

Advances in Bayesian Modeling: Applications and Methods

Avances en modelamiento bayesiano: aplicaciones y métodos

YIFEI YAN^{1,a}, JUAN SOSA^{2,b}, CARLOS ALBERTO MARTÍNEZ^{3,c}

¹DEPARTMENT OF STATISTICS, BASKIN SCHOOL OF ENGINEERING, UNIVERSITY OF CALIFORNIA, SANTA CRUZ, UNITED STATES

²DEPARTAMENTO DE ESTADÍSTICA, FACULTAD DE CIENCIAS, UNIVERSIDAD NACIONAL DE COLOMBIA, BOGOTÁ, COLOMBIA

³DEPARTAMENTO DE PRODUCCIÓN ANIMAL, FACULTAD DE MEDICINA VETERINARIA Y ZOOTECNIA, UNIVERSIDAD NACIONAL DE COLOMBIA, BOGOTÁ, COLOMBIA

Abstract

This paper explores the versatility and depth of Bayesian modeling by presenting a comprehensive range of applications and methods, combining Markov chain Monte Carlo techniques and variational approximations. Covering topics such as hierarchical modeling, spatial modeling, higher-order Markov chains, and Bayesian nonparametrics, the study emphasizes practical implementations across diverse fields, including oceanography, climatology, epidemiology, and astronomy. The aim is to bridge theoretical underpinnings with real-world applications, illustrating the formulation of Bayesian models, elicitation of priors, computational strategies, and posterior and predictive analyses. By leveraging different computational methods, this paper provides insights into model fitting, goodness-of-fit evaluation, and predictive accuracy, addressing computational efficiency and methodological challenges across various datasets and domains.

Keywords: Bayesian inference; Dirichlet process mixtures; Hierarchical modeling; Markov chain Monte Carlo; Model evaluation; Predictive analysis; Spatial modeling; Variational approximation.

^aPh.D. E-mail: yifeiyan@soe.ucsc.edu

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

^cPh.D. E-mail: camartinezn@unal.edu.co

Resumen

Este artículo explora la versatilidad y profundidad del modelado Bayesiano mediante la presentación de un amplio conjunto de aplicaciones y métodos, combinando técnicas de cadenas de Markov de Monte Carlo y aproximaciones variacionales. Al abarcar temas como modelamiento jerárquico, modelamiento espacial, cadenas de Markov de orden superior y métodos Bayesianos no paramétricos, el estudio enfatiza implementaciones prácticas en campos diversos, incluyendo oceanografía, climatología, epidemiología y astronomía. El objetivo es tender un puente entre los fundamentos teóricos y las aplicaciones del mundo real, ilustrando la formulación de modelos Bayesianos, la elicitación de distribuciones previas, las estrategias computacionales y los análisis posteriores y predictivos. Al aprovechar diferentes métodos computacionales, este trabajo ofrece una perspectiva sobre el ajuste de modelos, la evaluación de la bondad de ajuste y la precisión predictiva, abordando la eficiencia computacional y los desafíos metodológicos en distintos conjuntos de datos y dominios.

Palabras clave: Análisis predictivo; Aproximación variacional; Cadenas de Markov de Monte Carlo; Evaluación de modelos; Inferencia bayesiana; Mezclas de procesos de Dirichlet; Modelamiento espacial; Modelamiento jerárquico.

1. Introduction

Bayesian inference has established itself as a cornerstone of modern statistical methodology, offering a framework for inference, prediction, and decision making. Its versatility lies in its ability to handle complex data structures, implicit uncertainty quantification (hence the colloquial quote “uncertainty is for free”), and model comparison within a unified probabilistic framework. The rapid advancements in computational power and algorithms have enabled Bayesian methods to scale beyond their traditional limitations, fostering their application across a wide array of scientific disciplines, including oceanography, climatology, epidemiology, astronomy, and financial analysis, among many others.

At the heart of Bayesian computation lie two dominant approaches: Markov chain Monte Carlo (MCMC; e.g., [Gamerman & Lopes, 2006](#)) methods and variational inference (e.g., [Blei et al., 2017](#)). MCMC methods have been extensively used due to their asymptotic properties when sampling from complex posterior distributions. Techniques like the Gibbs sampler and Metropolis-Hastings algorithm provide the foundation for a wide range of Bayesian analyses, including hierarchical and spatial models. However, their reliance on iterative sampling often results in high computational costs, making them less suitable for large-scale problems. In contrast, variational inference offers a deterministic alternative by framing posterior approximation as an optimization problem, trading some accuracy for significant gains in computational efficiency. This duality between MCMC and variational approaches has expanded the frontier of Bayesian methods, allowing researchers to design their computational strategies to the specific needs of their problems.

