

Some Developments in Bayesian Hierarchical Linear Regression Modeling

Algunos desarrollos en modelos de regresión lineal jerárquicos bayesianos

JUAN SOSA^{1,a}, JEIMY-PAOLA ARISTIZABAL^{1,2,b}

¹DEPARTAMENTO DE ESTADÍSTICA, FACULTAD DE CIENCIAS, UNIVERSIDAD NACIONAL DE COLOMBIA, BOGOTÁ, COLOMBIA

²FACULTAD DE ADMINISTRACIÓN DE EMPRESAS, UNIVERSIDAD EXTERNADO DE COLOMBIA, BOGOTÁ, COLOMBIA

Abstract

Considering the flexibility and applicability of Bayesian modeling, in this work we revise the main characteristics of two hierarchical models in a regression setting. We study the full probabilistic structure of the models along with the full conditional distribution for each model parameter. Under our hierarchical extensions, we allow the mean of the second stage of the model to have a linear dependency on a set of covariates. The Gibbs sampling algorithms used to obtain samples when fitting the models are fully described and derived. In addition, we consider a case study in which the plant size is characterized as a function of nitrogen soil concentration and a grouping factor (farm).

Key words: Bayesian inference; Clustering; Gibbs Sampling; Hierarchical model; Linear regression.

Resumen

Considerando la flexibilidad y aplicabilidad del modelamiento Bayesiano, en este trabajo se revisan las principales características de dos modelos jerárquicos en un escenario de regresión. Se estudia la estructura probabilística completa de los modelos junto con la distribución condicional completa para cada parámetro del modelo. Las extensiones jerárquicas que se presentan permiten que la media de la segunda etapa del modelo tenga una dependencia lineal de un conjunto de covariables. Se describen y derivan completamente los algoritmos de muestreo de Gibbs para ajustar los modelos. Además, se considera un caso de estudio en el que se caracteriza el

^aPh.D. E-mail: jcsosam@unal.edu.co

^bPh.D. E-mail: jparistizabalr@gmail.com

tamaño de plantas en función de la concentración de nitrógeno en el suelo y un factor de agrupación (fincas).

Palabras clave: Agrupamiento; Inferencia bayesiana; Muestreador de Gibbs; Modelo jerárquico; Regresión lineal.

1. Introduction

A key characteristic of many problems is that the observed data can be used to estimate aspects of the population even though they are never observed. Often, it is quite natural to model such a problem hierarchically, with observable outcomes modeled conditionally on certain parameters, which themselves are assigned a probabilistic specification in terms of further random quantities.

Hierarchical models can have enough parameters to fit the data well, while using a population distribution to structure some dependence into the parameters, thereby avoiding problems of over-fitting. In addition, by establishing hierarchies we are not forced to choose between complete pooling and not pooling at all as the classic analysis of variance does (Gelman et al., 2013).

In this work we analyze observational continuous data arranged in groups. Specifically, we discuss hierarchical models for the comparison of group-specific parameters across groups in a regression setting. Our hierarchical approach is conceptually a straightforward generalization of a standard Normal model. Emulating Hoff (2009), we use an ordinary regression model to describe within-group heterogeneity of observations, and also, describe between-group heterogeneity using a sampling model for the group-specific regression parameters. Then, we take a step further and develop another hierarchical model adding a specific layer for carrying out clustering tasks. At this point, we explicitly note that the main difference between two given models relies on the stochastic structure of the regression parameters. For instance, the hierarchical approach (also known as HLRM later) embraces a way to borrow information across groups through a shrinkage effect promoted by the prior formulation, whereas the clustering specification (also known as CHLRM later) additionally aims to collect groups with similar fixed effects.

Since the intended audience for this article needs to be knowledgeable in statistical methods, the reader should be aware that hierarchical linear modeling from a Bayesian perspective is a well developed area within the Statistics literature. There are available many technical and methodological developments ranging from standard linear models (Gelman et al., 2013), generalized linear models (Dey et al., 2000) to even more sophisticated such as nonparametric linear models (Müller et al., 2015). Finally, the literature is also very broad in the treatment of these topics when applied to specific statistical disciplines such as longitudinal data analysis (Wakefield, 2013), spatio-temporal data analysis (Banerjee et al., 2014), and statistical analysis of network data (Kolaczyk & Csárdi, 2020), just to name a few.

Even though our modeling approach is quite conventional under the Bayesian paradigm, our main contribution mainly relies in how models are structured and developed. Thus, we provide for each model all the details regarding model and (hierarchical) prior specification, careful hyperparameter selection under little external information, and simulation-based algorithms for computation. Finally, we also consider several metrics to quantify the performance aspects of our models by means of both goodness-of-fit and in-sample model checking metrics. In every instance along the way, we go over several building blocks of the literature about hierarchical modeling techniques and clustering tasks.

This paper is structured as follows: Section 2 revisits all the details related to the Normal linear regression model (LRM). Sections 3 and 4 provide a full development of hierarchical Normal linear regression models (HLRM) as well as clustering hierarchical Normal linear regression models (CHLRM). Section 5 discusses in depth other modeling approaches. Section 6 shows several specifics related to computation and model fitting. Section 7 presents several specifics about model checking through tests statistics, and also, model selection through information criteria. Section 8 makes a complete analyzes of a case study. Finally, Section 9 discusses our findings and future developments.

2. Normal Linear Regression Model (LRM)

Here, we show some relevant aspects about linear regression modeling in a Bayesian setting, which is a powerful data analysis tool quite useful for carrying out many inferential tasks such as data characterization and prediction. Roughly speaking, our goal is to find a model for predicting the dependent variable (response) given one or more independent (predictor) variables.

2.1. Model Specification

First, we consider a simple scenario in which we want to characterize the sampling distribution of a random variable y through a set of explanatory variables $\mathbf{x} = [x_1, \dots, x_p]^T$. Thus, we look upon a Normal linear regression model of the form

$$y_{i,j} = \mathbf{x}_{i,j}^T \boldsymbol{\beta} + \epsilon_{i,j}, \quad \epsilon_{i,j} \mid \sigma^2 \stackrel{\text{iid}}{\sim} \text{N}(0, \sigma^2), \quad i = 1, \dots, n, \quad j = 1, \dots, m,$$

where $y_{i,j}$, $\mathbf{x}_{i,j}$, and $\epsilon_{i,j}$ are the response variable, the covariates, and the random error, respectively, corresponding to the i -th observation from the j -th group, and $\boldsymbol{\beta} = [\beta_1, \dots, \beta_p]^T$ are the regression parameters of the model. Note that the previous model can be re-expressed as

$$\mathbf{y} \mid \mathbf{X}, \boldsymbol{\beta}, \sigma^2 \sim \text{N}_{nm}(\mathbf{X}\boldsymbol{\beta}, \sigma^2 \mathbf{I})$$

where \mathbf{y} is the response vector given by $\mathbf{y} = [\mathbf{y}_1^T, \dots, \mathbf{y}_m^T]^T$, with $\mathbf{y}_j = [y_{j,1}, \dots, y_{j,n}]^T$, and \mathbf{X} is the design matrix arranged in a similar fashion.

In order to perform full Bayesian analysis using the likelihood given above, we consider a semi-conjugate prior distribution for β and σ^2 of the form

$$\beta \sim N(\beta_0, \Sigma_0) \quad \text{and} \quad \sigma^2 \sim \text{IG}(\nu_0/2, \nu_0\sigma_0^2/2)$$

as in most Normal sampling problems.

2.2. Prior Elicitation

In the absence of convincing external information to the data, it is customary using a diffuse prior distribution in order to be as minimally informative as possible. In the same spirit of the unit information prior as in [Kass & Wasserman \(1995\)](#), we let $g = nm$ and $\nu_0 = 1$, and set $\beta_0 = \hat{\beta}_{\text{OLS}}$, $\Sigma_0 = g\sigma_0^2(\mathbf{X}^\top \mathbf{X})^{-1}$, with $\sigma_0^2 = \hat{\sigma}_{\text{OLS}}^2$, where $\hat{\phi}_{\text{OLS}}$ stands for the ordinary least squares (OLS) estimate of ϕ . This choice of g makes the ratio $\frac{g}{g+1}$ very close to 1, and therefore, β is practically centered around $\mathbf{0}$; similarly, the prior distribution of σ^2 is weakly centered around $\hat{\sigma}_{\text{OLS}}^2$ since $\nu_0 = 1$. Note that large values of g as well as small values of ν_0 reflect weak prior beliefs. This prior distribution cannot be strictly considered a real prior distribution, as it requires knowledge of \mathbf{y} to be constructed. However, it only uses a small amount of the information given in \mathbf{y} , and can be loosely thought of as the prior distribution of a researcher with unbiased but weak prior information.

2.3. Posterior Inference

The posterior distribution can be explored using Markov chain Monte Carlo (MCMC) methods ([Gamerman & Lopes, 2006](#)) such as the Gibbs sampling. Implementing such an algorithm under the previous model is quite simple since the full conditional distributions are

$$\beta \mid \mathbf{y}, \mathbf{X}, \sigma^2 \sim N_p((\Sigma_0^{-1} + \sigma^{-2}\mathbf{X}^\top \mathbf{X})^{-1}(\Sigma_0^{-1}\beta_0 + \sigma^{-2}\mathbf{X}^\top \mathbf{y}), (\Sigma_0^{-1} + \sigma^{-2}\mathbf{X}^\top \mathbf{X})^{-1})$$

and

$$\sigma^2 \mid \mathbf{y}, \mathbf{X}, \beta \sim \text{IG}((\nu_0 + nm)/2, (\nu_0\sigma_0^2 + (\mathbf{y} - \mathbf{X}\beta)^\top (\mathbf{y} - \mathbf{X}\beta))/2).$$

See for example [Christensen et al. \(2011\)](#) for details about this result.

