\} \right] \right) \cdot \pi(\Omega), \end{aligned}$$

where  $\phi^i = \theta^i - B\mathbf{x}_i$  ( $i = 1, \dots, N$ ),  $\bar{\phi} \in \mathbb{R}^K$ , and  $\mathbf{G} \in \mathbb{R}^{K \times K}$  are defined by

$$\begin{aligned} \bar{\phi} &= N^{-1} \sum_{i=1}^N \phi^i = N^{-1} \sum_{i=1}^N (\theta^i - B\mathbf{x}_i), \quad \mathbf{G} = (N-1)^{-1} \sum_{i=1}^N (\phi^i - \bar{\phi})(\phi^i - \bar{\phi})^\top; \\ \sum_{i=1}^N (\phi^i - \alpha)(\phi^i - \alpha)^\top &= (N-1)\mathbf{G} + N(\bar{\phi} - \alpha)(\bar{\phi} - \alpha)^\top. \end{aligned}$$

In (35), we used the vector-form (b) (i.e.,  $i$ -indexing) to express the prior for  $\Theta$ . Notations  $\det(\mathbf{A})$  and  $\text{tr}(\mathbf{A})$  denote the determinant and trace of a square matrix  $\mathbf{A}$ , respectively. Matrix  $\mathbf{G}$  is the ‘latent’ covariance matrix based on the model matrix  $\Theta \in \mathbb{R}^{K \times K}$  (5) and coefficient matrix  $\mathbf{G} \in \mathbb{R}^{K \times P}$ , whose form resembles sample covariance matrix assuming  $\phi^i$  are observed [205].

Traditionally used priors for the covariance matrix  $\Omega$  (34) are the Jeffreys prior and the conjugate inverse Wishart prior (see [209,210] for the reviews of the earlier works):

- **Jeffreys prior.** The common non-informative prior has been the Jeffreys improper prior

$$\pi(\Omega) \propto (\det \Omega)^{-(K+1)/2}.$$

This prior was originally derived from an invariance argument by [211] for the case  $K = 1, 2$ ; and it was considered for arbitrary  $K$  by [208,212,213] to develop Bayesian multivariate theory.

- **Inverse-Wishart prior.** The common informative prior is the inverse-Wishart prior [214]

$$\pi(\Omega) = \mathcal{IW}(\mathbf{V}, d) = \frac{(\det \mathbf{V})^{d/2}}{2^{dK/2} \Gamma_K(d/2)} (\det \Omega)^{-(d+K+1)/2} \exp \left( -\frac{1}{2} \text{tr} \left[ \Omega^{-1} \mathbf{V} \right] \right),$$

where  $\Omega$  and  $\mathbf{V}$  are  $K$ -by- $K$  positive definite matrices, and  $\Gamma_K(\cdot)$  is the multivariate gamma function [215].  $\mathbf{V}$  is a scale matrix, and  $d(> K - 1)$  is the number of degrees of freedom. Conventionally,  $d$  is chosen to be as small as possible to reflect vague prior knowledge. A univariate specialization ( $K = 1$ ) is the inverse-gamma distribution.

The success of Bayesian computation and MCMC in the late 1980s opened up the potential of using more flexible non-conjugate priors for covariance matrices [216–218]. Limitations of the traditional priors studied by many statisticians also motivated them to develop a new prior. For example, some of them argued that the Jeffreys prior may not be really non-informative, particularly in high dimensional setting [207,219], and inverse Wishart prior is very restrictive and lacks flexibility [220]. Among many new priors developed for particular applications [221], a combination of *separation strategy* developed by [207] and LKJ prior [222] has been successful, heavily used in a variety of industrial problems, and relevant software has been developed, including R package STAN [17,148].

We illustrate a central idea of using the separation strategy [207] to estimate the population covariance matrix  $\Omega \in \mathbb{R}^{K \times K}$ . First, we decompose the matrix  $\Omega$  (34) into two components,  $K$  standard deviations  $\omega_l = \sqrt{\omega_{ll}} = \sqrt{\Omega[l, l]}$  ( $l = 1, \dots, K$ ) and correlation matrix  $\mathbf{R} \in \mathbb{R}^{K \times K}$ :

$$\Omega = \text{diag}(\omega) \mathbf{R} \text{diag}(\omega) \in \mathbb{R}^{K \times K}, \quad (36)$$

where  $\omega = (\omega_1, \dots, \omega_K)^\top$  is the  $K$ -dimensional vector of standard deviations,  $\text{diag}(\omega)$  is the diagonal matrix with diagonal elements  $\omega$ , and  $\mathbf{R}$  is the  $K$ -by- $K$  correlation matrix. Second, we specify priors independently for  $\omega$  and  $\mathbf{R}$ , denoted as  $\pi(\omega)$  and  $\pi(\mathbf{R})$ , respectively, so that we have the joint prior for  $(\omega, \mathbf{R})$ ,  $\pi(\omega, \mathbf{R}) = \pi(\omega) \cdot \pi(\mathbf{R})$ . Following notation from [207], let  $\mathcal{R}^K$  denote the correlation matrix space. Then, the priors  $\pi(\omega)$  and  $\pi(\mathbf{R})$  are supported on  $(0, \infty)^K$  and  $\mathcal{R}^K$ , respectively. Finally, draw sample from each of the full conditional posterior distributions  $\pi(\omega|\mathbf{R}, -)$  and  $\pi(\mathbf{R}|\omega, -)$  at time in *Step 5* within the Gibbs sampler in Subsection 5.2. (This means that, we are not directly sampling the covariance matrix  $\Omega$  from  $\pi(\Omega| -)$  (35) as we would do when Jeffreys or inverse-Wishart prior were used for the prior.)

Standardly used prior options for the scale vector  $\omega = (\omega_1, \dots, \omega_l, \dots, \omega_K)^\top$  are (i)  $\log(\omega) \sim \mathcal{N}_K(\mathbf{a}_\omega, \mathbf{B}_\omega)$ , where  $\log(\omega) = (\log \omega_1, \dots, \log \omega_l, \dots, \log \omega_K)^\top$ , with hyperparameters, mean  $\mathbf{a}_\omega \in \mathbb{R}^K$  and covariance  $\mathbf{B}_\omega \in \mathbb{R}^{K \times K}$  which is often diagonal [207]; and (ii)  $\omega_l \sim \mathcal{C}^+(0, b_{\omega_l})$  with the scale hyperparameter  $b_{\omega_l} > 0$  [17]. As for the prior distribution for the correlation  $\mathbf{R} \in \mathbb{R}^{K \times K}$ , the LKJ prior proposed by [222] is popularly used:

- **LKJ prior.** LKJ prior is supported over the correlation matrix space  $\mathcal{R}^K$ , or equivalently over the set of  $K \times K$  Cholesky factors of real symmetric positive definite matrices

$$\pi(\mathbf{R}) = \left[ 2^{\sum_{q=1}^{Q-1} (2\gamma - 2 + Q - q)(Q - q)} \prod_{q=1}^{Q-1} \mathcal{B}\left(\gamma + \frac{Q - q - 1}{2}, \gamma + \frac{Q - q - 1}{2}\right)^{Q - q} \right] (\det \mathbf{R})^{\gamma - 1}, \quad (37)$$

with the shape parameter  $\gamma > 0$ . The function  $\mathcal{B}(\alpha, \beta)$  is the beta function. If  $\gamma = 1$ , the density is uniform over the space  $\mathcal{R}^K$ ; for  $\gamma > 1$ , the density increasingly concentrates mass around the identity matrix  $\mathbf{I} \in \mathbb{R}^{K \times K}$  (i.e., favoring less correlation); for  $\gamma < 1$ , the density increasingly concentrates mass in the other direction, and has a trough at the identity matrix (i.e., favoring more correlation).

Note that the normalizing constant of the LKJ prior (37) is constant with respect to  $\gamma$ , therefore, we have  $\pi(\mathbf{R}) \propto (\det \mathbf{R})^{\gamma - 1}$ , with the shape hyperparameter  $\gamma > 0$ . The behavior of LKJ prior with  $\gamma = 1$  (i.e.,  $\pi(\mathbf{R}) \propto 1$ ) was studied by [207], where the author found that as  $K$  increases the marginal correlations tend to concentrate around zeros (see Figure 1 from [207]), hence, model matrix  $\Theta$  (5) are more likely to be treated as in the uncorrelated design setting.



As for the hyperparameter specification for the LKJ prior, Stan Development Team [17] recommends to use  $\gamma \geq 1$ . This suggestion is also well-aligned with the original intention of using the separation strategy to make a variance-correlation structure by [207] in that: (1) the authors intend to choose a diffuse prior for  $\pi(\mathbf{R})$  to reflect weak knowledge about the correlation  $\mathbf{R}$ , while (2) prior knowledge, possibly informative, shall be put on the scale parameters by specifying  $\pi(\omega)$ , as most statisticians are normally trained to do. Computational algorithm and theory concerning the LKJ prior can be found in [222–224].

## 7. Model selection

### 7.1. Setting

The recent development of MCMC methods has made it possible to fit enormously large classes of models with the aim of exploring real world complexities of data [225]. This ability naturally led us to wish to compare several candidate models that vary substantially in the model complexities and choose the best model out of them. For example, authors [31] tried to compare four different rate decline curves to fit the production data from the 360 wells in Figure 2. Indeed, upstream petroleum engineers endeavor to find a rate decline curve describing the production trajectories as accurately as possible so that EUR can be accurately estimated. Another application of model selection can be found in PK analysis. Taking the theophylline data in Figure 1 as an example, PK modelers may debate whether they need to use a two or three-compartment model with a nonlinear clearance to describe the PK exposure, or just one-compartment model with a linear clearance is sufficient.

In the current section, our primary focus is to illustrate a Bayesian approach to compare multiple Bayesian nonlinear mixed effects models explaining the data  $\{(\mathbf{y}_i, \mathbf{t}_i, \mathbf{x}_i)\}_{i=1}^N$  introduced in Subsection 2.5. To that end, we want to lay out the set-up that underlies our model selection procedure. Consider Stage 1 and 2 of the basic model (i.e., the hierarchy (2) – (3)), endowed with a joint prior  $\pi(\sigma^2, \alpha, \mathbf{B}, \Omega)$ . Very importantly, we do *not* assume the independent prior assumption as we did in Stage 3 (4): any prior assumption on the parameters  $(\sigma^2, \alpha, \mathbf{B}, \Omega)$  works fine in our framework for the model comparison. Therefore, the basic model (2) – (4) is considered as a subclass of candidate models that we want to compare. In our framework, we can also consider the correlated design setting discussed in Subsection 6.3, where the covariance matrix  $\Omega$  (34) is allowed to be any positive-definite matrix (i.e., does not need to be a diagonal matrix), as one of candidates. Eventually, in our framework, modelers have freedom to choose (i) the nonlinear function  $f$  to describe the temporal profile  $\mathbf{y}_i$  and (ii) prior distribution  $\pi(\sigma^2, \alpha, \mathbf{B}, \Omega)$ .