This paper explores Bayesian inference across a spectrum of applications and theoretical constructs, showcasing its adaptability to diverse data types and analytical challenges. The examples span hierarchical modeling for multi-level data structures, spatial modeling for geographic data, and higher-order Markov chains for complex sequential dependencies. Additionally, advanced topics such as Dirichlet process mixtures and Bayesian mixtures of Brownian motion are discussed to illustrate the ability of Bayesian methods to handle non-parametric and continuous-time processes. These examples highlight the dual role of Bayesian inference: As a theoretical tool for model development and as a practical solution for analyzing real-world data.

The applications in this work draw from datasets across various domains, each presenting unique challenges in data structure, computational requirements, and interpretability. For example, spatial modeling applied to climatology emphasizes the handling of geographic dependencies, while hierarchical models in oceanography address nested data structures. Bayesian methods are shown to provide robust solutions for model formulation, prior elicitation, computation, posterior analysis, predictive accuracy, and goodness-of-fit evaluation. Through this breadth of applications, the paper aims to bridge theoretical concepts with practical implementations, illustrating the versatility and depth of Bayesian methods.

The rest of the paper is structured as follows. Section 2 outlines the theoretical foundations of Bayesian inference and computational techniques, covering Markov chain Monte Carlo (MCMC) methods and variational approaches that underpin the subsequent methods–case study pairs. Sections 3 to 6 then develop and apply these ideas in a sequence of domain-specific examples: we first consider hierarchical and spatial models in environmental applications from oceanography and climatology, followed by dynamic formulations for epidemiological data; we then turn to nonparametric Dirichlet process mixture models in an astronomical case study on galaxy data, and to higher-order Markov chain constructions in a further applied example. Each section details the model formulation, computational strategies (comparing MCMC and variational inference implementations where appropriate), and empirical results on the corresponding datasets. Finally, the discussion section synthesizes the key findings, draws out practical lessons from the case studies, and highlights directions for future research.

2. A Brief Review on Bayesian Statistics

This section introduces key concepts in Bayesian inference, focusing on the interpretation of probability and its role in statistical modeling. It explores the distinction between frequentist and Bayesian perspectives, highlighting how Bayesian methods incorporate prior beliefs and update them with observed data. Additionally, it discusses the computational challenges that arise in Bayesian analysis and the numerical techniques, such as Monte Carlo methods, that enable practical implementation in complex problems.

2.1. Interpretations of Probability

A probability measure is a sigma-additive measure defined on a sigma algebra of subsets of the sample space, taking values in the interval $[0, 1]$. This axiomatic foundation has facilitated the development of a vast family of probability models applied across numerous fields. However, the interpretation of probability is crucial in the foundations of Bayesian inference. One perspective, known as the frequentist view, defines probability as the long-run relative frequency of an event, implying a conceptual framework where an experiment is repeated indefinitely under homogeneous conditions. In contrast, the subjective interpretation regards probability as a measure of personal belief or uncertainty. This perspective is particularly relevant in scenarios where repeated experimentation is infeasible, such as predicting the performance of a specific cyclist in an upcoming race.

The philosophical foundation of Bayesian inference has been shaped by seminal contributions from [Cox \(1946\)](#), [De Finetti \(1970\)](#), and [Savage \(1972\)](#). A comprehensive exposition of the formal framework underpinning rational decision-making under uncertainty is provided by [Bernardo & Smith \(2009, Chap. 2\)](#). The key principle of rationality in this context lies in ensuring internal consistency within the decision-making process. Their work demonstrates that Bayes' theorem provides a principled mechanism for updating beliefs about the state of nature as new evidence becomes available.

2.2. The trinity of Bayesian Inference: Notation and description of the problem

To infer a (potentially multidimensional) parameter θ from observed data \mathbf{x} , a Bayesian approach involves identifying plausible values for the components of θ while quantifying the associated uncertainty. This uncertainty is represented through a probability distribution known as the prior distribution, which is specified *before* any data is observed. The prior can incorporate existing knowledge and beliefs about the state of nature, in which case it is referred to as a “subjective” prior. Alternatively, it may encode minimal or no prior information about the parameter, in which case it is derived using general principles such as invariance under one-to-one transformations and is termed an “objective,” “non-subjective,” or “default” prior. In many cases, such priors are improper, meaning they do not integrate to one, but this does not necessarily hinder the validity of Bayesian inference. A thorough introduction to objective priors can be found in [Ghosh \(2011\)](#).