Under the previous setting, we need to carefully choose values for the set of model hyperparameters, namely, β_0 , Σ_0 , ν_0 , and σ_0^2 . Often, the analysis must be done in the absence of prior information, so we should use a prior distribution as minimally informative as possible. The so-called g -priors (see for example [Albert \(2009\)](#) for a brief discussion) offer this possibility and the desirable feature of invariance to changes in the scale of the regressors. A popular alternative in this direction consists in letting $\beta_0 = \mathbf{0}$ and $\Sigma_0 = g\sigma^2(\mathbf{X}^\top \mathbf{X})^{-1}$ for some positive value g that reflects the amount of information in the data relative to the prior distribution (choosing a large value of g naturally induces a diffuse prior). It can be shown that under the g -prior specification, $p(\sigma^2 \mid \mathbf{y}, \mathbf{X})$ is an Inverse Gamma distribution, which means that we can use direct sampling as follows:

1. Sample $\sigma^2 \sim \text{IG}\left((\nu_0 + nm)/2, (\nu_0\sigma_0^2 + \mathbf{y}^\top(\mathbf{I} - \frac{g}{g+1}\mathbf{X}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top)\mathbf{y})/2\right)$.
2. Sample $\boldsymbol{\beta} \sim \text{N}_p\left(\frac{g}{g+1}(\mathbf{X}^\top\mathbf{X})^{-1}\mathbf{X}^\top\mathbf{y}, \frac{g}{g+1}\sigma^2(\mathbf{X}^\top\mathbf{X})^{-1}\right)$.

See for example Hoff (2009) for more details.

3. Hierarchical Normal Linear Regression Model (HLRM)

We present the treatment of a hierarchical model, in which the observed data is assumed to be normally distributed with both group-specific fixed effects (and therefore subject-specific means) and group-specific variances. The model specification provided below is quite convenient because in addition to a global assessment of the mean relationship between the covariates and the response variable (as allowed by a standard linear regression model), it gives the means to carry out specific inferences within each group as well as comparisons among groups.

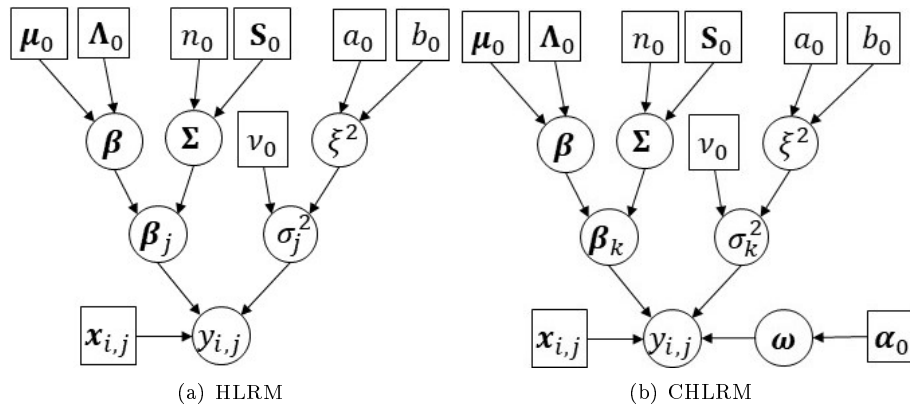


FIGURE 1: DAG representations.

3.1. Model Specification

We consider m independent groups, each one of them with n independent normally distributed data points (i.e., a balanced experiment), $y_{i,j}$, each of which with subject-specific mean $\mu_{i,j} = \mathbf{x}_{i,j}^\top\boldsymbol{\beta}_j$, with $\boldsymbol{\beta}_j = (\beta_{1,j}, \dots, \beta_{p,j})$, and group-specific variance σ_j^2 ; i.e.,

$$y_{i,j} \mid \mathbf{x}_{i,j}, \boldsymbol{\beta}_j, \sigma_j^2 \stackrel{\text{ind}}{\sim} \text{N}(\mathbf{x}_{i,j}^\top\boldsymbol{\beta}_j, \sigma_j^2), \quad i = 1, \dots, n, \quad j = 1, \dots, m.$$

In addition, we propose a hierarchical prior distribution with the following stages:

$$\boldsymbol{\beta}_j \mid \boldsymbol{\beta}, \boldsymbol{\Sigma} \stackrel{\text{iid}}{\sim} \text{N}_p(\boldsymbol{\beta}, \boldsymbol{\Sigma}) \quad \text{and} \quad \sigma_j^2 \mid \xi^2 \stackrel{\text{iid}}{\sim} \text{IG}(\nu_0/2, \nu_0\xi^2/2),$$

with

$$\boldsymbol{\beta} \sim N_p(\boldsymbol{\mu}_0, \boldsymbol{\Lambda}_0), \quad \boldsymbol{\Sigma} \sim IW(n_0, \mathbf{S}_0^{-1}), \quad \xi^2 \sim G(a_0, b_0),$$

where $\boldsymbol{\zeta} = (\boldsymbol{\beta}_1, \dots, \boldsymbol{\beta}_m)$, $\boldsymbol{\sigma}^2 = (\sigma_1^2, \dots, \sigma_m^2)$, $\boldsymbol{\beta}$, $\boldsymbol{\Sigma}$, and ξ^2 are the unknown model parameters, and $\boldsymbol{\mu}_0$, $\boldsymbol{\Lambda}_0$, n_0 , \mathbf{S}_0 , ν_0 , a_0 , and b_0 are hyperparameters carefully selected according to external information. Figure 1 provides a directed acyclic graph (DAG) representation of the model. As a final remark, note that fitting this hierarchical model is not equivalent to fitting regular regression models to each group independently since the information shared across groups (shrinkage effect) would be lost.

3.2. Prior Elicitation

Following the same unit-information-prior-inspired approach considered to select the hyperparameters of the Normal linear regression model (once again we refer the reader to Kass & Wasserman (1995)), we again let $g = nm$ and $\nu_0 = 1$, and set $\boldsymbol{\mu}_0 = \hat{\boldsymbol{\beta}}_{\text{OLS}}$, $\boldsymbol{\Lambda}_0 = g\sigma_0^2(\mathbf{X}^\top \mathbf{X})^{-1}$, with $\sigma_0^2 = \hat{\sigma}_{\text{OLS}}^2$ (see Section 2 for details). In addition, aiming to establishing a diffuse and reasonable centered prior for $\boldsymbol{\Sigma}$, we let $n_0 = p + 2$ and $\mathbf{S}_0 = \boldsymbol{\Lambda}_0$ because such a specification produces a mean vague concentration of $\boldsymbol{\beta}$ around $\boldsymbol{\mu}_0$ since $\mathbf{E}(\boldsymbol{\Sigma}) = \boldsymbol{\Lambda}_0$ a priori. Finally, we let $a_0 = 1$ and $b_0 = 1/\sigma_0^2$ because this choice leads to a diffuse prior for ξ^2 such that $\mathbf{E}(\xi^2) = \sigma_0$ with $\text{CV}(\xi^2) = 1$, which clearly emulates the prior elicitation in a regular regression setting for which $\mathbf{E}(\sigma^2) = \sigma_0^2$.

3.3. Posterior Inference

Joint posterior inference for the model parameters can be achieved by constricting a Gibbs sampling algorithm (Gamerman & Lopes, 2006), which requires iteratively sampling each parameter from its full conditional distribution.

Let $\boldsymbol{\Theta} = (\boldsymbol{\zeta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta}, \boldsymbol{\Sigma}, \xi^2)$ be the full set of parameters in the model. The posterior distribution of $\boldsymbol{\Theta}$ is

$$p(\boldsymbol{\Theta} \mid \mathbf{y}, \mathbf{X}) \propto p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\zeta}, \boldsymbol{\sigma}^2) p(\boldsymbol{\zeta} \mid \boldsymbol{\beta}, \boldsymbol{\Sigma}) p(\boldsymbol{\beta}) p(\boldsymbol{\Sigma}) p(\boldsymbol{\sigma}^2 \mid \xi^2) p(\xi^2),$$

which leads to

$$\begin{aligned} p(\boldsymbol{\Theta} \mid \mathbf{y}, \mathbf{X}) &\propto \prod_{j=1}^m \prod_{i=1}^n \sigma_j^{-1/2} \exp \left\{ -\frac{1}{2\sigma_j^2} \left(y_{i,j} - \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_j \right)^2 \right\} \\ &\times \prod_{j=1}^m |\boldsymbol{\Sigma}|^{-1/2} \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_j - \boldsymbol{\beta})^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\beta}_j - \boldsymbol{\beta}) \right\} \\ &\times \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta} - \boldsymbol{\mu}_0)^\top \boldsymbol{\Lambda}_0^{-1} (\boldsymbol{\beta} - \boldsymbol{\mu}_0) \right\} \times |\boldsymbol{\Sigma}|^{-(n_0+p+1)/2} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}_0 \boldsymbol{\Sigma}^{-1}) \right\} \\ &\times \prod_{j=1}^m (\xi^2)^{\nu_0/2} (\sigma_j^2)^{-(\nu_0/2+1)} \exp \left\{ -\frac{\nu_0 \xi^2 / 2}{\sigma_j^2} \right\} \times (\xi^2)^{a_0-1} \exp \{-b_0 \xi^2\}. \end{aligned}$$

Let ϕ represent any parameter in $\boldsymbol{\Theta}$. The full conditional distribution (fcd) of ϕ given the rest of the parameters, the design matrix \mathbf{X} , and the data \mathbf{y} is denoted

by $p(\phi \mid \text{rest})$. We derived these distributions looking at the dependencies in the full posterior distribution. Thus, we have that:

- The fcd of β_j , for $j = 1, \dots, m$, is

$$\beta_j \mid \text{rest} \sim N_p \left(\left(\Sigma^{-1} + \sigma_j^{-2} \mathbf{X}_j^T \mathbf{X}_j \right)^{-1} \left(\Sigma^{-1} \beta + \sigma_j^{-2} \mathbf{X}_j^T \mathbf{y}_j \right), \left(\Sigma^{-1} + \sigma_j^{-2} \mathbf{X}_j^T \mathbf{X}_j \right)^{-1} \right)$$

where $\mathbf{X}_j = [\mathbf{x}_{1,j}, \dots, \mathbf{x}_{n,j}]^T$.