Assume that a researcher wants to consider  $H$  functions, denoted as  $f(t; \theta_{[\mathcal{M}_1]}), f(t; \theta_{[\mathcal{M}_2]}), \dots, f(t; \theta_{[\mathcal{M}_H]})$ , as a possible option for the use of  $f$  in Stage 1. Here, the subscript ' $[\mathcal{M}_h]$ ' on the model parameter  $\theta$  is noted to indicate for the model index ( $h = 1, \dots, H$ ). Obviously, the candidate functions  $f(t; \theta_{[\mathcal{M}_h]})(h = 1, \dots, H)$  can have different functional forms dictated by different dimensions for the model parameters  $\theta_{[\mathcal{M}_h]} \in \mathbb{R}^{K_{[\mathcal{M}_h]}}$ . This will consecutively change the dimensions of the parameter blocks,  $\alpha_{[\mathcal{M}_h]} \in \mathbb{R}^{K_{[\mathcal{M}_h]}}$ ,  $\mathbf{B}_{[\mathcal{M}_h]} \in \mathbb{R}^{K_{[\mathcal{M}_h]} \times P}$ , and  $\Omega_{[\mathcal{M}_h]} \in \mathbb{R}^{K_{[\mathcal{M}_h]} \times K_{[\mathcal{M}_h]}}$ , accordingly, yet the support of the  $\sigma^2$  remains the same with  $(0, \infty)$  because we still consider the additive error model. After that, she now has the freedom to choose a prior  $\pi(\sigma^2, \alpha_{[\mathcal{M}_h]}, \mathbf{B}_{[\mathcal{M}_h]}, \Omega_{[\mathcal{M}_h]})(h = 1, \dots, H)$ . There are infinitely many choices for the prior, and one can use priors discussed in Section 6.

Now, the fundamental question naturally arising at this point is "what is the best model among the  $H$  candidate models?" To illustrate, we write the situation more technically. With above specifications, Stage 1 of each of the  $H$  candidate models is given as

$$\begin{aligned}\mathcal{M}_1: & y_{ij} = f(t_{ij}; (\boldsymbol{\theta}_{[\mathcal{M}_1]})^i) + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma^2), \quad (i = 1, \dots, N; j = 1, \dots, M_i); \\ & \vdots \\ \mathcal{M}_h: & y_{ij} = f(t_{ij}; (\boldsymbol{\theta}_{[\mathcal{M}_h]})^i) + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma^2), \quad (i = 1, \dots, N; j = 1, \dots, M_i); \\ & \vdots \\ \mathcal{M}_H: & y_{ij} = f(t_{ij}; (\boldsymbol{\theta}_{[\mathcal{M}_H]})^i) + \epsilon_{ij}, \quad \epsilon_{ij} \sim \mathcal{N}(0, \sigma^2), \quad (i = 1, \dots, N; j = 1, \dots, M_i),\end{aligned}$$

where each of the models has the corresponding model matrix

$$\begin{aligned}\mathcal{M}_1: & \boldsymbol{\Theta}_{[\mathcal{M}_1]} = \begin{bmatrix} (\boldsymbol{\theta}_{[\mathcal{M}_1]})^1 & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_1]})^i & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_1]})^N \end{bmatrix} \in \mathbb{R}^{K_{[\mathcal{M}_1]} \times N}; \\ & \vdots \\ \mathcal{M}_h: & \boldsymbol{\Theta}_{[\mathcal{M}_h]} = \begin{bmatrix} (\boldsymbol{\theta}_{[\mathcal{M}_h]})^1 & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_h]})^i & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_h]})^N \end{bmatrix} \in \mathbb{R}^{K_{[\mathcal{M}_h]} \times N}; \\ & \vdots \\ \mathcal{M}_H: & \boldsymbol{\Theta}_{[\mathcal{M}_H]} = \begin{bmatrix} (\boldsymbol{\theta}_{[\mathcal{M}_H]})^1 & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_H]})^i & \dots & (\boldsymbol{\theta}_{[\mathcal{M}_H]})^N \end{bmatrix} \in \mathbb{R}^{K_{[\mathcal{M}_H]} \times N},\end{aligned}$$

obtained by stacking individual-based vector horizontally as we did to obtain  $\boldsymbol{\Theta}$  (5). Again, the number of rows of the matrix  $\boldsymbol{\Theta}_{[\mathcal{M}_h]}$ , that is,  $K_{[\mathcal{M}_h]}$ , depends on the choice of the function  $f$ . Stage 2 of each of the  $H$  models will then be

$$\begin{aligned}\mathcal{M}_1: & (\boldsymbol{\theta}_{[\mathcal{M}_1]})^i = \boldsymbol{\alpha}_{[\mathcal{M}_1]} + \mathbf{B}_{[\mathcal{M}_1]}\mathbf{x}_i + (\boldsymbol{\eta}_{[\mathcal{M}_1]})^i, \quad (\boldsymbol{\eta}_{[\mathcal{M}_1]})^i \sim \mathcal{N}_K(\mathbf{0}, \boldsymbol{\Omega}_{[\mathcal{M}_1]}), \quad (i = 1, \dots, N); \\ & \vdots \\ \mathcal{M}_h: & (\boldsymbol{\theta}_{[\mathcal{M}_h]})^i = \boldsymbol{\alpha}_{[\mathcal{M}_h]} + \mathbf{B}_{[\mathcal{M}_h]}\mathbf{x}_i + (\boldsymbol{\eta}_{[\mathcal{M}_h]})^i, \quad (\boldsymbol{\eta}_{[\mathcal{M}_h]})^i \sim \mathcal{N}_K(\mathbf{0}, \boldsymbol{\Omega}_{[\mathcal{M}_h]}), \quad (i = 1, \dots, N); \\ & \vdots \\ \mathcal{M}_H: & (\boldsymbol{\theta}_{[\mathcal{M}_H]})^i = \boldsymbol{\alpha}_{[\mathcal{M}_H]} + \mathbf{B}_{[\mathcal{M}_H]}\mathbf{x}_i + (\boldsymbol{\eta}_{[\mathcal{M}_H]})^i, \quad (\boldsymbol{\eta}_{[\mathcal{M}_H]})^i \sim \mathcal{N}_K(\mathbf{0}, \boldsymbol{\Omega}_{[\mathcal{M}_H]}), \quad (i = 1, \dots, N).\end{aligned}$$

Finally, Stage 3 of each of the  $H$  models is comprised of the prior:

$$\begin{aligned}\mathcal{M}_1: & (\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_1]}, \mathbf{B}_{[\mathcal{M}_1]}, \boldsymbol{\Omega}_{[\mathcal{M}_1]}) \sim \pi(\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_1]}, \mathbf{B}_{[\mathcal{M}_1]}, \boldsymbol{\Omega}_{[\mathcal{M}_1]}); \\ & \vdots \\ \mathcal{M}_h: & (\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_h]}, \mathbf{B}_{[\mathcal{M}_h]}, \boldsymbol{\Omega}_{[\mathcal{M}_h]}) \sim \pi(\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_h]}, \mathbf{B}_{[\mathcal{M}_h]}, \boldsymbol{\Omega}_{[\mathcal{M}_h]}); \\ & \vdots \\ \mathcal{M}_H: & (\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_H]}, \mathbf{B}_{[\mathcal{M}_H]}, \boldsymbol{\Omega}_{[\mathcal{M}_H]}) \sim \pi(\sigma^2, \boldsymbol{\alpha}_{[\mathcal{M}_H]}, \mathbf{B}_{[\mathcal{M}_H]}, \boldsymbol{\Omega}_{[\mathcal{M}_H]}).\end{aligned}$$

We describe three model comparison criteria that are popularly used in the literature: deviance information criterion (DIC) [226,227], widely applicable information criterion (WAIC) [228], and posterior predictive loss criterion (PPLC) [229]. As in frequentist information criteria [230–232], formulation of the DIC, WAIC, and PPLC also takes the two terms into a consideration: *goodness-of-fit* and *penalty for model complexity*. Because increasing (or decreasing) model complexity is accompanied by the risk of over-fitting (or under-fitting), models should be compared by trading-off these two terms. Particularly, as we are currently discussing about a Bayesian hierarchical model, these criteria are obviously depending on what part of the model specification is considered to be part of the likelihood, and what

924  
925  
926  
927  
928  
929  
930  
931  
932

is not. [226] refer to this as the *focus issue*. For example, in the general form of a Bayesian hierarchical model consisting of a top-level likelihood  $p(\mathbf{y}|\Psi)$  for data  $\mathbf{y}$ , a prior model  $\pi(\Psi|\eta)$ , and a hyperprior  $\pi(\eta)$ , one might choose as the likelihood either the conditional density  $p(\mathbf{y}|\Psi)$ , or the marginal density  $p(\mathbf{y}|\eta) = \int p(\mathbf{y}|\Psi)\pi(\Psi|\eta)d\Psi$ . Based on [226], the former situation is referred to as “focus on  $\Psi$ ”, while the latter situation is referred to as “focus on  $\eta$ ”, respectively.

In our case, we shall “focus on parameters  $(\Theta_{[\mathcal{M}_h]}, \sigma^2)$ ” ( $h = 1, \dots, H$ ) used in the conditional density in Stage 1. For notational simplicity, we denote  $\Psi_{[\mathcal{M}_h]} = (\Theta_{[\mathcal{M}_h]}, \sigma^2)$ . Then, likelihood of each of the  $H$  models based on  $N$  observations  $\{(\mathbf{y}_i, \mathbf{t}_i, \mathbf{x}_i)\}_{i=1}^N$  is

$$\begin{aligned}\mathcal{M}_1: \quad \mathcal{L}(\Psi_{[\mathcal{M}_1]}|\mathbf{y}_{1:N}) &= \prod_{i=1}^N \mathcal{L}((\Psi_{[\mathcal{M}_1]})^i|\mathbf{y}_i) = \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, (\theta_{[\mathcal{M}_1]})^i), \sigma^2\mathbf{I}); \\ &\vdots \\ \mathcal{M}_h: \quad \mathcal{L}(\Psi_{[\mathcal{M}_h]}|\mathbf{y}_{1:N}) &= \prod_{i=1}^N \mathcal{L}((\Psi_{[\mathcal{M}_h]})^i|\mathbf{y}_i) = \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, (\theta_{[\mathcal{M}_h]})^i), \sigma^2\mathbf{I}); \\ &\vdots \\ \mathcal{M}_H: \quad \mathcal{L}(\Psi_{[\mathcal{M}_H]}|\mathbf{y}_{1:N}) &= \prod_{i=1}^N \mathcal{L}((\Psi_{[\mathcal{M}_H]})^i|\mathbf{y}_i) = \prod_{i=1}^N \mathcal{N}_{M_i}(\mathbf{y}_i|f_i(\mathbf{t}_i, (\theta_{[\mathcal{M}_H]})^i), \sigma^2\mathbf{I}),\end{aligned}$$

where  $\mathcal{L}((\Psi_{[\mathcal{M}_h]})^i|\mathbf{y}_i)$  is the likelihood (i.e., data distribution) based on an individual data with parameter  $(\Psi_{[\mathcal{M}_h]})^i = ((\theta_{[\mathcal{M}_h]})^i, \sigma^2)$  ( $i = 1, \dots, N$ ). One caveat of the practical calculation to obtain the three criteria is that, we mainly need the posterior samples of  $\Psi_{[\mathcal{M}_h]}$  drawn from the joint posterior density  $\pi(\Psi_{[\mathcal{M}_h]}, \alpha_{[\mathcal{M}_h]}, \mathbf{B}_{[\mathcal{M}_h]}, \Omega_{[\mathcal{M}_h]}|\mathbf{y}_{1:N})$  and the explicit form the likelihood function  $\mathcal{L}(\Psi_{[\mathcal{M}_h]}|\mathbf{y}_{1:N})$  from each of the models  $\mathcal{M}_h$  ( $h = 1, \dots, H$ ) due to our assumption of the focus.

In the next subsections, we provide some brief summaries of the criteria and then adapt them to our context. In what follows, to simplify the notation, we suppress the arguments  $[\mathcal{M}_h]$  in the parameters. For a detailed explanation of the criteria, refer to [233,234].