The density or probability mass function of the prior distribution is denoted as $p(\theta)$, with its associated parameters referred to as hyperparameters. In the Bayesian framework, all unknown model parameters are treated as random variables, making the specification of their probability distributions a fundamental step in model formulation. This process, known as prior elicitation, involves selecting a prior distribution that reflects existing knowledge, assumptions, or desired properties of θ before observing any data.

Once data containing information about the parameter is observed, the focus shifts to updating knowledge and beliefs about it. This information is captured by the likelihood function, denoted as $f(\mathbf{x} \mid \boldsymbol{\theta})$, $L(\boldsymbol{\theta})$, or $L(\mathbf{x}; \boldsymbol{\theta})$, depending on the notation used. The core principle of Bayesian inference establishes Bayes' rule as the optimal method for updating prior beliefs by combining the prior and the likelihood, resulting in the posterior distribution, which represents the updated distribution of the parameter given the observed data. The posterior density or mass function is denoted as $p(\boldsymbol{\theta} \mid \mathbf{x})$. Using this notation, Bayes' theorem is expressed as:

$$p(\boldsymbol{\theta} \mid \mathbf{x}) = \frac{p(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta})}{p(\mathbf{x})} \propto f(\mathbf{x} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta}),$$

where $p(\mathbf{x})$ is the marginal density or mass function of the data. According to [Bernardo & Smith \(2009, Chap. 1\)](#), although Bayes' theorem is a straightforward probability result that provides a principled way to learn about $\boldsymbol{\theta}$ from \mathbf{x} while accounting for uncertainty, it has been a subject of controversy due to differing interpretations and the scope of its inputs. Nevertheless, the Bayesian approach has been shown to be both theoretically sound and highly applicable to real-world problems, as explored in this paper.

De Finetti's representation theorem plays a fundamental role in Bayesian methods, as it establishes the existence of a prior distribution when the data are assumed to be exchangeable. This concept of exchangeability is particularly important because it naturally leads to key statistical modeling principles, such as independent random samples and parameter estimation ([Bernardo & Smith, 2009, Chap. 4](#)). The posterior distribution serves as the central object of inference in the Bayesian framework. It enables the estimation of parameters through point estimators, the construction of credible intervals that allow for probabilistic statements about $\boldsymbol{\theta}$, the testing of statistical hypotheses, the prediction of future observations, and the evaluation of model fit.

Large sample theory establishes that, under suitable regularity conditions, the posterior distribution exhibits asymptotic normality. Specifically, as the sample size increases, the posterior distribution can be approximated by a Normal distribution centered at the maximum likelihood estimator (MLE), with covariance given by the inverse of the observed Fisher information $\hat{\mathbf{I}}_n$. One of the most significant results in this context is the Bernstein-von Mises theorem, which is often regarded as a bridge between frequentist and Bayesian inference. This theorem implies that for large samples, the posterior mean closely approximates the MLE. More formally, under the regularity conditions of the Bernstein-von Mises theorem and assuming a prior with finite expectation, it can be shown that ([Ghosh et al., 2007](#)):

$$\sqrt{n}(\boldsymbol{\theta}_n^* - \hat{\boldsymbol{\theta}}_n^{\text{ML}}) \rightarrow 0, \quad \text{almost surely as } n \rightarrow \infty,$$

where $\boldsymbol{\theta}_n^*$ and $\hat{\boldsymbol{\theta}}_n^{\text{ML}}$ the posterior mean and the MLE of $\boldsymbol{\theta}$, respectively.

This result highlights the asymptotic equivalence of Bayesian and frequentist estimators. More precisely, in the Bayesian framework, we have:

$$\hat{\mathbf{I}}_n^{1/2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n^{\text{ML}}) \mid \mathbf{x} \rightarrow N_p(0, \mathbf{I}), \quad \text{almost surely,}$$

whereas in the frequentist setting:

$$\hat{\mathbf{I}}_n^{1/2}(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}_n^{\text{ML}}) \mid \boldsymbol{\theta} \rightarrow \mathbf{N}_p(0, \mathbf{I}), \quad \text{almost surely.}$$

These results illustrate that, under regularity conditions, Bayesian and frequentist inference converge asymptotically, reinforcing the robustness of Bayesian methods in large-sample scenarios.