- The fcd for β is

$$\beta \mid \text{rest} \sim N_p \left(\left(\Lambda_0^{-1} + m \Sigma^{-1} \right)^{-1} \left(\Lambda_0^{-1} \mu_0 + \Sigma^{-1} \sum_{j=1}^m \beta_j \right), \left(\Lambda_0^{-1} + m \Sigma^{-1} \right)^{-1} \right).$$

- The fcd of Σ is

$$\Sigma \mid \text{rest} \sim IW \left(n_0 + m, \left(\mathbf{S}_0 + \sum_{j=1}^m (\beta_j - \beta)(\beta_j - \beta)^T \right)^{-1} \right).$$

- The fcd of σ_j^2 , for $j = 1, \dots, m$, is

$$\sigma_j^2 \mid \text{rest} \sim IG \left((\nu_0 + n)/2, (\nu_0 \xi^2 + \sum_{i=1}^n (y_{i,j} - \mathbf{x}_{i,j}^T \beta_j)^2) / 2 \right).$$

- The fcd of ξ^2 is

$$\xi^2 \mid \text{rest} \sim G \left(a_0 + m \nu_0 / 2, b_0 + \frac{\nu_0}{2} \sum_{j=1}^m \sigma_j^{-2} \right).$$

Let $\phi^{(b)}$ denote the state of parameter ϕ in the b -th iteration of the Gibbs sampling algorithm, for $b = 1, \dots, B$. Then, such an algorithm in this case is as follows:

1. Choose an initial configuration for each parameter in the model, say $\beta_1^{(0)}, \dots, \beta_m^{(0)}, \beta^{(0)}, \Sigma^{(0)}, (\sigma_1^2)^{(0)}, \dots, (\sigma_m^2)^{(0)}$, and $(\xi^2)^{(0)}$.
2. Update $\beta_1^{(b-1)}, \dots, \beta_m^{(b-1)}, \beta^{(b-1)}, \Sigma^{(b-1)}, (\sigma_1^2)^{(b-1)}, \dots, (\sigma_m^2)^{(b-1)}$, and $(\xi^2)^{(b-1)}$ until convergence:
 - a) Sample $\beta_j^{(b)}$ from the fcd $p(\beta_j \mid (\sigma_j^2)^{(b-1)}, \beta^{(b-1)}, \Sigma^{(b-1)}, \mathbf{y}_j, \mathbf{X}_j)$, for $j = 1, \dots, m$.
 - b) Sample $\beta^{(b)}$ from the fcd $p(\beta \mid \{\beta_j^{(b)}\}, \Sigma^{(b-1)})$.
 - c) Sample $\Sigma^{(b)}$ from the fcd $p(\Sigma \mid \{\beta_j^{(b)}\}, \beta^{(b)})$.
 - d) Sample $(\sigma_j^2)^{(b)}$ from the fcd $p(\sigma_j^2 \mid \beta_j^{(b)}, (\xi^2)^{(b-1)}, \mathbf{y}_j, \mathbf{X}_j)$, for $j = 1, \dots, m$.
 - e) Sample $(\xi^2)^{(b)}$ from the fcd $p(\xi^2 \mid \{(\sigma_j^2)^{(b)}\})$.
3. Cycle until achieve convergence.

4. Clustering Hierarchical Normal Linear Regression Model (CHLRM)

Here we extend the hierarchical approach provided in the previous section by adding more structure to the model in order to perform clustering task at a group level (clusters whose elements are groups). Specifically, we consider a mixture model instead of a regular Normal distribution in the likelihood, and then, introduce a new set of parameters known as cluster assignments to “break” the mixture and then be able to identify those groups belonging to the same cluster.

4.1. Model Specification

A natural way to extend the standard hierarchical model consists in relaxing the normality assumption about the response variable by replacing it with a finite mixture of Normal components in such a way that

$$y_{i,j} \mid \mathbf{x}_{i,j}, \{\beta_k\} \{\sigma_k^2\} \stackrel{\text{iid}}{\sim} \sum_{k=1}^K \omega_k \mathbf{N}(y_{i,j} \mid \mathbf{x}_{i,j}^\top \beta_k, \sigma_k^2), \quad i = 1, \dots, n, \quad j = 1, \dots, m,$$

where K is a positive fixed integer that represents the number of clusters in which groups can be classified, β_1, \dots, β_K and $\sigma_1^2, \dots, \sigma_K^2$ are the cluster-specific regression parameters and cluster-specific variances of the mixture components, and $\omega_1, \dots, \omega_K$ are mixture probabilities such that $0 < \omega_k < 1$ and $\sum_{k=1}^K \omega_k = 1$. Note that under this formulation we recover the Normal linear regression model by setting $K = 1$.

According to the previous mixture, the probability that the group j is part of the cluster k is ω_k , i.e., $\Pr(\gamma_j = k \mid \omega_k) = \omega_k$, for $j = 1, \dots, m$ and $k = 1, \dots, K$, where γ_j is a categorical variable (known as either cluster assignment or cluster indicator) that takes integer values in $1, \dots, K$ with probabilities $\omega_1, \dots, \omega_K$, respectively. Thus, we can use the cluster assignments $\gamma_1, \dots, \gamma_m$ to “break” the mixture and write the model as

$$y_{i,j} \mid \mathbf{x}_{i,j}, \gamma_j, \beta_{\gamma_j}, \sigma_{\gamma_j}^2 \stackrel{\text{iid}}{\sim} \mathbf{N}(y_{i,j} \mid \mathbf{x}_{i,j}^\top \beta_{\gamma_j}, \sigma_{\gamma_j}^2) \quad i = 1, \dots, n, \quad j = 1, \dots, m.$$

In addition, a parsimonious way to formulate a hierarchical prior distribution can be achieved by letting

$$\gamma \mid \omega \sim \text{Cat}(\omega), \quad \beta_k \mid \beta, \Sigma \stackrel{\text{iid}}{\sim} \mathbf{N}_p(\beta, \Sigma) \quad \text{and} \quad \sigma_k^2 \mid \xi^2 \stackrel{\text{iid}}{\sim} \text{IG}(\nu_0/2, \nu_0 \xi^2/2),$$

with

$$\omega \sim \text{Dir}(\alpha_0), \quad \beta \sim \mathbf{N}_p(\mu_0, \Lambda_0), \quad \Sigma \sim \text{IW}(n_0, \mathbf{S}_0^{-1}), \quad \xi^2 \sim \mathbf{G}(a_0, b_0),$$

where $\gamma = (\gamma_1, \dots, \gamma_m)$, $\omega = (\omega_1, \dots, \omega_K)$, $\zeta = (\beta_1, \dots, \beta_K)$, $\sigma^2 = (\sigma_1^2, \dots, \sigma_K^2)$, β , Σ , and ξ^2 are the unknown model parameters, and α_0 , μ_0 , Λ_0 , n_0 , \mathbf{S}_0 , ν_0 , a_0 , and b_0 are hyperparameters carefully selected according to external information. Finally, note that the DAG representation of the model is very similar to that of HLRM, but including an extra random node corresponding to the clustering process (see Figure 1).

4.2. Prior Elicitation

Under the same approach as before, we let once again $g = nm$ and $\nu_0 = 1$, and set $\boldsymbol{\mu}_0 = \hat{\boldsymbol{\beta}}_{\text{OLS}}$, $\boldsymbol{\Lambda}_0 = g \sigma_0^2 (\mathbf{X}^\top \mathbf{X})^{-1}$, with $\sigma_0^2 = \hat{\sigma}_{\text{OLS}}^2$, $n_0 = p + 2$, $\mathbf{S}_0 = \boldsymbol{\Lambda}_0$, $a_0 = 1$, and $b_0 = 1/\sigma_0^2$. Thus, the only hyperparameter that remains unspecified is $\boldsymbol{\alpha}_0$. To do so in a sensible way, we let $\boldsymbol{\alpha}_0 = (\frac{1}{K}, \dots, \frac{1}{K})$, which has a direct connection with a Chinese restaurant process prior (Ishwaran & Zarepour, 2000) and places a diffuse prior distribution for the number of occupied clusters in the data.