## 7.2. Deviance information criterion

[227] suggested DIC, a generalized version of Akaike information criterion [230] for a Bayesian hierarchical model, given by

$$\text{DIC} = D(\bar{\Psi}) + 2 \cdot p_D. \quad (38)$$

In (38), the function  $D(\Psi) = -2\log \mathcal{L}(\Psi|\mathbf{y}_{1:N})$  is referred to as deviance. Deviance is a goodness-of-fit statistics whose lower value indicates a better fitting [235]. Goodness-of-fit term of DIC (i.e.,  $D(\bar{\Psi})$ ) is the value of deviance evaluated at the posterior mean of  $\Psi$ , denoted by  $\bar{\Psi} = \mathbb{E}[\Psi|\mathbf{y}_{1:N}] = \int \Psi \pi(\Psi|\mathbf{y}_{1:N})d\Psi$ , where  $\pi(\Psi|\mathbf{y}_{1:N})$  represents the posterior distribution of  $\Psi = (\Theta, \sigma^2)$ . Effective number of parameters (i.e, penalty term for model complexity) of DIC in (38) is obtained by  $p_D = \text{Var}[D(\Psi)|\mathbf{y}_{1:N}]/2 = 2 \cdot \text{Var}[\log \mathcal{L}(\Psi|\mathbf{y}_{1:N})|\mathbf{y}_{1:N}]$ . A model with a smaller value for DIC indicates a better predictive performance among considered models.

Some intuition behind having two competing additive terms in (38) is as follows. Typically, complex models get rewards in terms of the deviance than simple models: therefore, over-fitted models normally have a preference over under-fitted models when only the deviance is considered in model comparison, which is undesirable. By adding a penalty term for the model complexity to the deviance term, we hope that the resulting criterion produces a reasonable value based on fair comparison regardless of model complexity.

Roughly speaking, this principle (i.e., a trade-off between the goodness-of-fit and penalty terms) is commonly manifested in the three criteria, DIC, WAIC, and PPLC.

Going back to our examples, we can obtain the DIC corresponding to each of the  $H$  candidate models  $\mathcal{M}_h$  ( $h = 1, \dots, H$ ):

$$D(\Psi) = -2 \sum_{i=1}^N \log \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}),$$

where only mean function  $f_i(\mathbf{t}_i, \theta^i)$  differs across  $H$  models  $\mathcal{M}_h$ ,  $h = 1, \dots, H$ . In practice, posterior mean  $\bar{\Psi}$  and effective number of parameters  $p_D$  are not expressed in closed-forms, hence, the DIC (38) is stochastically approximated though MCMC techniques [236].

### 7.3. Widely applicable information criterion

[228] introduced WAIC which is regarded as a fully Bayesian version of the DIC (38) in the sense that a goodness-of-fit term exploits the entire posterior distribution. Note that the goodness-of-fit term of the DIC (38) is obtained by plugging the posterior mean  $\bar{\Psi}$  into the deviance  $D(\Psi)$ , which lacks a fully Bayesian sense. It is known that WAIC is asymptotically equivalent to Bayesian cross-validation [237], and also applicable to singular models.

WAIC is defined by

$$\text{WAIC} = -2 \cdot \text{LPPD} + 2 \cdot p_W, \quad (39)$$

where the goodness-of-fit term is called the log posterior predictive density (LPPD), which is defined as  $\text{LPPD} = \sum_{i=1}^N \log \mathbb{E}[\mathcal{L}(\Psi^i | \mathbf{y}_i) | \mathbf{y}_{1:N}]$ , and the effective number of parameter in the penalty term is defined by  $p_W = \sum_{i=1}^N \text{Var}[\log \mathcal{L}(\Psi^i | \mathbf{y}_i) | \mathbf{y}_{1:N}]$ .

In practice, as similar to DIC (38), WAIC (39) is obtained by stochastic approximations. Given posterior samples  $\{(\Psi)^{(s)}\}_{s=1}^S \sim \pi(\Psi | \mathbf{y}_{1:N})$ , the LPPD and  $p_W$  terms may be approximated by

$$\widehat{\text{LPPD}} = \sum_{i=1}^N \log \left( \frac{1}{S} \sum_{s=1}^S \mathcal{L}((\Psi^i)^{(s)} | \mathbf{y}_i) \right), \quad (40)$$

$$\widehat{p}_W = \sum_{i=1}^N \left\{ \frac{1}{S-1} \sum_{s=1}^S \left( \log \mathcal{L}((\Psi^i)^{(s)} | \mathbf{y}_i) - \frac{1}{S} \sum_{s=1}^S \log \mathcal{L}((\Psi^i)^{(s)} | \mathbf{y}_i) \right)^2 \right\}. \quad (41)$$

Returning to our examples, we can approximate the value of WAIC corresponding to each of the  $H$  models as follows. First, replace  $\mathcal{L}(\Psi^i | \mathbf{y}_i)$  in (40) and (41) with the individual-based data distribution  $p(\mathbf{y}_i | \Psi^i) = \mathcal{N}_{M_i}(\mathbf{y}_i | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I})$  ( $i = 1, \dots, N$ ), where only the mean function  $f_i(\mathbf{t}_i, \theta^i)$  differs across the  $H$  candidate models, and second, approximate  $\widehat{\text{LPPD}}$  and  $\widehat{p}_W$  by using a MCMC method, and finally, obtain an approximation of WAIC (39) corresponding to each model.

### 7.4. Posterior predictive loss criterion

[229] introduced PPLC as an alternative to DIC (38) or WAIC (39). A notable feature of PPLC different from DIC and WAIC is its use of replicated observations, denoted by  $\mathbf{y}_i^{\text{rep}} = (y_{i1}^{\text{rep}}, y_{i2}^{\text{rep}}, \dots, y_{iM_i}^{\text{rep}})^{\top} \in \mathbb{R}^{M_i}$ , corresponding to the actual observations  $\mathbf{y}_i = (y_{i1}, y_{i2}, \dots, y_{iM_i})^{\top} \in \mathbb{R}^{M_i}$ , for each  $i = 1, \dots, N$ . Here, the replicate  $\mathbf{y}_i^{\text{rep}}$  for the subject  $i$  is drawn from its posterior predictive density

$$f(\mathbf{y}_i^{\text{rep}} | \mathbf{y}_{1:N}) = \int p(\mathbf{y}_i^{\text{rep}} | \Psi^i) \cdot \pi(\Psi | \mathbf{y}_{1:N}) d\Psi, \quad (i = 1, \dots, N), \quad (42)$$

where  $p(\mathbf{y}_i^{\text{rep}} | \Psi^i)$  is the data density for the  $i$ -th subject and  $\pi(\Psi | \mathbf{y}_{1:N})$  is posterior distribution. The idea of using replicates  $\{\mathbf{y}_i^{\text{rep}}\}_{i=1}^N$  for a criticism of the model in light of the observed data  $\{\mathbf{y}_i\}_{i=1}^N$  is also purported by [238].

A general rule of the PPLC is principled on a balanced loss function [239]. Given any loss function  $l(\cdot)$  and a positive real number  $k$ , a balanced loss function is defined by

$$l(\mathbf{y}_i^{rep}, \mathbf{a}_i; \mathbf{y}_{1:N}) = l(\mathbf{y}_i^{rep}, \mathbf{a}_i) + k \cdot l(\mathbf{y}_i, \mathbf{a}_i), \quad k > 0, i = 1, \dots, N, \quad (43)$$

where  $\mathbf{a}_i$  is a non-stochastic action vector,  $k$  is a weight, and  $\mathbf{y}_i^{rep}$  is a replicate for its observed counterpart  $\mathbf{y}_i$ . Conceptually, the role of action vector  $\mathbf{a}_i$  is to accommodate both  $\mathbf{y}_i$ , and what we predict for  $\mathbf{y}_i^{rep}$ . Note that the loss function on the left-hand side of (43) penalizes actions  $\mathbf{a}_i$  both for departure from the corresponding observed value (fit) as well as for departure from what we expect the replicate to be (smoothness) [233]. A generic version of PPLC is defined by  $D_k = \sum_{i=1}^N \min_{\mathbf{a}_i} \mathbb{E}[l(\mathbf{y}_i^{rep}, \mathbf{a}_i; \mathbf{y}_{1:N}) | \mathbf{y}_{1:N}]$ , where the expectation  $\mathbb{E}[\cdot | \mathbf{y}_{1:N}]$  is taken with respect to the predictive density  $f(\mathbf{y}_i^{rep} | \mathbf{y}_{1:N})$  (42) for some specified  $k \geq 0$ . Note that the resulting value  $D_k$  does not depend on the action vector  $\mathbf{a}_i$  and replicates  $\{\mathbf{y}_i^{rep}\}_{i=1}^N$  as they are marginalized out by the minimization and expectation, respectively, but is dependent on the constant  $k > 0$ .

By choosing the quadratic loss  $l(\mathbf{y}, \mathbf{a}) = \|\mathbf{y} - \mathbf{a}\|_2^2$  in (43), the generic PPLC  $D_k$  may be simplified as

$$D_k = \frac{k}{k+1}G + P, \quad k \geq 0, \quad (44)$$

where  $G = \sum_{i=1}^N \|\mathbf{v}_i - \mathbf{y}_i\|_2^2$  and  $P = \sum_{i=1}^N \zeta_i^2$  represent the goodness-of-fit and penalty terms, respectively, with  $\mathbf{v}_i = \mathbb{E}[\mathbf{y}_i^{rep} | \mathbf{y}_{1:N}]$  and  $\zeta_i^2 = \mathbb{E}[\|\mathbf{y}_i^{rep} - \mathbf{v}_i\|_2^2 | \mathbf{y}_{1:N}]$ ,  $i = 1, \dots, N$ . Eventually, a model with a smaller value for the  $D_k$  (44) is preferable. It is known that ordering of models is insensitive to the particular choice of  $k$  [233].

Finally, we adapt the PPLC (44) to our examples. Due to the definition of notation  $\Psi^i = (\theta^i, \sigma^2)$  ( $i = 1, \dots, N$ ), the posterior predictive distribution of  $\mathbf{y}_i^{rep}$  (42) can be detailed as follows

$$f(\mathbf{y}_i^{rep} | \mathbf{y}_{1:N}) = \int \mathcal{N}_{M_i}(\mathbf{y}_i^{rep} | f_i(\mathbf{t}_i, \theta^i), \sigma^2 \mathbf{I}) \cdot \pi(\Psi^i | \mathbf{y}_{1:N}) d\Psi^i, \quad (i = 1, \dots, N).$$

To approximate  $D_k$  (44) for each model, first, choose a number  $k$ , saying  $k = 1$ , and second, approximate  $\mathbf{v}_i$  and  $\zeta_i^2$  through replicates  $\mathbf{y}_i^{rep}$  drawn from the predictive density  $f(\mathbf{y}_i^{rep} | \mathbf{y}_{1:N})$  (42) for each  $i = 1, \dots, N$ , and finally, complete the  $G$  and  $P$  to get an approximation to the  $D_k$  (44).

## 8. Extensions and recent developments

### 8.1. Residual error models

In the basic version of the Bayesian nonlinear mixed effects model (2) – (4), we assume that residual errors in the individual-level model are additive to the mean function  $f$  across all subjects and times. Under this assumption, the (conditional) variance of the  $i$ -th subject's trajectory  $\mathbb{V}[y_{ij} | \theta^i]$  is constant with  $\sigma^2$  over time  $t_{ij}$  ( $j = 1, \dots, M_i$ ). The additive error model is the most standard assumption used in a variety of problems arising from many industrial and academic researches [28,30,31,240,241]. However, when there exists some systematic temporal trend in the volatility of individual trajectories, for instance, the variance  $\mathbb{V}[y_{ij} | \theta^i]$  seems to decrease over times  $t_{ij}$  as shown in Figure 1 and 2, the additive residual assumption may not be adequate to fully account for the reality of the data.