2.3. Toy Examples and the Need for Numerical Methods

We present two toy examples to illustrate the computation of the posterior distribution and highlight the necessity of numerical methods in most real-life problems. Consider the case where $x_1, \dots, x_n \mid \theta \stackrel{\text{ind}}{\sim} \text{Bernoulli}(\theta)$ and the prior distribution is $\theta \sim \text{Beta}(\alpha, \beta)$. This setup yields a product Bernoulli likelihood and a Beta prior with hyperparameters $\alpha > 0$ and $\beta > 0$. Applying Bayes' rule, the posterior distribution follows:

$$\theta \mid \mathbf{x} \sim \text{Beta}(\alpha + \sum_{i=1}^n x_i, \beta + n - \sum_{i=1}^n x_i).$$

A commonly used objective prior that corresponds to both Jeffreys' prior (which is invariant under one-to-one transformations) and Bernardo's reference prior (which maximizes the Kullback-Leibler divergence between the prior and the posterior) is $\text{Beta}(1/2, 1/2)$. Another widely used prior is the Laplace prior, $\text{Beta}(1, 1)$, which corresponds to a uniform distribution on $(0, 1)$. The Haldane prior, defined by the improper distribution $\alpha = \beta = 0$, is also frequently considered. However, in this case, when all observed data points are either successes or failures, the resulting posterior distribution is improper. Therefore, when employing improper priors, it is essential to verify the properness of the posterior distribution to ensure valid inference. Now, suppose $x_1, \dots, x_n \mid \theta \stackrel{\text{ind}}{\sim} \mathbf{N}(\theta, \sigma^2)$ distribution, where σ^2 is known. Assuming a Normal prior $\theta \sim \mathbf{N}(\mu, \tau^2)$, the posterior distribution is given by

$$\theta \mid \mathbf{x} \sim \mathbf{N}\left(B\mu + (1 - B)\bar{x}, (1 - B)\frac{\sigma^2}{n}\right),$$

where \bar{x} is the sample mean and $B = \sigma^2/(n\tau^2 + \sigma^2)$.

In both examples, the posterior distribution is available in closed form, allowing for direct computation of moments and other relevant quantities. These posterior characteristics can be used to derive point estimates, construct credible intervals, and assess the dispersion of the parameter of interest. From a decision-theoretic perspective, the choice of a point estimator—whether the posterior mean, median, or mode—is justified by minimizing the expected loss under different loss functions (Ghosh et al., 2007). For many real-life problems, posterior moments and even the posterior distribution itself cannot be derived in closed form. This limitation restricted the practical use of Bayesian inference for a long time. However, the advent of numerical algorithms and increased computing power enabled the approximation of posterior moments and other quantities of interest, such as the posterior probability of the parameter falling within a given set.

3. Bayesian Hierarchical Modeling of Sea Surface Temperature in the Mediterranean

Sea surface temperature (SST) is an important indicator of environmental change, particularly in the context of global warming and marine feedback in the carbon cycle. The Mediterranean Sea, as one of the largest semi-closed European seas, serves as a valuable model for understanding broader oceanic changes. Nykjaer (2009) indicated that the Mediterranean's upper-layer temperature has been increasing by $0.03 \pm 0.008^{\circ}\text{C}$ per year from 1985 to 2006, consistent with global warming trends. By analyzing SST data from different devices in December 2003, this section aims to provide insights into environmental changes in the Mediterranean and their implications for climate change (Joos et al., 1999).

Here, we present an analysis of SST measurements collected in the Mediterranean in December 2003 using four different device types across various locations. The dataset consists of 336 observations from 86 devices, including bucket, engine room intake (ERI), fixed buoys, and drifting buoys, with some devices providing replicated observations. The goal is to fit a Bayesian hierarchical model to account for the variability associated with device types, using MCMC to explore the posterior distribution. The analysis shows that the commonality of SST measurements varies depending on the device type, with the overall mean SST estimated to be 19.72°C with a standard deviation of 0.65°C .

TABLE 1: Summary statistics of SST by device type.

	SST			No of Records per device		
	N	Mean	SD	Mean	Min	Max
Bucket	36	19.3	1.5	1.0	1	1
D.Buoy	10	19.9	0.9	2.1	1	3
ERI	35	19.6	1.4	1.1	1	2
F.Buoy	5	20.4	0.9	48.0	23	94
Overall	86	19.6	1.4	3.9	1	94

Among the 86 devices, 34 are bucket devices and 35 are ERIs, each representing approximately 40% of the total. The remaining devices include 10 drifting buoys (d.buoys) and 5 fixed buoys (f.buoys). According to the descriptive statistics (Table 1) and the box plots (not shown here), the SSTs associated with different device types show slight variations. Fixed buoys have the highest mean and median SST among all groups. However, it is important to note that fixed buoys also have the fewest devices (only five) but the highest number of records per device, averaging 48 records. The box plots indicate an outlier in the ERI group, originating from device 59. The SST reading of device 59 is 15.0°C , the lowest among all devices.