4.3. Posterior Inference

Once again we appeal to MCMC methods as in Section 3 to explore the posterior distribution of the model parameters $\boldsymbol{\Theta} = (\boldsymbol{\gamma}, \boldsymbol{\omega}, \boldsymbol{\zeta}, \boldsymbol{\sigma}^2, \boldsymbol{\beta}, \boldsymbol{\Sigma}, \boldsymbol{\xi}^2)$. The posterior distribution of $\boldsymbol{\Theta}$ is such that

$$p(\boldsymbol{\Theta} \mid \mathbf{y}, \mathbf{X}) \propto p(\mathbf{y} \mid \mathbf{X}, \boldsymbol{\gamma}, \boldsymbol{\zeta}, \boldsymbol{\sigma}^2) p(\boldsymbol{\gamma} \mid \boldsymbol{\omega}) p(\boldsymbol{\omega}) p(\boldsymbol{\zeta} \mid \boldsymbol{\beta}, \boldsymbol{\Sigma}) p(\boldsymbol{\beta}) p(\boldsymbol{\Sigma}) p(\boldsymbol{\sigma}^2 \mid \boldsymbol{\xi}^2) p(\boldsymbol{\xi}^2),$$

which leads to

$$\begin{aligned} p(\boldsymbol{\Theta} \mid \mathbf{y}, \mathbf{X}) &\propto \prod_{j=1}^m \prod_{i=1}^n \sigma_j^{-1/2} \exp \left\{ -\frac{1}{2\sigma_j^2} \left(y_{i,j} - \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_{\gamma_j} \right)^2 \right\} \times \prod_{j=1}^m \prod_{k=1}^K \omega_k^{[\gamma_j=k]} \times \prod_{k=1}^K \omega_k^{\alpha_{0k}} \\ &\times \prod_{k=1}^K |\boldsymbol{\Sigma}|^{-1/2} \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta}_k - \boldsymbol{\beta})^\top \boldsymbol{\Sigma}^{-1} (\boldsymbol{\beta}_k - \boldsymbol{\beta}) \right\} \\ &\times \exp \left\{ -\frac{1}{2} (\boldsymbol{\beta} - \boldsymbol{\mu}_0)^\top \boldsymbol{\Lambda}_0^{-1} (\boldsymbol{\beta} - \boldsymbol{\mu}_0) \right\} \times |\boldsymbol{\Sigma}|^{-(n_0+p+1)/2} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}_0 \boldsymbol{\Sigma}^{-1}) \right\} \\ &\times \prod_{k=1}^K (\boldsymbol{\xi}^2)^{\nu_0/2} (\sigma_k^2)^{-(\nu_0/2+1)} \exp \left\{ -\frac{\nu_0 \boldsymbol{\xi}^2/2}{\sigma_k^2} \right\} \times (\boldsymbol{\xi}^2)^{a_0-1} \exp \{-b_0 \boldsymbol{\xi}^2\}, \end{aligned}$$

where $[x = i]$ is the Iverson bracket. Such a posterior distribution is quite reminiscent of the one that we derived for the hierarchical model in Section 3, but this time it includes a portion related with the clustering process, and also, group-specific parameters cycle over K terms instead of m .

Once again, we derive the fcd's from the posterior distribution, obtaining that:

- The fcs of γ_j , for $j = 1, \dots, m$, is a Categorical distribution such that

$$\Pr(\gamma_j = k \mid \text{rest}) \propto \omega_k \prod_{i=1}^n \mathbf{N}(y_{i,j} \mid \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_k, \sigma_k^2), \quad \text{for } k = 1, \dots, K.$$

- The fcs of $\boldsymbol{\omega}$ is

$$\boldsymbol{\omega} \mid \text{rest} \sim \text{Dir}(\alpha_{01} + n_1, \dots, \alpha_{0K} + n_K)$$

where $n_k = \#\{j : \gamma_j = k\}$ is the number of elements in cluster k , for $k = 1, \dots, K$.

- The fcd of β_k , for $k = 1, \dots, K$, is

$$\beta_k | \text{rest} \sim N_p \left(\left(\Sigma^{-1} + \sigma_k^{-2} \mathbf{X}_{(k)}^T \mathbf{X}_{(k)} \right)^{-1} \left(\Sigma^{-1} \beta + \sigma_k^{-2} \mathbf{X}_{(k)}^T \mathbf{y}_{(k)} \right), \left(\Sigma^{-1} + \sigma_k^{-2} \mathbf{X}_{(k)}^T \mathbf{X}_{(k)} \right)^{-1} \right)$$

where $\mathbf{X}_{(k)} = [\mathbf{X}_j^T : \gamma_j = k]^T$ and $\mathbf{y}_{(k)} = [\mathbf{y}_j^T : \gamma_j = k]^T$. Note that if cluster k is empty, then the fcd of β_k is just $\beta_k | \text{rest} \sim N_p(\beta, \Sigma)$.

- The fcd for β is

$$\beta | \text{rest} \sim N_p \left(\left(\Lambda_0^{-1} + K^* \Sigma^{-1} \right)^{-1} \left(\Lambda_0^{-1} \mu_0 + \Sigma^{-1} \sum_{k:n_k > 0}^K \beta_k \right), \left(\Lambda_0^{-1} + K^* \Sigma^{-1} \right)^{-1} \right)$$

where K^* is the number of non-empty clusters.

- The fcd of Σ is

$$\Sigma | \text{rest} \sim IW \left(n_0 + K^*, \left(\mathbf{S}_0 + \sum_{k:n_k > 0}^K (\beta_k - \beta)(\beta_k - \beta)^T \right)^{-1} \right).$$

- The fcd of σ_k^2 , for $k = 1, \dots, K$, is

$$\sigma_k^2 | \text{rest} \sim \text{IG} \left((\nu_0 + n_k)/2, \left(\nu_0 \xi^2 + \sum_{j:\gamma_j=k}^m \sum_{i=1}^n (y_{i,j} - \mathbf{x}_{i,j}^T \beta_{\gamma_j})^2 \right) / 2 \right).$$

Again, note that if cluster k is empty, then the fcd of σ_k^2 is just $\sigma_k^2 | \text{rest} \sim \text{IG}(\nu_0/2, \nu_0 \xi^2/2)$.

- The fcd of ξ^2 is

$$\xi^2 | \text{rest} \sim G \left(a_0 + K^* \nu_0/2, b_0 + \frac{\nu_0}{2} \sum_{k:n_k > 0}^K \sigma_k^{-2} \right).$$

Thus, the Gibbs sampling algorithm in this case is as follows:

1. For a given value of K , choose an initial configuration for each parameter in the model, say $\gamma_1^{(0)}, \dots, \gamma_m^{(0)}$, $\omega^{(0)}$, $\beta_1^{(0)}, \dots, \beta_K^{(0)}$, $\beta^{(0)}$, $\Sigma^{(0)}$, $(\sigma_1^2)^{(0)}, \dots, (\sigma_K^2)^{(0)}$, and $(\xi^2)^{(0)}$.
2. Update $\gamma_1^{(b-1)}, \dots, \gamma_m^{(b-1)}$, $\omega^{(b-1)}$, $\beta_1^{(b-1)}, \dots, \beta_K^{(b-1)}$, $\beta^{(b-1)}$, $\Sigma^{(b-1)}$, $(\sigma_1^2)^{(b-1)}, \dots, (\sigma_K^2)^{(b-1)}$, and $(\xi^2)^{(b-1)}$ until convergence:
 - a) Sample $\gamma_j^{(b)}$ from the fcd $p(\gamma_j | \omega^{(b-1)}, \{\beta_k^{(b-1)}\}, \{\sigma_k^2\}^{(b-1)}, \mathbf{X}_j)$, for $j = 1, \dots, m$.
 - b) Sample $\omega^{(b)}$ from the fcd $p(\omega | \{\gamma_j^{(b)}\})$.
 - c) Sample $\beta_k^{(b)}$ from the fcd $p(\beta_k | \{\gamma_j^{(b)}\}, \{\sigma_k^2\}^{(b-1)}, \beta^{(b-1)}, \Sigma^{(b-1)}, \{\mathbf{y}_j\}, \{\mathbf{X}_j\})$, for $k = 1, \dots, K$.
 - d) Sample $\beta^{(b)}$ from the fcd $p(\beta | \{\beta_k^{(b)}\}, \Sigma^{(b-1)})$.
 - e) Sample $\Sigma^{(b)}$ from the fcd $p(\Sigma | \{\beta_k^{(b)}\}, \beta^{(b)})$.
 - f) Sample $(\sigma_k^2)^{(b)}$ from the fcd $p(\sigma_k^2 | \{\gamma_j^{(b)}\}, \{\beta_k^{(b)}\}, (\xi^2)^{(b-1)}, \{\mathbf{y}_j\}, \{\mathbf{X}_j\})$, for $k = 1, \dots, K$.
 - g) Sample $(\xi^2)^{(b)}$ from the fcd $p(\xi^2 | \{(\sigma_k^2)^{(b)}\})$.
3. Cycle until achieve convergence.

5. Further Extensions

In order to gain stochastic flexibility, yet another way of extending the model provided in Section 3 consists in completely relaxing the normality assumption about the response variable by assigning a prior distribution to it. Specifically, we consider a Dirichlet process (DP) mixture model of the form

$$y_{i,j} \mid \mathbf{x}_{i,j}, \sigma_j^2, G \stackrel{\text{ind}}{\sim} \int \mathbf{N}(y_{i,j} \mid \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_j, \sigma_j^2) dG(\boldsymbol{\beta}_j) \quad \text{and} \quad G \sim \text{DP}(\alpha, H)$$

where α is a positive scalar parameter and H is a base distribution function. In this case, the DP generates cumulative distribution functions on \mathbb{R} (see for example Müller et al. (2015) for a formal treatment of the DP). The model given above can be also written as

$$y_{i,j} \mid \mathbf{x}_{i,j}, \boldsymbol{\beta}_j, \sigma_j^2 \stackrel{\text{ind}}{\sim} \mathbf{N}(\mathbf{x}_{i,j}^\top \boldsymbol{\beta}_j, \sigma_j^2), \quad \boldsymbol{\beta}_j \mid G \sim G \quad \text{and} \quad G \sim \text{DP}(\alpha, H),$$


which makes evident why the $\mu_{i,j} = \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_j$ can be interpreted as subject-specific random effects. Such an extension is beyond the scope of this work and will be discussed in detail elsewhere.