List in Table 1 are popularly used residual error models that can be used in Stage 1 (2). Some of them are deployed as options for user to choose in industrial software such as MONOLIX [115] and NONMEM [41,102], and open source R package such as NLMIXR [18]. If we assume  $\epsilon_{ij} = \varepsilon_{ij} = 0$ , then all the error models leads to the same deterministic equation  $y_{ij} = f(t_{ij}; \theta^i)$ . That being said, if the variances of the residuals  $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$  and  $\varepsilon_{ij} \sim \mathcal{N}(0, \zeta^2)$ , that is,  $\sigma^2$  and  $\zeta^2$ , are quite small, then the inference outcome based on each of the error models will be similar each other.

**Table 1.** List of residual error models that can be used in Stage 1 (2).

Residual error type	Individual-level model	Mean $\mathbb{E}[y_{ij} \theta^i]$	Variance $\mathbb{V}[y_{ij} \theta^i]$
Additive	$y_{ij} = f(t_{ij}; \theta^i) + \epsilon_{ij}$	$f(t_{ij}; \theta^i)$	$\sigma^2$
Proportional	$y_{ij} = f(t_{ij}; \theta^i) \cdot (1 + \epsilon_{ij})$	$f(t_{ij}; \theta^i)$	$\{f(t_{ij}; \theta^i)\}^2 \cdot \sigma^2$
Exponential	$y_{ij} = f(t_{ij}; \theta^i) \cdot \exp(\epsilon_{ij})$	$f(t_{ij}; \theta^i) \cdot \exp(\sigma^2/2)$	$\{f(t_{ij}; \theta^i)\}^2 \cdot (\exp(\sigma^2) - 1) \cdot \exp(\sigma^2)$
Additive and proportional	$y_{ij} = f(t_{ij}; \theta^i) \cdot (1 + \epsilon_{ij}) + \epsilon_{ij}$	$f(t_{ij}; \theta^i)$	$\{f(t_{ij}; \theta^i)\}^2 \cdot \sigma^2 + \zeta^2$
Additive and exponential	$y_{ij} = f(t_{ij}; \theta^i) \cdot \exp(\epsilon_{ij}) + \epsilon_{ij}$	$f(t_{ij}; \theta^i) \cdot \exp(\sigma^2/2)$	$\{f(t_{ij}; \theta^i)\}^2 \cdot (\exp(\sigma^2) - 1) \cdot \exp(\sigma^2) + \zeta^2$

<sup>1</sup> Random errors are assumed to be distributed according to  $\epsilon_{ij} \sim \mathcal{N}(0, \sigma^2)$  and  $\epsilon_{ij} \sim \mathcal{N}(0, \zeta^2)$  ( $i = 1, \dots, N; j = 1, \dots, M_i$ ), with independence between  $\epsilon_{ij}$  and  $\epsilon_{ij}$ .

The exponential error model (i.e.,  $y_{ij} = f(t_{ij}; \theta^i) \cdot \exp(\epsilon_{ij})$ ) is a routine option, which is used when the ranges of the response  $y_{ij}$  and mean function  $f$  are positive real numbers, while we want to systematically describe the temporal volatility. In the practical implementation, we take the natural logarithm on the both sides of equation of the error model so that the model is converted to an additive error model in log-scale (i.e.,  $\log y_{ij} = \log(f(t_{ij}; \theta^i)) + \epsilon_{ij}$ ). That way, relevant Bayesian computation becomes much straightforward. [31] analyzed the shale oil data shown in Figure 2 in this formulation.

## 8.2. Bayesian Nonparametric methods

Recently, the use of the Bayesian nonparametric (BNP) statistical models has received increasing attention in the statistical literature because they allow modelers to gain model flexibility and robustness compared to its parametric counterpart [242,243]. BNP methods can be applied to the formulation of the basic model (2) – (4), when the parametric specification for the error distributions is too restrictive to achieve certain purpose of the analysis, or inference results lead to poor performance due to the inappropriate parametric form. Typically, BNP methods are applied to the population-level model, by extending or relaxing the parametric assumption on the random errors  $\eta_{li}$ , while retaining the individual-level model as fully parametric [244]. A Gaussian process prior [245,246] or a Dirichlet process prior [247–249] is popularly used for such extension and relaxation. Mathematical concepts of the processes are explained in [246,247].

To illustrate some motivation behind the application of BNP methods, we take the shale oil production data in Figure 2 researched by [31] as an example. Their goal was to predict EUR at a new location before the actual drilling takes place. Because the geological location is not stochastically incorporated into the basic model (2) – (4), authors extended the linear regression in Stage 2 into a spatial linear regression as follows

$$\theta_{li} = \theta_l(\mathbf{s}_i) = \alpha_l + \mathbf{x}_i^\top \boldsymbol{\beta}_l + v_l(\mathbf{s}_i) + \eta_l(\mathbf{s}_i), \quad (i = 1, \dots, N; l = 1, \dots, K),$$

where  $\eta_l(\cdot) \sim \mathcal{GP}(0, \omega_l^2 I(\cdot, \cdot))$  represents a Gaussian white noise process with the indicator function  $I(\cdot, \cdot)$  with variance  $\omega_l^2$ . The stochastic process  $v_l(\cdot) \sim \mathcal{GP}(0, \mathcal{K}(\cdot, \cdot))$  is the newly introduced Gaussian process with a radial basis function kernel  $\mathcal{K}(\mathbf{s}_{i_1}, \mathbf{s}_{i_2}) = \gamma_l^2 \exp[-\|\mathbf{s}_{i_1} - \mathbf{s}_{i_2}\|_2^2 / \{2\rho_l^2\}]$  with variance  $\gamma_l^2$  and range parameter  $\rho_l^2$ , and  $\mathbf{s}_i$  represents the (longitude, latitude) of the  $i$ -th shale oil well. The existence of  $v_l(\cdot)$  enables spatial prediction of EUR at a new location, taking an advantage of the geological proximity information, which is called the latent kriging technique.

Another motivation to use BNP methods is the situation when there exists multimodality in the distribution  $\mathcal{P}$  of model parameter vector  $\{\theta^i\}_{i=1}^N \sim \mathcal{P}$ . Note that, in the basic model (2) – (4), the distribution  $\mathcal{P}$  is assumed to be a single multivariate normal distribution  $\mathcal{N}_K(\boldsymbol{\alpha} + \mathbf{B}\mathbf{x}_i, \boldsymbol{\Omega})$ . In the multimodality case, the population may consist of disparate subpopulations, and the single multivariate normal distribution of the basic model can produce a poor model performance due to the lack of flexibility. A natural generalization to accommodate such multimodality is an extension to a finite mixture of multivariate normal distributions [250], or furthermore, to a countably infinite number of mixtures of



multivariate normal distributions [251]. Particularly, in the latter case, if a Dirichlet process prior [247] is placed on the mixture components, then the resulting infinite mixture models are generally called Dirichlet process mixture (DPM) model [252,253]. DPM model is one of the most studied topics in BNP methods in recent years [243]. See [254] for a survey of the posterior computations of using DPM models.

A number of authors have studied DPM models under the basic model (2) – (4) or similar forms with their own specifications [24,241,244,255]. For example, [244] placed a DPM model only for the model parameter vector  $\theta^i$  ( $i = 1, \dots, N$ ), while the covariates  $x_i$  are incorporated into the base measure of Dirichlet process. In contrast, [241,255] used a DPM model jointly for the model parameter vector and covariates,  $(\theta^i, x_i)$  ( $i = 1, \dots, N$ ), to induce a nonparametric regression function  $\mathbb{E}[\theta^i | x_i]$ . [24] used a Dirichlet process prior only for a certain component  $\theta_{li}$  ( $i = 1, \dots, N$ ) corresponding to a block indicator in an analysis-of-variance setup. Refer to the [256] for a review and references therein for more specifications.

8.3. Software development

Recent years have seen the great success of Bayesian nonlinear mixed effects models, or more generally, Bayesian hierarchical models (BHM), in a variety of disciplines such as biology, medical research, physics, social, and educational sciences [30,31,86,257,258]. This was partly due to the widespread introduction of software package that enabled applied researchers to answer substantive research questions through applications of BHM [17,102,259–262]. Most of the Bayesian software such as JAGS [259], BUGS [260], and STAN [17] are designed to require a reasonable understanding of the MCMC sampling scheme. From the perspective of implementation, spirits of most Bayesian software are similar in that, researchers only need to designate a DAG structure [263,264] of a BHM. Such a DAG structure can be abstractly represented as the collection  $\{ \text{data } y, \text{likelihood } p(y|\theta), \text{prior } \pi(\theta|\eta), \text{hyperprior } \pi(\eta) \}$  that should be programmed by textually or graphically, after which Bayesian software prints out simulated Markov chains from the posterior distribution  $\pi(\theta, \eta | y)$ . See [260,265] for an overall idea about how Bayesian software operates.

From the algorithmic perspective, the performance of Bayesian software may highly depend on two aspects: (i) whether the program has been designed to exploit a conditional independence structure arising from the hierarchy; and (ii) what sampling algorithms have been deployed to simulate Markov chains from a non-closed form distribution, possibly of high-dimensional. As discussed in Subsection 5.3, conditional independence is inherent in the formulation of BHM, of which proper exploitation can greatly improve the computational efficiency [128,266]. This can be mostly done by constructing a Gibbs sampling algorithm with a blocking strategy into the consideration [267]. A general rule is that the convergence of the Gibbs sampler can be improved by grouping correlated latent variables as a single parameter block to sample from as a whole [268]. On the other hand, in the task of sampling from a non-closed form distribution, we know that a naive MH algorithm [133] requires the specification of proposal density, which can be problematic in developing software. Therefore, fully automated sampling algorithms such as ESS (Algorithm 1) [120], NUTS [151], and slice sampler [121] are appreciated as general-purpose inference engines in the development of Bayesian software when conjugate-update is infeasible.

Most of the Bayesian software packages, for instance, WINBUGS [265,269], OPENBUGS [270], and JAGS [259], use three family of MCMC algorithms: Gibbs [266], MH [133], and slice sampling [121]. In contrast, STAN [17] implements HMC [119,138] and its extension, NUTS [151]. Perhaps, STAN is most extensively used Bayesian software package in recent years due to the fast converge of the inference engines regardless of whether the priors are conjugate or not. By that reason, and its great modeling flexibility, STAN has been used as a basic platform for other high-level packages like BRMS [262] and TORSTEN [261].

9. Discussion

This review of Bayesian nonlinear mixed-effects models is of necessity incomplete, as the literature is too vast to attempt even a moderate review. We have chosen to focus much of our attention on providing some of the most recent literature on the Bayesian analysis of the underlying basic model from the implementational viewpoint point. We wish that this review will offer readers familiar with frequentist analysis some pedagogical insight into the Bayesian approach, and provide those new to nonlinear mixed effects modeling a foundation of the implementation of Bayesian and frequentist computations for appreciating its idea and utility. We look forward to continuing methodological developments, software developments, and new applications of this rich class of models in industrial and academic research.

**Funding:** This research received no external funding.

**Institutional Review Board Statement:** Not applicable.

**Informed Consent Statement:** Not applicable.

**Data Availability Statement:** Not applicable.

**Conflicts of Interest:** The authors declare no conflict of interest.

References

1. Sterba, S.K. Fitting nonlinear latent growth curve models with individually varying time points. *Structural Equation Modeling: A Multidisciplinary Journal* **2014**, *21*, 630–647.