3.1. Modeling

To assess the mean and variability of SST with device type as a predictor, we formulate a two-level Bayesian hierarchical model and use MCMC methods in order to estimate the model parameters.

3.1.1. Bayesian Hierarchical Modeling

Let $y_{i,j}$ be the average SST measurement from the i -th device of the j -th device type, and $n_{i,j}$ be the corresponding number of records associated with this device. We define a two-level hierarchical model in which the SST measurement $y_{i,j}$ for each individual device is modeled as varying around the mean SST of its corresponding device type θ_j (level 1), and the mean SSTs of the four device types vary around a global mean μ (level 2). Specifically, we assume that $y_{i,j} = \theta_j + \epsilon_{i,j}$, where the $\epsilon_{i,j}$ are conditionally independent (as in the rest of the document), following $\epsilon_{i,j} \mid \sigma_j^2 \sim t_5(0, \sigma_j^2/n_{i,j})$, and $\theta_j \mid \mu, \tau^2 \sim N(\mu, \tau^2)$. Here, $j = 1 \dots 4$ represents the device types: Bucket devices, drifting buoys, ERI, and fixed buoys, respectively, and the error variance within each device group is denoted by σ_j^2 .

On the one hand, note that at level 1, we model the errors using a Student's t distribution with 5 degrees of freedom. This choice is motivated by the heavier tails of the t distribution compared to the Gaussian distribution, allowing for greater robustness to outliers and reducing their influence on the model estimates. To facilitate sampling, we introduce auxiliary variables to represent such t_5 distribution (see [Albert & Chib, 1993](#) and [Gelman et al., 2013](#), Chap. 12). Specifically, we introduce $\lambda_{i,j}$ for each $y_{i,j}$, such that $y_{i,j} \mid \theta_j, \sigma_j^2, \lambda_{i,j} \sim N(\theta_j, \sigma_j^2/(n_{i,j}\lambda_{i,j}))$, with $\lambda_{i,j} \sim G(5/2, 5/2)$. Upon marginalization, it can be shown that this structure ensures that the errors follow a t_5 distribution. On the other hand, we assume that all σ_j^2 share a common scale σ^2 and model them as $\sigma_j^2 \mid \sigma^2 \sim IG(\alpha + 1, \alpha\sigma^2)$, with $\sigma^2 \sim G(a, b)$. We set $\alpha = 2$ to allow some variability in each σ_j^2 relative to σ^2 , but not excessively. For instance, with $\sigma^2 = 1$, about 90% of the prior density for σ_j^2 lies between 0.3 and 3. Additionally, we use $a = b = 2$ to specify a relatively weak prior for σ^2 , centered at 1.

At level 2, we assume that the mean SSTs of each device type are drawn from a normal distribution centered around a global mean μ with variance τ^2 . Therefore, we model $\theta_j \mid \mu, \tau^2 \sim N(\mu, \tau^2)$. Finally, we assign a flat prior for (μ, τ^2) , meaning that $p(\mu, \log \tau^2) \propto 1$, or equivalently, $p(\mu, \tau^2) \propto 1/\tau^2$ ([Gelman et al., 2013](#), Chap. 3).

3.1.2. Computation

Let $J = 4$ be the number of device types, and let N_j be the number of devices of type j . Then, the joint posterior is given by:

$$p(\boldsymbol{\theta}, \boldsymbol{\sigma}^2, \boldsymbol{\lambda}, \mu, \tau^2, \sigma^2 \mid \mathbf{y}) \propto \prod_{j=1}^J \prod_{i=1}^{N_j} \frac{\lambda_{i,j}^{1/2}}{\sigma_j} \exp\left(-\frac{n_{i,j}\lambda_{i,j}}{2\sigma_j^2} (y_{i,j} - \theta_j)^2\right) \\ \cdot \prod_{j=1}^J \prod_{i=1}^{N_j} \lambda_{i,j}^{5/2-1} \exp\left(-\frac{5}{2} \lambda_{i,j}\right) \cdot \prod_{j=1}^J \frac{1}{\tau} \exp\left(-\frac{1}{2\tau^2} (\theta_j - \mu)^2\right) \\ \cdot \prod_{j=1}^J (\sigma_j^2)^{-(\alpha+1)} \exp\left(-\frac{1}{\sigma_j^2} \alpha \sigma^2\right) \cdot (\sigma^2)^{a-1} \exp(-b\sigma^2) \cdot \frac{1}{\tau^2},$$

where $\boldsymbol{\theta} = (\theta_j)$, $\boldsymbol{\sigma}^2 = (\sigma_j^2)$, $\boldsymbol{\lambda} = (\lambda_{i,j})$, and $\mathbf{y} = (y_{i,j})$.