Other straightforward parametric extensions are considering group-specific effects $\theta_1, \dots, \theta_m$ in a way that $\mathbf{E}(y_{i,j} \mid \mathbf{x}_{i,j}, \boldsymbol{\beta}_j, \theta_j) = \mathbf{x}_{i,j}^\top \boldsymbol{\beta}_j + \theta_j$, for $i = 1, \dots, n$, $j = 1, \dots, m$, as well as extra model hierarchies such as letting ν_0 to be a integer random value ranging from 1 to a fixed large upper bound in a way that $p(\nu) \propto e^{-\kappa_0 \nu}$, where κ_0 is a hyperparameter. For clustering tasks, more sophisticated extensions require the specification of nonparametric priors of the form $\gamma_j \mid \{\omega_k\} \stackrel{\text{iid}}{\sim} \sum_{k=1}^{\infty} \omega_k \delta_k$, for $j = 1, \dots, m$, where $\omega_k = u_k \prod_{h < k} (1 - u_h)$ are weights constructed from a sequence u_1, u_2, \dots , with $u_k \stackrel{\text{iid}}{\sim} \text{Beta}(1 - a, b + ka)$ for $0 < a < 1$ and $b > -a$. The joint distribution of the set of weights $\omega_1, \omega_2, \dots$ is called a stick-breaking prior with parameters a and b . This formulation is connected to the stick-breaking construction of the Poisson-Dirichlet process (Pitman & Yor, 1997). The stick-breaking representation associated with the Dirichlet process is a special case with $a = 0$.


6. Computation

We implement the models provided in Sections 2, 3, and 4 following the corresponding algorithms provided in each section (our code is available for those readers that explicitly ask it from the corresponding author). Every time our results are based on $B = 50,000$ samples of the posterior distribution obtained after thinning the original chains every 10 observations and a burn-in period of 10,000 iterations. In addition, before using the MCMC samples with inferential purposes, we determine first if there is any evidence of lack of convergence of any chain to its stationary distribution. Following standard practices, we produce log-likelihood traceplots of each model. Fitting the models to the data provided in Section 8, such plots strongly suggest that there are no stationary concerns


since the log-likelihoods move in a consistent direction (see Figure 2). Further, autocorrelation plots of model parameters (not shown here) indicate that there are no signs of strong dependence in the chains. Thus, we are very confident of our MCMC samples to perform inductive tasks.



(a) LRM



(b) HLRM



(c) CHLRM

FIGURE 2: Log-likelihood traceplots when fitting the models for the plant size data analyzed in Section 8.

7. Model Checking and Goodness-of-Fit

We check the in-sample performance of the model by generating new data from the posterior predictive distribution, and then, calculating a battery of test statistics (such as the mean), aiming to compare the corresponding empirical distributions with the actual values observed in the sample (Gelman et al., 2013). In this spirit, for each quantity of interest, we also compute the posterior predictive p -value (ppp), which can be calculated as

$$ppp = \Pr(t(\mathbf{y}^{\text{rep}}) > t(\mathbf{y}) \mid \mathbf{y})$$

where \mathbf{y}^{rep} is a predictive dataset and t a test statistic, in order to measure how good the model is in fitting the actual sample.

Finally, in order to assess the goodness-of-fit of each model as a measure of their predictive performance, in what follows we consider two metrics that account for both model fit and model complexity. The goal here is not necessarily picking the model with lowest estimated prediction error but to determine if improvements in fitting the model are large enough to justify the additional difficulty. That why such measures also serve as model-selection tools. The model-based literature has largely focused on the Bayesian Information Criteria (BIC) as a mechanism for model selection. However, the BIC is typically inappropriate for hierarchical models since the hierarchical structure implies that the effective number of parameters will typically be lower than the actual number of parameters in the likelihood (Gelman et al., 2014). Two popular alternatives to the BIC that address such an issue are the Deviance Information Criterion (DIC) (Spiegelhalter et al., 2002, 2014),

$$\text{DIC} = -2 \log p(\mathbf{y} \mid \hat{\Theta}) + 2p_{\text{DIC}},$$

with $p_{\text{DIC}} = 2 \log p(\mathbf{y} \mid \hat{\Theta}) - 2\text{E}(\log p(\mathbf{y} \mid \Theta))$ or $p_{\text{DIC}} = \text{Var}(\log p(\mathbf{y} \mid \Theta))$, and the Watanabe-Akaike Information Criterion (WAIC) (Watanabe, 2010, 2013),

$$\text{WAIC} = -2 \sum_{i,j} \log \text{E}(p(y_{i,j} \mid \Theta)) + 2p_{\text{WAIC}},$$

with $p_{\text{WAIC}} = 2 \sum_{i,j} \{ \log \text{E}(p(y_{i,j} \mid \Theta)) - \text{E}(\log p(y_{i,j} \mid \Theta)) \}$, where $\hat{\Theta}$ is the posterior mean of model parameters, and p_{DIC} and p_{WAIC} are penalty terms accounting for model complexity. Note that in the previous expressions all expectations, which are computed with respect to the posterior distribution, can be approximated by averaging over Markov chain Monte Carlo (MCMC) samples (see Section 6 for details). Next section we use the two versions of the DIC presented here.

Additional out-of-sample goodness-of-fit assessments can be carried out through a series of cross-validation experiments (randomly selected subsets of roughly equal size in the dataset are treated as missing and then predicted using the rest of the data) on several datasets exhibiting different kinds of grouping factors as well as samples sizes. We refer the reader to Gelman et al. (2013, Chap. 7) for such protocol.

8. Illustration

Nitrogen is an essential macro-nutrient needed by all plants to thrive. It is an important component of many structural, genetic, and metabolic compounds in plant cells. Increasing the levels of nitrogen during the vegetative stage can strengthen and support the plant roots, enabling them to take in more water nutrients. This allows a plant to grow more rapidly and produce large amounts of succulent, green foliage, which in turn can generate bigger yields, tastier vegetables, and a crop that is more resistant to pests, diseases, and other adverse conditions. Using too much nitrogen, however, can be just as harmful to plants as to little. A researcher took $n = 5$ measurements of nitrogen soil concentration (x) and plant sizes (y) within each of $m = 24$ farms. Thus, we have that $y_{i,j}$ and $x_{i,j}$ are the plant size and the nitrogen soil concentration values, respectively, associated the i -th plant from the j -th farm, $i = 1, \dots, 5$ and $j = 1, \dots, 24$. This dataset is given in [Crawley \(2012, p. 704\)](https://github.com/shifteight/R-lang/blob/master/TRB/data/farms.txt) and remains publicly available at <https://github.com/shifteight/R-lang/blob/master/TRB/data/farms.txt>.

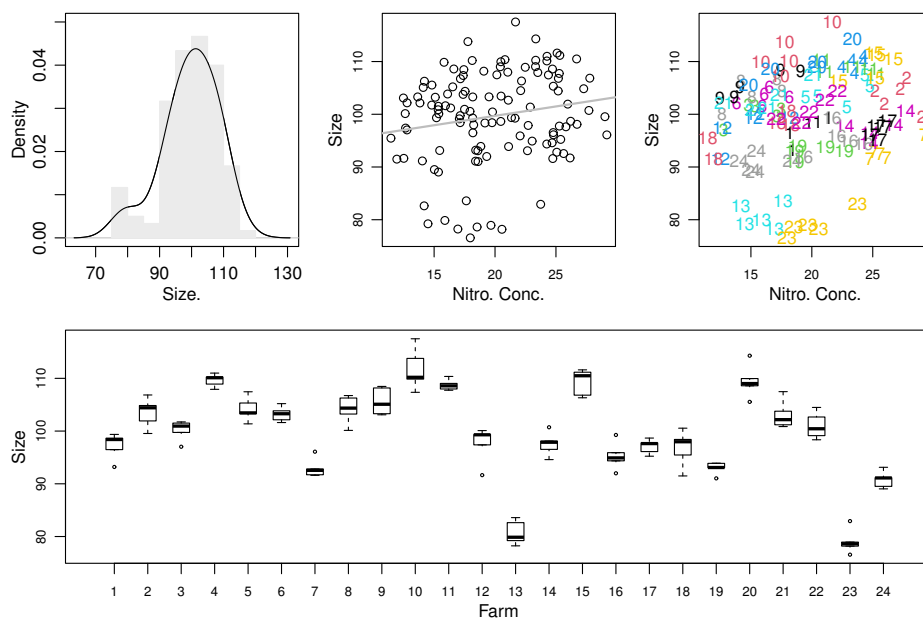


FIGURE 3: Descriptive plots. The second panel exhibit the ordinary least squares regression line for this data. Colors in the third panel correspond to different farms.