2. McArdle, J.J. Latent variable growth within behavior genetic models. *Behavior Genetics* **1986**, *16*, 163–200.

3. Cook, N.R.; Ware, J.H. Design and analysis methods for longitudinal research. *Annual Review of Public Health* **1983**, *4*, 1–23.

4. Mehta, P.D.; West, S.G. Putting the individual back into individual growth curves. *Psychological methods* **2000**, *5*, 23.

5. Zeger, S.L.; Liang, K.Y. An overview of methods for the analysis of longitudinal data. *Statistics in medicine* **1992**, *11*, 1825–1839.

6. Diggle, P.; Diggle, P.J.; Heagerty, P.; Liang, K.Y.; Zeger, S.; et al. *Analysis of longitudinal data*; Oxford university press, 2002.

7. Demidenko, E. *Mixed models: theory and applications with R*; John Wiley & Sons, 2013.

8. Snijders, T.A.; Bosker, R.J. *Multilevel analysis: An introduction to basic and advanced multilevel modeling*; sage, 2011.

9. Goldstein, H. *Multilevel statistical models*; Vol. 922, John Wiley & Sons, 2011.

10. Raudenbush, S.W.; Bryk, A.S. *Hierarchical linear models: Applications and data analysis methods*; Vol. 1, sage, 2002.

11. Efron, B. The future of indirect evidence. *Statistical science: a review journal of the Institute of Mathematical Statistics* **2010**, *25*, 145.

12. Sheiner, L.B.; Rosenberg, B.; Melmon, K.L. Modelling of individual pharmacokinetics for computer-aided drug dosage. *Computers and Biomedical Research* **1972**, *5*, 441–459.

13. Lindstrom, M.J.; Bates, D.M. Nonlinear mixed effects models for repeated measures data. *Biometrics* **1990**, pp. 673–687.

14. Davidian, M.; Giltinan, D.M. Nonlinear models for repeated measurement data: an overview and update. *Journal of agricultural, biological, and environmental statistics* **2003**, *8*, 387–419.

15. Davidian, M.; Giltinan, D.M. *Nonlinear models for repeated measurement data*; Routledge, 1995.

16. Beal, S.; et al. The NONMEM system **1980**.

17. Stan Development Team. RStan: the R interface to Stan, 2021. R package version 2.21.3.

18. Fidler, M.; Wilkins, J.J.; Hooijmaijers, R.; Post, T.M.; Schoemaker, R.; Trame, M.N.; Xiong, Y.; Wang, W. Nonlinear mixed-effects model development and simulation using nlmixr and related R open-source packages. *CPT: pharmacometrics & systems pharmacology* **2019**, *8*, 621–633.

19. Wang, W.; Hallow, K.; James, D. A tutorial on RxODE: simulating differential equation pharmacometric models in R. *CPT: pharmacometrics & systems pharmacology* **2016**, *5*, 3–10.

20. Stegmann, G.; Jacobucci, R.; Harring, J.R.; Grimm, K.J. Nonlinear mixed-effects modeling programs in R. *Structural Equation Modeling: A Multidisciplinary Journal* **2018**, *25*, 160–165.

21. Vonesh, E.; Chinchilli, V.M. *Linear and nonlinear models for the analysis of repeated measurements*; CRC press, 1996.

22. Lee, S.Y. *Structural equation modeling: A Bayesian approach*; Vol. 711, John Wiley & Sons, 2007.

23. Dellaportas, P.; Smith, A.F. Bayesian inference for generalized linear and proportional hazards models via Gibbs sampling. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* **1993**, *42*, 443–459.

24. Bush, C.A.; MacEachern, S.N. A semiparametric Bayesian model for randomised block designs. *Biometrika* **1996**, *83*, 275–285.

25. Zeger, S.L.; Karim, M.R. Generalized linear models with random effects; a Gibbs sampling approach. *Journal of the American statistical association* **1991**, *86*, 79–86.

26. Brooks, S.P. Bayesian computation: a statistical revolution. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **2003**, *361*, 2681–2697.

27. Bennett, J.; Wakefield, J. A comparison of a Bayesian population method with two methods as implemented in commercially available software. *Journal of pharmacokinetics and biopharmaceutics* **1996**, *24*, 403–432. 1162
28. Wakefield, J. The Bayesian analysis of population pharmacokinetic models. *Journal of the American Statistical Association* **1996**, *91*, 62–75. 1163
29. Gelman, A.; Bois, F.; Jiang, J. Physiological pharmacokinetic analysis using population modeling and informative prior distributions. *Journal of the American Statistical Association* **1996**, *91*, 1400–1412. 1164
30. Lee, S.Y.; Lei, B.; Mallick, B. Estimation of COVID-19 spread curves integrating global data and borrowing information. *PloS one* **2020**, *15*, e0236860. 1165
31. Lee, S.Y.; Mallick, B.K. Bayesian Hierarchical Modeling: Application Towards Production Results in the Eagle Ford Shale of South Texas. *Sankhya B* **2021**, pp. 1–43. 1166
32. Bonate, P.L.; et al. *Pharmacokinetic-pharmacodynamic modeling and simulation*; Vol. 20, Springer, 2011. 1167
33. Rowland, M.; Tozer, T.N. *Clinical pharmacokinetics/pharmacodynamics*; Lippincott Williams and Wilkins Philadelphia, 2005. 1170
34. Gabrielsson, J.; Weiner, D. *Pharmacokinetic and pharmacodynamic data analysis: concepts and applications*; CRC press, 2001. 1171
35. Dua, P.; Hawkins, E.; Van Der Graaf, P. A tutorial on target-mediated drug disposition (TMDD) models. *CPT: pharmacometrics & systems pharmacology* **2015**, *4*, 324–337. 1172
36. Xu, X.S.; Yuan, M.; Zhu, H.; Yang, Y.; Wang, H.; Zhou, H.; Xu, J.; Zhang, L.; Pinheiro, J. Full covariate modelling approach in population pharmacokinetics: understanding the underlying hypothesis tests and implications of multiplicity. *British journal of clinical pharmacology* **2018**, *84*, 1525–1534. 1173
37. Roses, A.D. Pharmacogenetics and the practice of medicine. *Nature* **2000**, *405*, 857–865. 1174
38. Guidance, F. Population Pharmacokinetics. *FDA guidance page* **1999**. 1175
39. Efron, B. *The jackknife, the bootstrap and other resampling plans*; SIAM, 1982. 1176
40. Broeker, A.; Wicha, S.G. Assessing parameter uncertainty in small-n pharmacometric analyses: value of the log-likelihood profiling-based sampling importance resampling (LLP-SIR) technique. *Journal of pharmacokinetics and pharmacodynamics* **2020**, *47*, 219–228. 1177
41. Bauer, R.J. NONMEM tutorial part I: description of commands and options, with simple examples of population analysis. *CPT: pharmacometrics & systems pharmacology* **2019**, *8*, 525–537. 1178
42. Giger, F.; Reiss, L.; Jourdan, A. The reservoir engineering aspects of horizontal drilling. SPE Annual Technical Conference and Exhibition. OnePetro, 1984. 1179
43. Al-Haddad, S.; Crafton, J. Productivity of horizontal wells. Low Permeability Reservoirs Symposium. OnePetro, 1991. 1180
44. Mukherjee, H.; Economides, M.J. A parametric comparison of horizontal and vertical well performance. *SPE Formation Evaluation* **1991**, *6*, 209–216. 1181
45. Joshi, S. Cost/benefits of horizontal wells. SPE western regional/AAPG Pacific section joint meeting. OnePetro, 2003. 1182
46. Valdes, A.; McVay, D.A.; Noynaert, S.F. Uncertainty quantification improves well construction cost estimation in unconventional reservoirs. SPE Unconventional Resources Conference Canada. OnePetro, 2013. 1183
47. Bellarby, J. *Well completion design*; Elsevier, 2009. 1184
48. Lewis, J.; Beal, C.H. Some New Methods for Estimating the Future Production of Oil Wells. *Transactions of the AIME* **1918**, *59*, 492–525. 1185
49. Currie, S.M.; Ilk, D.; Blasingame, T.A. Continuous estimation of ultimate recovery. SPE Unconventional Gas Conference. OnePetro, 2010. 1186
50. Arps, J.J.; et al. Analysis of decline curves. *Transactions of the AIME* **1945**, *160*, 228–247. 1187
51. Weibull, W.; et al. A statistical distribution function of wide applicability. *Journal of applied mechanics* **1951**, *18*, 293–297. 1188
52. Ilk, D.; Rushing, J.A.; Perego, A.D.; Blasingame, T.A.; et al. Exponential vs. hyperbolic decline in tight gas sands: understanding the origin and implications for reserve estimates using Arps' decline curves. SPE annual technical conference and exhibition. Society of Petroleum Engineers, 2008. 1189
53. Valkó, P.P.; Lee, W.J.; et al. A better way to forecast production from unconventional gas wells. SPE Annual technical conference and exhibition. Society of Petroleum Engineers, 2010. 1190
54. Clark, A.J. Decline curve analysis in unconventional resource plays using logistic growth models. PhD thesis, 2011. 1191
55. Duong, A.N.; et al. Rate-decline analysis for fracture-dominated shale reservoirs. *SPE Reservoir Evaluation & Engineering* **2011**, *14*, 377–387. 1192
56. Fetkovich, M.J.; et al. Decline curve analysis using type curves. *Journal of Petroleum Technology* **1980**, *32*, 1–065. 1193
57. Ali, T.A.; Sheng, J.J.; et al. Production Decline Models: A Comparison Study. SPE Eastern Regional Meeting. Society of Petroleum Engineers, 2015. 1194
58. Harris, S.; Lee, W.J.; et al. A Study of Decline Curve Analysis in the Elm Coulee Field. SPE Unconventional Resources Conference. Society of Petroleum Engineers, 2014. 1195
59. Miao, Y.; Li, X.; Lee, J.; Zhao, C.; Zhou, Y.; Li, H.; Chang, Y.; Lin, W.; Xiao, Z.; Wu, N.; et al. Comparison of Various Rate-Decline Analysis Models for Horizontal Wells with Multiple Fractures in Shale gas Reservoirs. SPE Trinidad and Tobago Section Energy Resources Conference. Society of Petroleum Engineers, 2018. 1196
60. Duffee, G. Forecasting interest rates. In *Handbook of economic forecasting*; Elsevier, 2013; Vol. 2, pp. 385–426. 1197

61. Gürkaynak, R.S.; Sack, B.; Wright, J.H. The US Treasury yield curve: 1961 to the present. *Journal of monetary Economics* **2007**, *54*, 2291–2304. 1220

62. Zaloom, C. How to read the future: the yield curve, affect, and financial prediction. *Public Culture* **2009**, *21*, 245–268. 1222

63. Nelson, C.R.; Siegel, A.F. Parsimonious modeling of yield curves. *Journal of business* **1987**, pp. 473–489. 1223

64. Diebold, F.X.; Li, C. Forecasting the term structure of government bond yields. *Journal of econometrics* **2006**, *130*, 337–364. 1224

65. Hays, S.; Shen, H.; Huang, J.Z. Functional dynamic factor models with application to yield curve forecasting. *The Annals of Applied Statistics* **2012**, pp. 870–894. 1225

66. Chen, Y.; Niu, L. Adaptive dynamic Nelson–Siegel term structure model with applications. *Journal of Econometrics* **2014**, *180*, 98–115. 1226

67. AC04987841, A. *Zero-coupon yield curves: Technical documentation*; Bank for Internat. Settlements, 2005. 1227

68. Hautsch, N.; Yang, F. Bayesian inference in a stochastic volatility Nelson–Siegel model. *Computational Statistics & Data Analysis* **2012**, *56*, 3774–3792. 1228