We implement an MCMC algorithm to estimate the model parameters. The chain starts with 10 000 burn-in iterations, followed by 50 000 iterations with a thinning factor of 10, as determined from a test run. Thus, the final sample consists of 5 000 thinned simulations. The algorithm is based on the following update steps:

1. Sample $\theta_j \mid \mu, \tau^2, \sigma_j^2, \lambda_{i,j} \sim \mathcal{N}(\mu_{\theta_j}, \sigma_{\theta_j}^2)$, for $j = 1, \dots, J$, where

$$\mu_{\theta_j} = \frac{\frac{\mu}{\tau^2} + \sum_{j=1}^{N_j} \frac{n_{i,j}}{\sigma_j^2} \lambda_{i,j} y_{i,j}}{\frac{1}{\tau^2} + \sum_{j=1}^{N_j} \frac{n_{i,j}}{\sigma_j^2} \lambda_{i,j}}, \quad \sigma_{\theta_j}^2 = \frac{1}{\frac{1}{\tau^2} + \sum_{j=1}^{N_j} \frac{n_{i,j}}{\sigma_j^2} \lambda_{i,j}}.$$

2. Sample $\sigma_j^2 \mid \theta_j, \lambda_{i,j}, \sigma^2 \sim \text{IG}(\alpha_{\sigma_j}, \beta_{\sigma_j})$, for $j = 1, \dots, J$, where

$$\alpha_{\sigma_j} = \alpha + \frac{N_j}{2}, \quad \beta_{\sigma_j} = \alpha \sigma^2 + \frac{1}{2} \sum_{i=1}^{N_j} n_{i,j} \lambda_{i,j} (y_{i,j} - \theta_j)^2.$$

3. Sample $\lambda_{i,j} \mid \theta_j, \sigma_j^2 \sim \mathcal{G}(\alpha_{\lambda_{i,j}}, \beta_{\lambda_{i,j}})$, for $i = 1 \dots, N_j$ and $j = 1 \dots J$, where

$$\alpha_{\lambda_{i,j}} = 3 \quad \beta_{\lambda_{i,j}} = \frac{5}{2} + \frac{1}{2} \frac{n_{i,j}}{\sigma_j^2} (y_{i,j} - \theta_j)^2.$$

4. Sample $\mu \mid \boldsymbol{\theta}, \tau^2 \sim \mathcal{N}(\mu_\mu, \sigma_\mu^2)$, where

$$\mu_\mu = \frac{1}{J} \sum_{j=1}^J \theta_j, \quad \sigma_\mu^2 = \frac{\tau^2}{J}.$$

5. Sample $\tau^2 \mid \mu, \boldsymbol{\theta} \sim \text{IG}(\alpha_\tau, \beta_\tau)$, where

$$\alpha_\tau = \frac{J-1}{2}, \quad \beta_\tau = \frac{1}{2} \sum_{j=1}^J (\theta_j - \mu)^2.$$

6. Sample $\sigma^2 \mid \boldsymbol{\sigma}^2 \sim \mathcal{G}(\alpha_\sigma, \beta_\sigma)$, where

$$\alpha_\sigma = a, \quad \beta_\sigma = b + \alpha \sum_{j=1}^J \frac{1}{\sigma_j^2}.$$

3.2. Illustration

In this case, the model quickly converges once the chain begins burning, and with a thinning factor of 20, the final sample comprises 5 000 simulations per model parameter. Across all variables in the final sample, the average effective sample size is 4 969, indicating negligible autocorrelation after thinning. Standard MCMC convergence diagnostics (including trace plots, potential scale reduction factors, and effective sample sizes for all monitored parameters) indicated good mixing and did not reveal any evidence of convergence problems.

TABLE 2: Posterior statistics of θ , σ , μ , τ , and σ .