8.1. Exploratory Data Analysis

A histogram of the plant size is shown in the first panel of Figure 3. The plant size ranges from 76.56 to 117.50 which seems quite large in comparison with the plant size range within each farm (see the bottom panel). The second and

third panels show the relationship between the plant size and the nitrogen soil concentration. In particular, the second panel exhibits the ordinary least squares (OLS) regression line for these data, which is given by $y = 92.57 + 0.36x$ with a standard error of $\hat{\sigma}_{\text{OLS}} = 8.46$ with 118 degrees of freedom, and an adjusted R-squared of $R_{\text{adj}}^2 = 2.56\%$. The third panel presents the same plot but taking into account the farm where the measurements belong to. These panels along with the OLS fit indicate two main features about this experiment: (1) there is an important relationship between the nitrogen soil concentration and the plant size (in our OLS fit the slope turns out to be significant, p -value = 0.04); and (2) there is a clear farm effect on plant size since colors in the scatter plot reveal clustering patterns, and also, the corresponding boxplots strongly suggest differences in terms of mean plant growth among farms.

8.2. Fitting a LRM

We fit the LRM given in Section 2 to these data without taking into account the farm information. Posterior summaries of the model parameters are provided in Table 1. Even though the posterior mean of the model parameters practically coincide with their corresponding OLS estimates, these results are again quite limited since they do not allow us to isolate any kind of effect over the plant size arising from the grouping factor. Such limitation strongly motivates the hierarchical approaches that we present in this paper.

TABLE 1: Posterior summaries of the model parameters in the linear regression model.

Parameter	Mean	SD	Q2.5%	Q97.5%
β_1	92.59	3.60	85.53	99.61
β_2	0.36	0.18	0.01	0.70
σ^2	72.89	9.67	56.35	94.07

8.3. Fitting a HLRM

Now we go further and fit the HLRM given in Section 3 to the plant size data considering both group-specific fixed effects β_1, \dots, β_m and group-specific variances $\sigma_1^2, \dots, \sigma_m^2$. Such a group-specific approach is very convenient because it allows us to carry out separate-group inferences as opposed to its non-hierarchical counterpart.

We present our main results in Figures 4 and 5 where we display 95% quantile-based credible intervals for β_1, \dots, β_m and $\sigma_1^2, \dots, \sigma_m^2$, respectively. At this point we are capable of making evident some important findings. The uncertainty about the group-specific parameters is not constant since the amplitude of the credible intervals clearly varies across farms. This effect is particularly evident for the variance components. In addition, point estimates (posterior means) of the group-specific parameters are also quite variable, which strongly suggests that for these data considering this approach is beneficial because it allows us to characterize

for each farm its own unique features. However, some farms show some signs of similar features, which was also evident before in Figure 3. We explore clustering patterns in the next subsection.

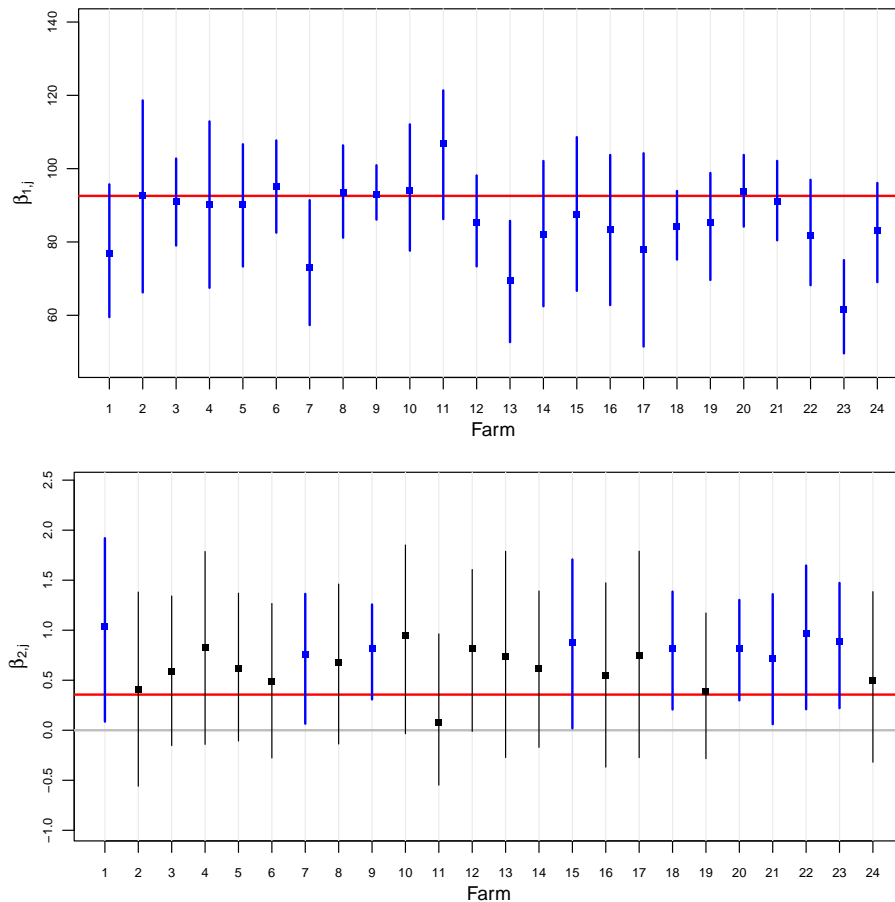


FIGURE 4: 95% quantile-based credible intervals and posterior means (squares) for the regressor fixed-effects regressor parameters in the hierarchical Normal linear regression model. Colored thicker lines correspond to credible intervals that do not contain zero. Top panel: $\beta_{1,1}, \dots, \beta_{1,m}$ (intercepts). Bottom panel: $\beta_{2,1}, \dots, \beta_{2,m}$ (slopes). OLS estimates are depicted through a red horizontal line.

As expected, we see that all the intercepts are statistically significant, but also highly variable (ranging from 61.48 to 106.81). On the other hand, the story behind the slopes (ranging from 0.08 to 1.04) is quite different. We see that just 9 out of 24 (37.5%) of such parameters turn out to be significant, namely, for farms 1, 7, 9, 15, 18, 20, 21, 22, and 23 (Figure 6 shows the corresponding estimated regression lines for these farms). The previous fact strongly suggests that the relevance of the relationship between the nitrogen soil concentration and the plant

size is not consistent across farms. Therefore, in this case, considering farms as a grouping factor makes a substantial impact on the analysis.




FIGURE 5: 95% quantile-based credible intervals and posterior means (squares) for the variance parameters $\sigma_1^2, \dots, \sigma_m^2$ in the hierarchical Normal linear regression model.

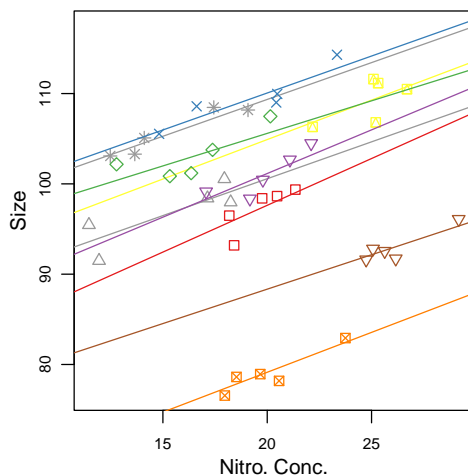


FIGURE 6: Estimated lines of the hierarchical Normal linear regression model for those farms whose nitrogen soil concentration is statistically significant. Colors are used to represent different farms.

8.4. Fitting a CHLRM

It is quite reasonable attempting to identify clusters composed of farms given the abundant evidence of similarities among groups and cluster formation detected

in both the exploratory data analysis and the hierarchical modeling stage of the analysis. Here, we fit the CHLRM model given in Section 4 in order to provide a formal partition of farms. To this end, we fit the model using $K = m$ as a default number of communities. Such an extreme case represents the prior belief that there are no clustering patterns at all. A moderate large value of K is convenient in situations like this because it allows the data to self-select how many non-empty clusters should be considered. We will see in what follows that out from the m clusters, many turn out to be empty and only a few remain.

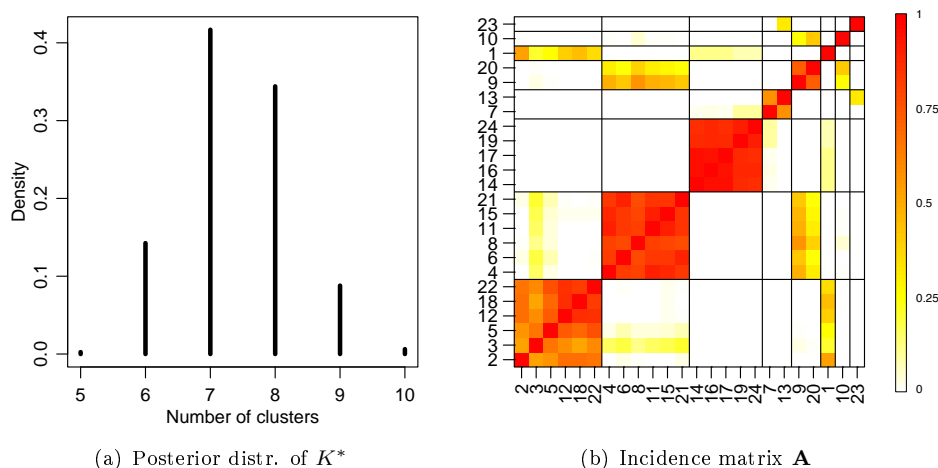


FIGURE 7: Posterior inference on the cluster assignments $\gamma_1, \dots, \gamma_m$.