69. Diebold, F.X.; Li, C.; Yue, V.Z. Global yield curve dynamics and interactions: a dynamic Nelson–Siegel approach. *Journal of Econometrics* **2008**, *146*, 351–363. 1229

70. Cruz-Marcelo, A.; Ensor, K.B.; Rosner, G.L. Estimating the term structure with a semiparametric Bayesian hierarchical model: an application to corporate bonds. *Journal of the American Statistical Association* **2011**, *106*, 387–395. 1230

71. Richards, F. A flexible growth function for empirical use. *Journal of experimental Botany* **1959**, *10*, 290–301. 1231

72. Nelder, J.A. 182. note: An alternative form of a generalized logistic equation. *Biometrics* **1962**, *18*, 614–616. 1232

73. Seber, G.A.; Wild, C.J. Nonlinear Regression. Hoboken. *New Jersey: John Wiley & Sons* **2003**, *62*, 63. 1233

74. Anton, H.; Herr, A. *Calculus with analytic geometry*; Wiley New York, 1988. 1234

75. Causton, D. A computer program for fitting the Richards function. *Biometrics* **1969**, pp. 401–409. 1235

76. Birch, C.P. A new generalized logistic sigmoid growth equation compared with the Richards growth equation. *Annals of Botany* **1999**, *83*, 713–723. 1236

77. Kahm, M.; Hasenbrink, G.; Lichtenberg-Fraté, H.; Ludwig, J.; Kschischo, M. grofit: fitting biological growth curves with RJ Stat. Softw. 33: 1–21, 2010. 1237

78. Cao, L.; Shi, P.J.; Li, L.; Chen, G. A New Flexible Sigmoidal Growth Model. *Symmetry* **2019**, *11*, 204. 1238

79. Wang, X.S.; Wu, J.; Yang, Y. Richards model revisited: Validation by and application to infection dynamics. *Journal of Theoretical Biology* **2012**, *313*, 12–19. 1239

80. Hsieh, Y.H.; Lee, J.Y.; Chang, H.L. SARS epidemiology modeling. *Emerging infectious diseases* **2004**, *10*, 1165. 1240

81. Hsieh, Y.H. Richards model: a simple procedure for real-time prediction of outbreak severity. In *Modeling and dynamics of infectious diseases*; World Scientific, 2009; pp. 216–236. 1241

82. Hsieh, Y.H.; Ma, S. Intervention measures, turning point, and reproduction number for dengue, Singapore, 2005. *The American journal of tropical medicine and hygiene* **2009**, *80*, 66–71. 1242

83. Hsieh, Y.H.; Chen, C. Turning points, reproduction number, and impact of climatological events for multi-wave dengue outbreaks. *Tropical Medicine & International Health* **2009**, *14*, 628–638. 1243

84. Hsieh, Y.H. Pandemic influenza A (H1N1) during winter influenza season in the southern hemisphere. *Influenza and Other Respiratory Viruses* **2010**, *4*, 187–197. 1244

85. Wu, K.; Darcet, D.; Wang, Q.; Sornette, D. Generalized logistic growth modeling of the COVID-19 outbreak in 29 provinces in China and in the rest of the world. *arXiv preprint arXiv:2003.05681* **2020**. 1245

86. Lee, S.Y.; Munafo, A.; Girard, P.; Goteti, K. Optimization of dose selection using multiple surrogates of toxicity as a continuous variable in phase I cancer trial. *Contemporary clinical trials* **2021**, p. 106657. 1246

87. Dugel, P.U.; Koh, A.; Ogura, Y.; Jaffe, G.J.; Schmidt-Erfurth, U.; Brown, D.M.; Gomes, A.V.; Warburton, J.; Weichselberger, A.; Holz, F.G.; et al. HAWK and HARRIER: phase 3, multicenter, randomized, double-masked trials of brolicizumab for neovascular age-related macular degeneration. *Ophthalmology* **2020**, *127*, 72–84. 1247

88. Willyard, C. New human gene tally reignites debate. *Nature* **2018**, *558*, 354–356. 1248

89. Bishop, C.M. *Pattern recognition and machine learning*; springer, 2006. 1249

90. Goodfellow, I.; Bengio, Y.; Courville, A. *Deep learning*; MIT press, 2016. 1250

91. Boyd, S.; Boyd, S.P.; Vandenberghe, L. *Convex optimization*; Cambridge university press, 2004. 1251

92. James, W.; Stein, C. Estimation with quadratic loss. In *Breakthroughs in statistics*; Springer, 1992; pp. 443–460. 1252

93. Dawid, A.P. Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B (Methodological)* **1979**, *41*, 1–15. 1253

94. Liu, Q.; Pierce, D.A. A note on Gauss–Hermite quadrature. *Biometrika* **1994**, *81*, 624–629. 1254

95. Hedeker, D.; Gibbons, R.D. A random-effects ordinal regression model for multilevel analysis. *Biometrics* **1994**, pp. 933–944. 1255

96. Vonesh, E.F.; Wang, H.; Nie, L.; Majumdar, D. Conditional second-order generalized estimating equations for generalized linear and nonlinear mixed-effects models. *Journal of the American Statistical Association* **2002**, *97*, 271–283. 1256

97. Hinrichs, A.; Novak, E.; Ullrich, M.; Woźniakowski, H. The curse of dimensionality for numerical integration of smooth functions II. *Journal of Complexity* **2014**, *30*, 117–143. 1257

98. Beal, S.L.; Sheiner, L.B. Estimating population kinetics. *Critical reviews in biomedical engineering* **1982**, *8*, 195–222. 1258

99. Vonesh, E.F.; Carter, R.L. Mixed-effects nonlinear regression for unbalanced repeated measures. *Biometrics* **1992**, pp. 1–17. 1259



100. Goldstein, H. Nonlinear multilevel models, with an application to discrete response data. *Biometrika* **1991**, pp. 45–51. 1279
101. Vonesh, E.F. A note on the use of Laplaces approximation for nonlinear mixed-effects models. *Biometrika* **1996**, *83*, 447–452. 1280
102. Beal, S.L.; Sheiner, L.B.; Boeckmann, A.; Bauer, R.J. NONMEM users guides. *NONMEM Project Group, University of California, San Francisco* **1992**. 1281
103. Institute, S. SAS OnlineDoc, version 8., 1999. 1282
104. Marsden, J.E.; Hoffman, M.J.; et al. *Elementary classical analysis*; Macmillan, 1993. 1283
105. Lindley, D.V.; Smith, A.F. Bayes estimates for the linear model. *Journal of the Royal Statistical Society: Series B (Methodological)* **1972**, *34*, 1–18. 1284
106. Dempster, A.P.; Laird, N.M.; Rubin, D.B. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)* **1977**, *39*, 1–22. 1285
107. Meng, X.L.; Rubin, D.B. Using EM to obtain asymptotic variance-covariance matrices: The SEM algorithm. *Journal of the American Statistical Association* **1991**, *86*, 899–909. 1286
108. Walker, S. An EM algorithm for nonlinear random effects models. *Biometrics* **1996**, pp. 934–944. 1287
109. Allasonnière, S.; Chevallier, J. A new class of stochastic EM algorithms. Escaping local maxima and handling intractable sampling. *Computational Statistics & Data Analysis* **2021**, *159*, 107159. 1288
110. Kuhn, E.; Lavielle, M. Maximum likelihood estimation in nonlinear mixed effects models. *Computational statistics & data analysis* **2005**, *49*, 1020–1038. 1289
111. Samson, A.; Lavielle, M.; Mentré, F. The SAEM algorithm for group comparison tests in longitudinal data analysis based on non-linear mixed-effects model. *Statistics in medicine* **2007**, *26*, 4860–4875. 1290
112. Delyon, B.; Lavielle, M.; Moulines, E. Convergence of a stochastic approximation version of the EM algorithm. *Annals of statistics* **1999**, pp. 94–128. 1291
113. Kuhn, E.; Lavielle, M. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM: Probability and Statistics* **2004**, *8*, 115–131. 1292
114. Allasonnière, S.; Kuhn, E.; Trouvé, A. Construction of Bayesian deformable models via a stochastic approximation algorithm: a convergence study. *Bernoulli* **2010**, *16*, 641–678. 1293
115. Lavielle, M. Monolix user guide manual, 2005. 1294
116. Bernardo, J.M.; Smith, A.F. *Bayesian theory*; Vol. 405, John Wiley & Sons, 2009. 1295
117. Lindley, D.V. *Bayesian statistics, a review*; Vol. 2, SIAM, 1972. 1296
118. Casella, G.; George, E.I. Explaining the Gibbs sampler. *The American Statistician* **1992**, *46*, 167–174. 1297
119. Neal, R.M.; et al. MCMC using Hamiltonian dynamics. *Handbook of markov chain monte carlo* **2011**, *2*, 2. 1298
120. Murray, I.; Prescott Adams, R.; MacKay, D.J. Elliptical slice sampling **2010**. 1299
121. Neal, R.M. Slice sampling. *The annals of statistics* **2003**, *31*, 705–767. 1300
122. Ranganath, R.; Gerrish, S.; Blei, D. Black box variational inference. *Artificial Intelligence and Statistics*, 2014, pp. 814–822. 1301
123. Wang, C.; Blei, D.M. Variational inference in nonconjugate models. *Journal of Machine Learning Research* **2013**, *14*, 1005–1031. 1302
124. Minka, T.P. Expectation propagation for approximate Bayesian inference. *arXiv preprint arXiv:1301.2294* **2013**. 1303
125. Blei, D.M.; Kucukelbir, A.; McAuliffe, J.D. Variational inference: A review for statisticians. *Journal of the American statistical Association* **2017**, *112*, 859–877. 1304
126. Andrieu, C.; De Freitas, N.; Doucet, A.; Jordan, M.I. An introduction to MCMC for machine learning. *Machine learning* **2003**, *50*, 5–43. 1305
127. Zhang, C.; Büttepage, J.; Kjellström, H.; Mandt, S. Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence* **2018**, *41*, 2008–2026. 1306
128. Lee, S.Y. Gibbs sampler and coordinate ascent variational inference: A set-theoretical review. *Communications in Statistics-Theory and Methods* **2021**, pp. 1–21. 1307
129. Team, R.C. A language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing. URL <http://www.R-project.org> **2013**. 1308
130. Lee, A.; Yau, C.; Giles, M.B.; Doucet, A.; Holmes, C.C. On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods. *Journal of computational and graphical statistics* **2010**, *19*, 769–789. 1309
131. Suchard, M.A.; Wang, Q.; Chan, C.; Frelinger, J.; Cron, A.; West, M. Understanding GPU programming for statistical computation: Studies in massively parallel massive mixtures. *Journal of computational and graphical statistics* **2010**, *19*, 419–438. 1310
132. Hastings, W.K. Monte Carlo sampling methods using Markov chains and their applications **1970**. 1311
133. Metropolis, N.; Rosenbluth, A.W.; Rosenbluth, M.N.; Teller, A.H.; Teller, E. Equation of state calculations by fast computing machines. *The journal of chemical physics* **1953**, *21*, 1087–1092. 1312
134. Robert, C.P., The Metropolis–Hastings Algorithm. In *Wiley StatsRef: Statistics Reference Online*; John Wiley and Sons, Ltd; pp. 1–15. 1313
135. Chib, S.; Greenberg, E. Understanding the metropolis-hastings algorithm. *The american statistician* **1995**, *49*, 327–335. 1314
136. Dwivedi, R.; Chen, Y.; Wainwright, M.J.; Yu, B. Log-concave sampling: Metropolis-Hastings algorithms are fast! Conference on Learning Theory. PMLR, 2018, pp. 793–797. 1315
137. Ma, Y.A.; Chen, Y.; Jin, C.; Flammarion, N.; Jordan, M.I. Sampling can be faster than optimization. *Proceedings of the National Academy of Sciences* **2019**, *116*, 20881–20885. 1316
138. Duane, S.; Kennedy, A.D.; Pendleton, B.J.; Roweth, D. Hybrid monte carlo. *Physics letters B* **1987**, *195*, 216–222. 1317