Parameter	Mean	SD	2.5%	50%	97.5%
θ_1	19.49	0.23	18.99	19.49	19.91
θ_2	19.75	0.25	19.28	19.73	20.29
θ_3	19.61	0.20	19.20	19.61	20.00
θ_4	20.04	0.38	19.39	20.01	20.80
σ_1	1.27	0.18	0.96	1.25	1.67
σ_2	1.04	0.27	0.64	0.99	1.69
σ_3	1.27	0.18	0.89	1.16	1.55
σ_4	3.95	1.52	1.75	3.70	7.50
μ	19.72	0.85	18.88	19.69	20.70
τ	0.65	1.16	0.02	0.42	2.67
σ	0.50	0.19	0.18	0.49	0.89

3.2.1. Parameter Estimation

Table 2 summarizes the posterior distributions of the key parameters: $\theta_1, \dots, \theta_J$, $\sigma_1, \dots, \sigma_J$, μ , τ , and σ . The posterior means for θ_j range from 19.49°C to 20.04°C, reflecting a pattern consistent with Table 1. Specifically, the bucket device records the lowest mean, followed by ERI and drifting buoys, while floating buoys exhibit the highest mean. The posterior means for μ and τ are 19.72°C and 0.65°C, respectively. Moreover, histograms of the model parameters (not shown here) suggest that most parameters exhibit approximately symmetric, unimodal posterior distributions with relatively low variance. However, the variance components τ^2 and σ^2 deviate from this pattern, displaying right-skewed posterior distributions with the majority of the probability mass concentrated near their respective posterior means.

Figure 1 presents the posterior means and 95% quantile-based credible intervals for $\lambda_{i,j}$ across all devices. The $\lambda_{i,j}$ values are generally on the same level, with reasonable 95% credible intervals. The posterior means of $\lambda_{i,j}$ range from 0.30 to 1.20, with an average of 0.97 and a median of 0.98. No notable differences are observed between the $\lambda_{i,j}$ values across different device types.

3.2.2. Commonality Among Device Types

Table 2 highlights the heterogeneity in within-type variances. For the first three device types (bucket devices, drifting buoys, and ERI), the σ_j values are relatively close to 1. In contrast, σ_4 for fixed buoys is notably higher, with a posterior mean of 3.95. This disparity is expected, as $y_{i,j}$ represents the mean of all $n_{i,j}$ records collected by device (i, j) , and its variance is given by $\sigma_j^2/n_{i,j}$. Consequently, larger σ_j values suggest greater variability in SST measurements within that device type, particularly for fixed buoys.

Figure 2 displays the $y_{i,j}$ values (represented by the centers of the circles) and the corresponding $n_{i,j}$ values (indicated by the area of the circles) for each device group. In the fixed buoy group, the large circle areas indicate that each $y_{i,j}$ is derived from averaging a substantial number of records. As noted in Table 1, the average number of records per $y_{i,j}$ in the fixed buoy group is 48, nearly 50 times




FIGURE 1: Posterior means and 95% quantile-based credible intervals for $\lambda_{i,j}$ across all devices, with the reference line indicating $\bar{\lambda} = 0.97$.

greater than that of the other three device types. Thus, while the $y_{i,j}$ values for fixed buoy devices seem closely clustered, this is due to averaging over large $n_{i,j}$ values. In fact, the variance within the fixed buoy group is not inherently small, as the $y_{i,j}$ values represent means calculated from a significant number of records.

A large variance within a device type indicates low commonality between individual devices of the same type. For device types such as bucket, drifting bucket, and ERI, where σ_j values are approximately 1°C , devices within each type exhibit some degree of shared measurement characteristics. In contrast, devices within the fixed buoy type show very weak commonality, as evidenced by σ_4 being around 4°C .

3.2.3. Outlier Detection

By employing the auxiliary variables $\lambda_{i,j}$ to reparameterize the t_5 distribution associated with the error terms $\epsilon_{i,j}$, we can utilize these $\lambda_{i,j}$ values for each device to identify potential outliers. Under this parameterization, the variance of $y_{i,j}$ is expressed as $\sigma_j^2/(n_{i,j}\lambda_{i,j})$. Consequently, a very small $\lambda_{i,j}$ corresponds to a large variance, indicating a significant deviation from the mean and thus signaling an outlier.