Let K^* be the number of non-empty clusters in the partition induced by $\gamma_1, \dots, \gamma_m$. The left panel of Figure 7 shows the posterior distribution of K^* . We see that the most highly probable values are $K^* = 7$ and $K^* = 8$. In fact, $\Pr(K^* \in \{7, 8\} \mid \mathbf{y}, \mathbf{X}) = 0.76$, which means that around three quarters of the posterior partitions are composed of either 7 or 8 clusters of farms. The estimated number of non-empty clusters in the data is $K^* = 7$ (maximum a posteriori with a reference value very close to 0.4).

On the other hand, the right panel in Figure 7 shows the $m \times m$ incidence matrix $\mathbf{A} = [a_{j,j'}]$ obtained from the posterior distribution of the cluster assignments $\gamma_1, \dots, \gamma_m$. The incidence matrix is a pairwise-probability matrix whose elements are given by $a_{j,j'} = \Pr(\gamma_j = \gamma_{j'} \mid \mathbf{y}, \mathbf{X})$, for $j, j' = 1, \dots, m$ (note that $a_{j,j} = 1$ for all j). Thus, $a_{j,j'}$ simply represents the posterior probability that farms j and j' belong to the same community. Such probabilities are indeed identifiable, however labels themselves are not since the likelihood is invariant to relabelling of the mixture components (this is known as the label switching problem; see Stephens (2000) and references therein). On top the incidence matrix, we also present a point estimate of the partition induced by such a matrix (represented by black lines), which can be obtained by employing the clustering methodology proposed in Lau & Green (2007) with a relative error cost of 0.5. As expected, we see that

eight clusters are formed from the data. Labels in both rows and columns are deliberately arranged in a way that the corresponding partition can be visualize easily. The corresponding cluster sizes are 6, 6, 5, 2, 2, 1, 1, and 1. Specifically, the clusters are composed of the following farms: $C_1 = \{2, 3, 5, 12, 18, 22\}$, $C_2 = \{4, 6, 8, 11, 15, 21\}$, and $C_3 = \{14, 16, 17, 19, 24\}$, $C_4 = \{7, 13\}$, $C_5 = \{9, 20\}$, $C_6 = \{1\}$, $C_7 = \{10\}$, and $C_8 = \{23\}$. These clusters along with their estimated regression lines are also represented in Figure 6. We see that the estimated clusters make sense spatially according to the data points. The slope parameters turn out to be significant for every cluster.

8.5. Model Checking and Goodness-of-Fit

First, we evaluate in-sample predictive performance of each model by means of the mean square error (MSE) of replicated data as well as the posteriori predictive p -values (ppp's) associated with the predictive distribution of a set of test statistics (mean, median, interquartile range, and the standard deviation), both locally and globally (i.e., with and without considering the farm as grouping factor, respectively). Furthermore, as discussed in Section 7, we also consider the deviance information criterion (DIC) in order to assess the overall goodness-of-fit of the models.

In what follows we examine the performance of all the fitted models, namely, the linear regression model (LRM), the hierarchical linear regression model (HLRM), and the clustering hierarchical linear regression model (CHLRM). Our main findings at a global level are presented in Table 2. As expected, the performance of LRM at predicting new data is the worst. Interestingly, both HLRM and CHLRM practically have the same behavior in this regard. A similar result is encountered again in terms of model fit, but this time both versions of the DIC favor HLRM over CHLRM. These results strongly suggest that considering a hierarchical structure when building a model clearly favor both in-sample predictive performance as well as goodness-of-fit.

TABLE 2: Global measures associated with the mean square error, posterior predictive p -values, and deviance information criterion for all the models fitted to the plant size data. LRM: Normal linear regression model. HLRM: Hierarchical Normal linear regression model. CHLRM: Clustering hierarchical Linear regression model.

Model	MSE	p_{DIC1}	DIC1	p_{DIC2}	DIC2
LRM	144.528	2.972	857.158	0.792	854.381
HLRM	10.206	30.469	526.450	22.804	511.122
CHLRM	10.068	32.167	555.020	18.929	528.545

At global level, all the models seem to predict adequately the test statistics since there are no evidence of extreme ppp values close to either 0 or 1 (table not shown here). However, the story locally is quite different. Figure 8 show boxplots summarizing the ppp distribution at a local level (i.e., within each farm). As opposed to its early global conduct, LRM misbehaves and fails at capturing the the

test statistics because the ppp's are too extreme. On the other hand, both HLRM and CHLRM fit the data properly since the ppp distribution is approximately centered around 0.5. However, the ppp's for HLRM are less spread than those for CHLRM, which indicates that HLRM tends to have a mild improvement in predictive performance in this case.

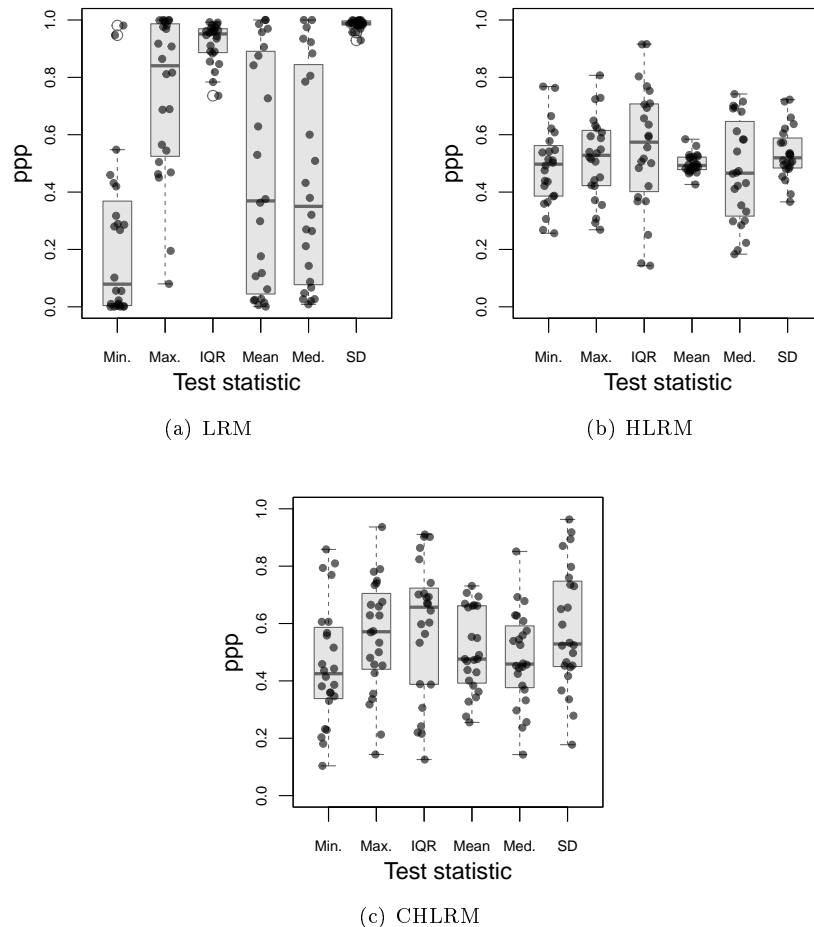


FIGURE 8: Local ppp's for a battery of test statistics.

9. Discussion

Hierarchical models provide a strong alternative to analyze complex and realistic settings. Their parameter flexibility allow us to describe many characteristics of a given dataset that a regular single-level model does not provide. The ability to model within and between means and variances yields to better

knowledge of the problem (even if we want to predict future values at any stage). For even more detail and deep types of complexity, this kind of models provide many alternatives for generalizations as discussed in Section 5. In addition, the information criteria (e.g., DIC) discussed in Section 7 was proven to be successful to be a powerful as a model selection tool.

If additional time-dependent covariates were available, we would be able to extend the linear dependency in the model by letting $E(y(t) | \mathbf{x}, \boldsymbol{\beta}(t)) = \mathbf{x}^\top \boldsymbol{\beta}(t)$, where $\boldsymbol{\beta}(t) = (\beta_1(t), \dots, \beta_d(t))$ is a vector of arbitrary real smooth functions called dynamic parameters. This model plays a fundamental role identifying and characterizing dynamic tendencies and patterns over time in many scientific areas, such as biology, epidemiology and medical science, among others (Sosa & Buitrago, 2021).

Finally, we recommend consider variational methods as alternative approaches for parameter estimation in order to consider larger datasets. Such techniques currently constitute an active research area in computational statistics (Blei et al., 2017).

[Received: October 2021 — Accepted: February 2022]

References

- Albert, J. (2009), *Bayesian Computation with R*, Use R!, Springer. http://books.google.com/books?id=AALhk_mt7SYC
- Banerjee, S., Carlin, B. & Gelfand, A. (2014), *Hierarchical modeling and analysis for spatial data*, CRC press.
- Blei, D., Kucukelbir, A. & McAuliffe, J. (2017), ‘Variational inference: A review for statisticians’, *Journal of the American statistical Association* **112**(518), 859–877.
- Christensen, R., Johnson, W., Branscum, A. & Hanson, T. (2011), *Bayesian Ideas and Data Analysis: An Introduction for Scientists and Statisticians*, Chapman & Hall/CRC Texts in Statistical Science, Taylor & Francis. <http://books.google.com/books?id=qPERhCbePNcC>
- Crawley, M. J. (2012), *The R book*, John Wiley & Sons.
- Dey, D., Ghosh, S. & Mallick, B. (2000), *Generalized linear models: A Bayesian perspective*, CRC Press.
- Ferguson, T. (1973), ‘A bayesian analysis of some nonparametric problems’, *The annals of statistics* pp. 209–230.
- Gamerman, D. & Lopes, H. (2006), *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*, CRC Press.