139. Mengersen, K.L.; Tweedie, R.L. Rates of convergence of the Hastings and Metropolis algorithms. *The annals of Statistics* **1996**, *24*, 101–121. 1338
140. Chen, T.; Fox, E.; Guestrin, C. Stochastic gradient hamiltonian monte carlo. International conference on machine learning. PMLR, 2014, pp. 1683–1691. 1339
141. Aicher, C.; Ma, Y.A.; Foti, N.J.; Fox, E.B. Stochastic gradient mcmc for state space models. *SIAM Journal on Mathematics of Data Science* **2019**, *1*, 555–587. 1340
142. Griewank, A.; Walther, A. *Evaluating derivatives: principles and techniques of algorithmic differentiation*; SIAM, 2008. 1341
143. Øksendal, B. Stochastic differential equations. In *Stochastic differential equations*; Springer, 2003; pp. 65–84. 1342
144. Uhlenbeck, G.E.; Ornstein, L.S. On the theory of the Brownian motion. *Physical review* **1930**, *36*, 823. 1343
145. Roberts, G.O.; Tweedie, R.L. Geometric convergence and central limit theorems for multidimensional Hastings and Metropolis algorithms. *Biometrika* **1996**, *83*, 95–110. 1344
146. Asai, Y.; Kloeden, P.E. Numerical schemes for random ODEs via stochastic differential equations. *Commun. Appl. Anal* **2013**, *17*, 521–528. 1345
147. Casella, G.; Robert, C.P. Monte Carlo statistical methods, 1999. 1346
148. Carpenter, B.; Gelman, A.; Hoffman, M.D.; Lee, D.; Goodrich, B.; Betancourt, M.; Brubaker, M.; Guo, J.; Li, P.; Riddell, A. Stan: A probabilistic programming language. *Journal of statistical software* **2017**, *76*, 1–32. 1347
149. Abadi, M.; Barham, P.; Chen, J.; Chen, Z.; Davis, A.; Dean, J.; Devin, M.; Ghemawat, S.; Irving, G.; Isard, M.; et al. Tensorflow: A system for large-scale machine learning. 12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16), 2016, pp. 265–283. 1348
150. Leimkuhler, B.; Reich, S. *Simulating hamiltonian dynamics*; Number 14, Cambridge university press, 2004. 1349
151. Hoffman, M.D.; Gelman, A.; et al. The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **2014**, *15*, 1593–1623. 1350
152. Zou, D.; Gu, Q. On the convergence of Hamiltonian Monte Carlo with stochastic gradients. International Conference on Machine Learning. PMLR, 2021, pp. 13012–13022. 1351
153. Meza, C.; Osorio, F.; De la Cruz, R. Estimation in nonlinear mixed-effects models using heavy-tailed distributions. *Statistics and Computing* **2012**, *22*, 121–139. 1352
154. Jeffreys, H. An invariant form for the prior probability in estimation problems. *Proc. R. Soc. Lond. A* **1946**, *186*, 453–461. 1353
155. Makalic, E.; Schmidt, D.F. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* **2015**, *23*, 179–182. 1354
156. Hoff, P.D. *A first course in Bayesian statistical methods*; Vol. 580, Springer, 2009. 1355
157. Castillo, I.; Schmidt-Hieber, J.; Van der Vaart, A. Bayesian linear regression with sparse priors. *The Annals of Statistics* **2015**, *43*, 1986–2018. 1356
158. Lee, S.Y.; Pati, D.; Mallick, B.K. Tail-adaptive Bayesian shrinkage. *arXiv preprint arXiv:2007.02192* **2020**. 1357
159. Makalic, E.; Schmidt, D.F. A simple sampler for the horseshoe estimator. *IEEE Signal Processing Letters* **2016**, *23*, 179–182. 1358
160. Tibshirani, R. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)* **1996**, *58*, 267–288. 1359
161. Zou, H.; Hastie, T. Regularization and variable selection via the elastic net. *Journal of the royal statistical society: series B (statistical methodology)* **2005**, *67*, 301–320. 1360
162. Fan, J.; Samworth, R.; Wu, Y. Ultrahigh dimensional feature selection: beyond the linear model. *The Journal of Machine Learning Research* **2009**, *10*, 2013–2038. 1361
163. Lu, Y.; Stuart, A.; Weber, H. Gaussian Approximations for Probability Measures on  $\mathbb{R}^d$ . *SIAM/ASA Journal on Uncertainty Quantification* **2017**, *5*, 1136–1165. 1362
164. Wang, Y.; Blei, D.M. Frequentist consistency of variational Bayes. *Journal of the American Statistical Association* **2019**, *114*, 1147–1161. 1363
165. Johnstone, I.M. High dimensional Bernstein-von Mises: simple examples. *Institute of Mathematical Statistics collections* **2010**, *6*, 87. 1364
166. Le Cam, L.; LeCam, L.M.; Yang, G.L. *Asymptotics in statistics: some basic concepts*; Springer Science & Business Media, 2000. 1365
167. Davidian, M.; Gallant, A.R. Smooth nonparametric maximum likelihood estimation for population pharmacokinetics, with application to quinidine. *Journal of Pharmacokinetics and Biopharmaceutics* **1992**, *20*, 529–556. 1366
168. Wei, Y.; Higgins, J.P. Bayesian multivariate meta-analysis with multiple outcomes. *Statistics in medicine* **2013**, *32*, 2911–2934. 1367
169. Zellner, A. On assessing prior distributions and Bayesian regression analysis with g-prior distributions. *Bayesian inference and decision techniques* **1986**. 1368
170. Pirmohamed, M. Pharmacogenetics and pharmacogenomics. *British journal of clinical pharmacology* **2001**, *52*, 345. 1369
171. Weinshilboum, R.M.; Wang, L. Pharmacogenetics and pharmacogenomics: development, science, and translation. *Annu. Rev. Genomics Hum. Genet.* **2006**, *7*, 223–245. 1370
172. Arab-Alameddine, M.; Di Iulio, J.; Buclin, T.; Rotger, M.; Lubomirov, R.; Cavassini, M.; Fayet, A.; Décosterd, L.; Eap, C.B.; Biollaz, J.; et al. Pharmacogenetics-based population pharmacokinetic analysis of efavirenz in HIV-1-infected individuals. *Clinical Pharmacology & Therapeutics* **2009**, *85*, 485–494. 1371
173. Hastie, T.; Tibshirani, R.; Wainwright, M. *Statistical learning with sparsity: the lasso and generalizations*; Chapman and Hall/CRC, 2015. 1372
174. Mitchell, T.J.; Beauchamp, J.J. Bayesian variable selection in linear regression. *Journal of the American Statistical Association* **1988**, *83*, 1023–1032. 1373



175. George, E.I.; McCulloch, R.E. Stochastic search variable selection. *Markov chain Monte Carlo in practice* **1995**, *68*, 203–214. 1397
176. Johnson, V.E.; Rossell, D. On the use of non-local prior densities in Bayesian hypothesis tests. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2010**, *72*, 143–170. 1398  
1399
177. Yang, Y.; Wainwright, M.J.; Jordan, M.I.; et al. On the computational complexity of high-dimensional Bayesian variable selection. *The Annals of Statistics* **2016**, *44*, 2497–2532. 1400  
1401
178. Castillo, I.; van der Vaart, A.; et al. Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics* **2012**, *40*, 2069–2101. 1402  
1403
179. Park, T.; Casella, G. The bayesian lasso. *Journal of the American Statistical Association* **2008**, *103*, 681–686. 1404
180. Tipping, M.E. Sparse Bayesian learning and the relevance vector machine. *Journal of machine learning research* **2001**, *1*, 211–244. 1405
181. Griffin, J.E.; Brown, P.J.; et al. Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis* **2010**, *5*, 171–188. 1406  
1407
182. Carvalho, C.M.; Polson, N.G.; Scott, J.G. The horseshoe estimator for sparse signals. *Biometrika* **2010**, *97*, 465–480. 1408
183. Carvalho, C.M.; Polson, N.G.; Scott, J.G. Handling sparsity via the horseshoe. *Artificial Intelligence and Statistics*, 2009, pp. 73–80. 1409  
1410
184. Polson, N.G.; Scott, J.G. Shrink globally, act locally: Sparse Bayesian regularization and prediction. *Bayesian statistics* **2010**, *9*, 105. 1411
185. George, E.I.; McCulloch, R.E. Approaches for Bayesian variable selection. *Statistica sinica* **1997**, pp. 339–373. 1412
186. Johnstone, I.M.; Silverman, B.W. Needles and straw in haystacks: Empirical Bayes estimates of possibly sparse sequences. *The Annals of Statistics* **2004**, *32*, 1594–1649. 1413  
1414
187. Pati, D.; Bhattacharya, A.; Pillai, N.S.; Dunson, D.; et al. Posterior contraction in sparse Bayesian factor models for massive covariance matrices. *The Annals of Statistics* **2014**, *42*, 1102–1130. 1415  
1416
188. Song, Q.; Liang, F. Nearly optimal Bayesian shrinkage for high dimensional regression. *arXiv preprint arXiv:1712.08964* **2017**. 1417
189. Martin, R.; Mess, R.; Walker, S.G.; et al. Empirical Bayes posterior concentration in sparse high-dimensional linear models. *Bernoulli* **2017**, *23*, 1822–1847. 1418  
1419
190. Bai, R.; Ghosh, M. High-dimensional multivariate posterior consistency under global–local shrinkage priors. *Journal of Multivariate Analysis* **2018**, *167*, 157–170. 1420  
1421
191. Zhang, R.; Ghosh, M. Ultra High-dimensional Multivariate Posterior Contraction Rate Under Shrinkage Priors. *arXiv preprint arXiv:1904.04417* **2019**. 1422  
1423
192. Lee, S.; Kim, J.H. Exponentiated generalized Pareto distribution: Properties and applications towards extreme value theory. *Communications in Statistics-Theory and Methods* **2019**, *48*, 2014–2038. 1424  
1425
193. Armagan, A.; Dunson, D.B.; Lee, J. Generalized double Pareto shrinkage. *Statistica Sinica* **2013**, *23*, 119. 1426
194. O'Hara, R.B.; Sillanpää, M.J. A review of Bayesian variable selection methods: what, how and which. *Bayesian analysis* **2009**, *4*, 85–117. 1427  
1428
195. Bhadra, A.; Datta, J.; Polson, N.G.; Willard, B. Lasso meets horseshoe: A survey. *Statistical Science* **2019**, *34*, 405–427. 1429
196. Gelman, A.; et al. Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian analysis* **2006**, *1*, 515–534. 1430  
1431
197. Fan, J.; Liao, Y.; Liu, H. An overview of the estimation of large covariance and precision matrices. *The Econometrics Journal* **2016**, *19*, C1–C32. 1432  
1433
198. Bickel, P.J.; Levina, E. Covariance regularization by thresholding. *The Annals of Statistics* **2008**, *36*, 2577–2604. 1434
199. Lam, C.; Fan, J. Sparsistency and rates of convergence in large covariance matrix estimation. *Annals of statistics* **2009**, *37*, 4254. 1435
200. El Karoui, N. High-dimensionality effects in the Markowitz problem and other quadratic programs with linear constraints: Risk underestimation. *The Annals of Statistics* **2010**, *38*, 3487–3566. 1436  
1437
201. Stein, C. Estimation of a covariance matrix, Rietz Lecture. 39th Annual Meeting IMS, Atlanta, GA, 1975, 1975. 1438
202. Pourahmadi, M. *High-dimensional covariance estimation: with high-dimensional data*; Vol. 882, John Wiley & Sons, 2013. 1439
203. Ledoit, O.; Wolf, M. A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis* **2004**, *88*, 365–411. 1440  
1441
204. Rajaratnam, B.; Massam, H.; Carvalho, C.M. Flexible covariance estimation in graphical Gaussian models. *The Annals of Statistics* **2008**, *36*, 2818–2849. 1442  
1443
205. Won, J.H.; Lim, J.; Kim, S.J.; Rajaratnam, B. Condition-number-regularized covariance estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2013**, *75*, 427–450. 1444  
1445
206. Liu, C. Bartlett's Decomposition of the Posterior Distribution of the Covariance for Normal Monotone Ignorable Missing Data. *Journal of Multivariate Analysis* **1993**, *46*, 198–206. 1446  
1447
207. Barnard, J.; McCulloch, R.; Meng, X.L. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* **2000**, pp. 1281–1311. 1448  
1449
208. Geisser, S. Bayesian estimation in multivariate analysis. *The Annals of Mathematical Statistics* **1965**, *36*, 150–159. 1450
209. Lin, S.P. A Monte Carlo comparison of four estimators of a covariance matrix. *Multivariate Analysis* **1985**, pp. 411–429. 1451
210. Brown, P.J.; Le, N.D.; Zidek, J.V. Inference for a covariance matrix. *Aspects of Uncertainty: A Tribute to DV Lindley* **1994**, pp. 77–92. 1452
211. Jeffreys, H. *The theory of probability*; OUP Oxford, 1998. 1453
212. Geisser, S.; Cornfield, J. Posterior distributions for multivariate normal parameters. *Journal of the Royal Statistical Society. Series B (Methodological)* **1963**, pp. 368–376. 1454  
1455