- Gelman, A., Carlin, J., Stern, H., Dunson, D., Vehtari, A. & Rubin, D. (2013), *Bayesian Data Analysis, Third Edition*, Chapman & Hall/CRC Texts in Statistical Science, Taylor & Francis. <http://books.google.com/books?id=ZXL6AQAAQBAJ>
- Gelman, A., Hwang, J. & Vehtari, A. (2014), ‘Understanding predictive information criteria for bayesian models’, *Statistics and computing* **24**(6), 997–1016.
- Hoff, P. (2009), *A First Course in Bayesian Statistical Methods*, Springer Texts in Statistics, Springer. <http://books.google.com/books?id=V8jT2SimGR0C>
- Ishwaran, H. & Zarepour, M. (2000), ‘Markov chain monte carlo in approximate dirichlet and beta two-parameter process hierarchical models’, *Biometrika* **87**(2), 371–390.
- Kass, R. & Wasserman, L. (1995), ‘A reference bayesian test for nested hypotheses and its relationship to the schwarz criterion’, *Journal of the american statistical association* **90**(431), 928–934.
- Kolaczyk, E. & Csárdi, G. (2020), *Statistical analysis of network data with R*, Vol. 65, Springer.
- Lau, J. & Green, P. (2007), ‘Bayesian model-based clustering procedures’, *Journal of Computational and Graphical Statistics* **16**(3), 526–558.
- Müller, P., Quintana, F. A., Jara, A. & Hanson, T. (2015), *Bayesian nonparametric data analysis*, Springer.
- Pitman, J. & Yor, M. (1997), ‘The two-parameter poisson-dirichlet distribution derived from a stable subordinator’, *The Annals of Probability* pp. 855–900.
- Sethuraman, J. (1994), ‘A constructive definition of dirichlet priors’, *Statistica sinica* pp. 639–650.
- Sosa, J. & Buitrago, L. (2021), ‘Time-varying coefficient model estimation through radial basis functions’, *Journal of Applied Statistics* pp. 1–25.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & Linde, A. (2014), ‘The deviance information criterion: 12 years on’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **76**(3), 485–493.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P. & Van Der Linde, A. (2002), ‘Bayesian measures of model complexity and fit’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **64**(4), 583–639.
- Stephens, M. (2000), ‘Dealing with label switching in mixture models’, *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **62**(4), 795–809.
- Wakefield, J. (2013), *Bayesian and frequentist regression methods*, Springer Science & Business Media.

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

Watanabe, S. (2013), ‘A widely applicable Bayesian information criterion’, *Journal of Machine Learning Research* **14**(Mar), 867–897.

Appendix A. Dirichlet process

A random distribution function F is generated from a Dirichlet Process with parameters $\alpha > 0$ and G a distribution function on \mathbb{R} , denoted by $F \sim \text{DP}(\alpha, G)$, if for any finite measurable partition B_1, \dots, B_k of \mathbb{R} ,

$$(F(B_1), \dots, F(B_k)) \sim \text{Dir}(\alpha G(B_1), \dots, \alpha G(B_k)).$$

G plays the role of the center of the DP (also referred to as base probability measure, or base distribution), where as α can be viewed as a precision parameter (the larger α is, the closer we expect a realization F from the process to be to G). See [Ferguson \(1973\)](#) for the role of G on more technical properties of the DP.

Alternatively, the constructive definition of the DP ([Sethuraman, 1994](#)) states that $F \sim \text{DP}(\alpha, G)$ if F is (almost surely) of the form

$$F(\cdot) = \sum_{k=1}^{\infty} \omega_k \delta_{\vartheta_k}(\cdot)$$

where $\delta_{\vartheta}(\cdot)$ denotes a point mass at ϑ (degenerate distribution putting probability one on ϑ), $\vartheta_k \stackrel{\text{i.i.d.}}{\sim} G$, $\omega_k = z_k \prod_{\ell=1}^{k-1} (1 - z_{\ell})$, $z_k \stackrel{\text{i.i.d.}}{\sim} \text{Beta}(1, \alpha)$, for $k = 1, 2, \dots$. Hence, the DP generates distributions that can be represented as countable mixtures of point masses (the locations ϑ_k arise i.i.d. from the base distribution G), whose weights ω_k arise through a stick-breaking construction (it can be shown that $\sum_{k=1}^{\infty} \omega_k = 1$ almost surely).

Based on its constructive definition, it is evident that the DP generates (almost surely) discrete distributions on \mathbb{R} .

Appendix B. Notation

The cardinality of a set A is denoted by $|A|$. If P is a logical proposition, then $1_P = 1$ if P is true, and $1_P = 0$ if P is false. $\lfloor x \rfloor$ denotes the floor of x , whereas $[n]$ denotes the set of all integers from 1 to n , i.e., $\{1, \dots, n\}$. The Gamma function is given by $\Gamma(x) = \int_0^{\infty} u^{x-1} e^{-u} du$.

Matrices and vectors with entries consisting of subscripted variables are denoted by a boldfaced version of the letter for that variable. For example, $\mathbf{x} = (x_1, \dots, x_n)$ denotes an $n \times 1$ column vector with entries x_1, \dots, x_n . We use $\mathbf{0}$ and $\mathbf{1}$ to denote the column vector with all entries equal to 0 and 1, respectively,

and \mathbf{I} to denote the identity matrix. A subindex in this context refers to the corresponding dimension; for instance, \mathbf{I}_n denotes the $n \times n$ identity matrix. The transpose of a vector \mathbf{x} is denoted by \mathbf{x}^\top ; analogously for matrices. Moreover, if \mathbf{X} is a square matrix, we use $\text{tr}(\mathbf{X})$ to denote its trace and \mathbf{X}^{-1} to denote its inverse. The norm of \mathbf{x} , given by $\sqrt{\mathbf{x}^\top \mathbf{x}}$, is denoted by $\|\mathbf{x}\|$.

Now, we present the form of some standard probability distributions:

- Multivariate normal:

A $d \times 1$ random vector $\mathbf{X} = (X_1, \dots, X_d)$ has a multivariate Normal distribution with parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$, denoted by $\mathbf{X} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma} \sim \mathcal{N}_d(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, if its density function is

$$p(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}) = (2\pi)^{-d/2} |\boldsymbol{\Sigma}|^{-1/2} \exp \left\{ -\frac{1}{2} (\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \right\}.$$

- Inverse Wishart:

A $d \times d$ random matrix \mathbf{W} has a Inverse Wishart distribution with parameters ν y \mathbf{S}^{-1} , i.e., $\mathbf{W} \sim \text{WI}(\nu, \mathbf{S}^{-1})$, if its density function is

$$p(\mathbf{W}) \propto |\mathbf{W}|^{-(\nu+d+1)/2} \exp \left\{ -\frac{1}{2} \text{tr}(\mathbf{S}\mathbf{W}^{-1}) \right\}, \quad \nu > 0, \quad \mathbf{S} > 0.$$

- Gamma:

A random variable X has a Gamma distribution with parameters $\alpha, \beta > 0$, denoted by $X \mid \alpha, \beta \sim \mathcal{G}(\alpha, \beta)$, if its density function is

$$p(x \mid \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} \exp \{-\beta x\}, \quad x > 0.$$

- Inverse Gamma:

A random variable X has an Inverse Gamma distribution with parameters $\alpha, \beta > 0$, denoted by $X \mid \alpha, \beta \sim \text{IG}(\alpha, \beta)$, if its density function is

$$p(x \mid \alpha, \beta) = \frac{\beta^\alpha}{\Gamma(\alpha)} x^{-(\alpha+1)} \exp \{-\beta/x\}, \quad x > 0.$$

- Beta:

A random variable X has a Beta distribution with parameters $\alpha, \beta > 0$, denoted by $X \mid \alpha, \beta \sim \text{Beta}(\alpha, \beta)$, if its density function is

$$p(x \mid \alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} x^{\alpha-1} (1-x)^{\beta-1}, \quad 0 < x < 1.$$

- Dirichlet:

A $K \times 1$ random vector $\mathbf{X} = (X_1, \dots, X_K)$ has a dirichlet distribution with parameter vector $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)$, where each $\alpha_k > 0$, denoted by $\mathbf{X} \mid \boldsymbol{\alpha} \sim \text{Dir}(\boldsymbol{\alpha})$, if its density function is

$$p(\mathbf{x} \mid \boldsymbol{\alpha}) = \begin{cases} \frac{\Gamma(\sum_{k=1}^K \alpha_k)}{\prod_{k=1}^K \Gamma(\alpha_k)} \prod_{k=1}^K x_k^{\alpha_k-1}, & \text{if } \sum_{k=1}^K x_k = 1; \\ 0, & \text{otherwise.} \end{cases}$$

- Categorical:

A $K \times 1$ random vector $\mathbf{X} = (X_1, \dots, X_K)$ has a categorical distribution with parameter vector $\mathbf{p} = (p_1, \dots, p_K)$, where $\sum_{k=1}^K p_k = 1$, denoted by $\mathbf{X} \mid \mathbf{p} \sim \text{Cat}(\mathbf{p})$, if its probability mass function is

$$p(x \mid \mathbf{p}) = \begin{cases} \prod_{k=1}^K p_k^{\mathbf{1}\{x=k\}}, & \text{if } \sum_{k=1}^K x_k = 1; \\ 0, & \text{otherwise.} \end{cases}$$