213. Villegas, C.; et al. On the a priori distribution of the covariance matrix. *Annals of Mathematical Statistics* **1969**, *40*, 1098–1099. 1456

214. Schervish, M.J. *Theory of statistics*; Springer Science & Business Media, 2012. 1457

215. James, A.T. Distributions of matrix variates and latent roots derived from normal samples. *The Annals of Mathematical Statistics* **1964**, *35*, 475–501. 1458

216. Yang, R.; Berger, J.O. Estimation of a covariance matrix using the reference prior. *The Annals of Statistics* **1994**, pp. 1195–1211. 1459

217. Daniels, M.J.; Kass, R.E. Shrinkage estimators for covariance matrices. *Biometrics* **2001**, *57*, 1173–1184. 1460

218. Wong, F.; Carter, C.K.; Kohn, R. Efficient estimation of covariance selection models. *Biometrika* **2003**, *90*, 809–830. 1461

219. Sun, D.; Berger, J.O. Objective Bayesian analysis for the multivariate normal model. *Bayesian Statistics* **2007**, *8*, 525–562. 1462

220. Daniels, M.J.; Pourahmadi, M. Bayesian analysis of covariance matrices and dynamic models for longitudinal data. *Biometrika* **2002**, *89*, 553–566. 1463

221. Smith, M.; Kohn, R. Parsimonious covariance matrix estimation for longitudinal data. *Journal of the American Statistical Association* **2002**, *97*, 1141–1153. 1464

222. Lewandowski, D.; Kurowicka, D.; Joe, H. Generating random correlation matrices based on vines and extended onion method. *Journal of multivariate analysis* **2009**, *100*, 1989–2001. 1465

223. Ghosh, S.; Henderson, S.G. Behavior of the NORTA method for correlated random vector generation as the dimension increases. *ACM Transactions on Modeling and Computer Simulation (TOMACS)* **2003**, *13*, 276–294. 1466

224. Joe, H. Generating random correlation matrices based on partial correlations. *Journal of Multivariate Analysis* **2006**, *97*, 2177–2189. 1467

225. Gilks, W.R.; Richardson, S.; Spiegelhalter, D. *Markov chain Monte Carlo in practice*; CRC press, 1995. 1468

226. Spiegelhalter, D.J.; Best, N.G.; Carlin, B.P.; Van Der Linde, A. Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2002**, *64*, 583–639. 1469

227. Gelman, A.; Carlin, J.B.; Stern, H.S.; Rubin, D.B. *Bayesian data analysis*; Chapman and Hall/CRC, 2004. 1470

228. Watanabe, S. Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research* **2010**, *11*, 3571–3594. 1471

229. Gelfand, A.E.; Ghosh, S.K. Model choice: a minimum posterior predictive loss approach. *Biometrika* **1998**, *85*, 1–11. 1472

230. Akaike, H. Information theory and an extension of the maximum likelihood principle. In *Selected papers of hirotugu akaike*; Springer, 1998; pp. 199–213. 1473

231. Efron, B. How biased is the apparent error rate of a prediction rule? *Journal of the American statistical Association* **1986**, *81*, 461–470. 1474

232. Burnham, K.P.; Anderson, D.R. Practical use of the information-theoretic approach. In *Model selection and inference*; Springer, 1998; pp. 75–117. 1475

233. Banerjee, S.; Carlin, B.P.; Gelfand, A.E. *Hierarchical modeling and analysis for spatial data*; CRC press, 2014. 1476

234. Gelman, A.; Hwang, J.; Vehtari, A. Understanding predictive information criteria for Bayesian models. *Statistics and computing* **2014**, *24*, 997–1016. 1477

235. Celeux, G.; Forbes, F.; Robert, C.P.; Titterton, D.M.; et al. Deviance information criteria for missing data models. *Bayesian analysis* **2006**, *1*, 651–673. 1478

236. Robert, C.; Casella, G. *Monte Carlo statistical methods*; Springer Science & Business Media, 2013. 1479

237. Vehtari, A.; Gelman, A. WAIC and cross-validation in Stan. *Helsinki: Aalto University* **2014**. 1480

238. Box, G.E. Sampling and Bayes’ inference in scientific modelling and robustness. *Journal of the Royal Statistical Society. Series A (General)* **1980**, pp. 383–430. 1481

239. Zellner, A. Bayesian and non-Bayesian estimation using balanced loss functions. In *Statistical decision theory and related topics V*; Springer, 1994; pp. 377–390. 1482

240. Vonesh, E.F. Non-linear models for the analysis of longitudinal data. *Statistics in medicine* **1992**, *11*, 1929–1954. 1483

241. Müller, P.; Rosner, G.L. A Bayesian population model with hierarchical mixture priors applied to blood count data. *Journal of the American Statistical Association* **1997**, *92*, 1279–1292. 1484

242. Müller, P.; Quintana, F.A. Nonparametric Bayesian data analysis. *Statistical science* **2004**, *19*, 95–110. 1485

243. Hjort, N.L.; Holmes, C.; Müller, P.; Walker, S.G. *Bayesian nonparametrics*; Vol. 28, Cambridge University Press, 2010. 1486

244. Walker, S.; Wakefield, J. Population models with a nonparametric random coefficient distribution. *Sankhyā: The Indian Journal of Statistics, Series B* **1998**, pp. 196–214. 1487

245. MacKay, D.J.; et al. Introduction to Gaussian processes. *NATO ASI series F computer and systems sciences* **1998**, *168*, 133–166. 1488

246. Rasmussen, C.E. Gaussian processes in machine learning. Summer school on machine learning. Springer, 2003, pp. 63–71. 1489

247. Ferguson, T.S. Prior distributions on spaces of probability measures. *The annals of statistics* **1974**, *2*, 615–629. 1490

248. Escobar, M.D. Estimating normal means with a Dirichlet process prior. *Journal of the American Statistical Association* **1994**, *89*, 268–277. 1491

249. Escobar, M.D.; West, M. Bayesian density estimation and inference using mixtures. *Journal of the american statistical association* **1995**, *90*, 577–588. 1492

250. McLachlan, G.J.; Lee, S.X.; Rathnayake, S.I. Finite mixture models. *Annual review of statistics and its application* **2019**, *6*, 355–378. 1493

251. Rasmussen, C.E.; et al. The infinite Gaussian mixture model. *NIPS. Citeseer*, 1999, Vol. 12, pp. 554–560. 1494

252. Antoniak, C.E. Mixtures of Dirichlet processes with applications to Bayesian nonparametric problems. *The annals of statistics* **1974**, pp. 1152–1174. 1495

253. Teh, Y.W.; Jordan, M.I.; Beal, M.J.; Blei, D.M. Hierarchical dirichlet processes. *Journal of the american statistical association* **2006**, *101*, 1566–1581. 1514

254. Jara, A. Theory and computations for the Dirichlet process and related models: an overview. *International Journal of Approximate Reasoning* **2017**, *81*, 128–146. 1515

255. Rosner, G.L.; Müller, P. Bayesian population pharmacokinetic and pharmacodynamic analyses using mixture models. *Journal of pharmacokinetics and biopharmaceutics* **1997**, *25*, 209–233. 1516

256. Müller, P.; Quintana, F.; Rosner, G. A method for combining inference across related nonparametric Bayesian models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **2004**, *66*, 735–749. 1517

257. Brown, H.; Prescott, R. *Applied mixed models in medicine*; John Wiley & Sons, 2015. 1518

258. Congdon, P.D. *Applied Bayesian hierarchical methods*; CRC Press, 2010. 1519

259. Plummer, M.; et al. JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. Proceedings of the 3rd international workshop on distributed statistical computing. Vienna, Austria., 2003, Vol. 124, pp. 1–10. 1520

260. Lunn, D.; Spiegelhalter, D.; Thomas, A.; Best, N. The BUGS project: Evolution, critique and future directions. *Statistics in medicine* **2009**, *28*, 3049–3067. 1521

261. Margossian, C.; Gillespie, B. Torsten: A Prototype Model Library for Bayesian PKPD Modeling in Stan User Manual: Version 0.81. 1522

262. Bürkner, P.C. brms: An R package for Bayesian multilevel models using Stan. *Journal of statistical software* **2017**, *80*, 1–28. 1523

263. Jordan, M.I. Graphical models. *Statistical science* **2004**, *19*, 140–155. 1524

264. Lauritzen, S.L.; Dawid, A.P.; Larsen, B.N.; Leimer, H.G. Independence properties of directed Markov fields. *Networks* **1990**, *20*, 491–505. 1525

265. Lunn, D.J.; Thomas, A.; Best, N.; Spiegelhalter, D. WinBUGS-a Bayesian modelling framework: concepts, structure, and extensibility. *Statistics and computing* **2000**, *10*, 325–337. 1526

266. Geman, S.; Geman, D. Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on pattern analysis and machine intelligence* **1984**, pp. 721–741. 1527

267. Liu, J.S. The collapsed Gibbs sampler in Bayesian computations with applications to a gene regulation problem. *Journal of the American Statistical Association* **1994**, *89*, 958–966. 1528

268. Park, T.; Lee, S. Improving the Gibbs sampler. *Wiley Interdisciplinary Reviews: Computational Statistics* **2021**, p. e1546. 1529

269. Spiegelhalter, D.J.; Thomas, A.; Best, N.; Lunn, D. WinBUGS version 1.4 user manual. MRC Biostatistics Unit, Cambridge. URL <http://www.mrc-bsu.cam.ac.uk/bugs> **2003**. 1530

270. Spiegelhalter, D.; Thomas, A.; Best, N.; Lunn, D. OpenBUGS user manual. Version **2007**, 3, 2007. 1531