Figure 3 shows the posterior mean and median of the $\lambda_{i,j}$ values. Only three devices have $\lambda_{i,j}$ values below 0.5: Device 59 (circled in black) and devices 13 and 44 (circled in red). Device 59 stands out with the lowest $\lambda_{i,j}$ value, having a mean

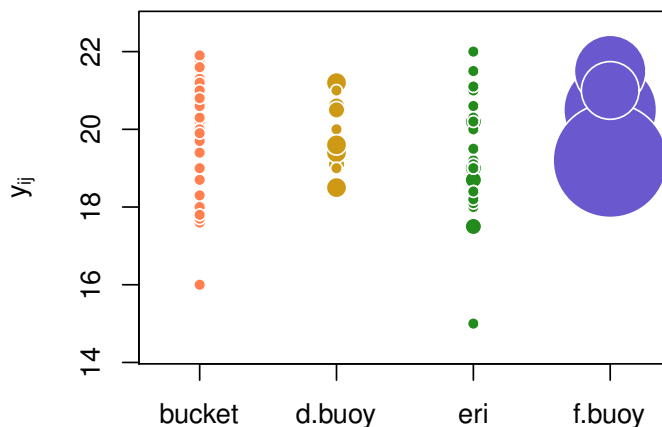


FIGURE 2: $y_{i,j}$ (represented by the center of the circle) and the corresponding $n_{i,j}$ (proportional to the area of the circles) by device group.

of 0.30 and a median of 0.25, and is notably separated from all other $\lambda_{i,j}$ values, marking it as a clear outlier. In contrast, devices 13 and 44 also have very low $\lambda_{i,j}$ values but are not substantially distant from other low $\lambda_{i,j}$ values, classifying them as potential outliers rather than definitive ones.

3.2.4. Goodness-of-fit

We compute the Bayesian p -values (Gelman et al. (2013, Chap. 6)) in order to evaluate the goodness-of-fit of the model. The Bayesian p -value evaluates the fit of a Bayesian model by comparing observed data \mathbf{y} to the posterior predictive distribution of replicated data \mathbf{z} . Such value is defined as the probability that \mathbf{z} is more extreme than \mathbf{y} based on a test quantity T , calculated as:

$$p_B = \Pr(T(\mathbf{z}, \boldsymbol{\theta}) \geq T(\mathbf{y}, \boldsymbol{\theta}) \mid \mathbf{y}) = \iint I(T(\mathbf{z}, \boldsymbol{\theta}) \geq T(\mathbf{y}, \boldsymbol{\theta})) p(\mathbf{z} \mid \boldsymbol{\theta}) p(\boldsymbol{\theta} \mid \mathbf{y}) d\mathbf{z} d\boldsymbol{\theta},$$

where $I(\cdot)$ is the indicator function and $\boldsymbol{\theta}$ are the model parameters. This probability incorporates uncertainty in both the posterior distribution of $\boldsymbol{\theta}$ and the posterior predictive distribution of \mathbf{z} , which provides a measure of model fit. Here, extreme values of the p -values are a sign of lack of model fit.

In this case, we predict the new observations of each device using the MCMC samples. In this way, for each $y_{i,j}$ in the k -th MCMC sample, we generate $z_{i,j}$ from

$$z_{i,j}^{(k)} \sim N(\theta_j^{(k)}, \sigma_j^{2(k)} / (n_{i,j} \lambda_{i,j}^{(k)})),$$




FIGURE 3: Posterior mean and median of the λ_{ij} values, with the reference line indicating $\bar{\lambda} = 0.97$.

and we define the test quantity as $T(z_{i,j}^{(k)}, \theta_j^{(k)}) = |z_{i,j}^{(k)} - \theta_j^{(k)}|$. Based on the predicted $z_{i,j}$'s for each device and each sample from the MCMC simulation, the average Bayesian p -value turns out to be 0.44. Therefore, we consider that our model fits the data relatively well.

Alternatively, for a graphical check of the goodness-of-fit, we use a residual-predicted value plot, where the x -axis represents the average predicted value $z_{i,j}$ and the y -axis represents the average residual $y_{i,j} - z_{i,j}$. A healthy residual plot should have randomly scattered points with no specific patterns. Figure 4 shows the residual-predicted value plot, in which we observe that the residuals exhibit a healthy elliptical distribution with no discernible pattern. The three points detached from the others correspond to the outlier (device 59, displayed in green) and the potential outliers (devices 13 and 44, displayed in coral) that we identified previously.

3.2.5. Remarks

Under our Bayesian hierarchical model, the mean SST of certain areas of the Mediterranean in December 2013 is estimated to be 19.72 ± 0.65 °C. The variability among devices of the same type differs depending on the device type. More detailed data would enable more accurate estimation of both the between- and within-device-type variation in SST measurements.




FIGURE 4: Residual-predicted value plot.

4. Bayesian Spatial Modeling of Rainfall Data

The temporal and geographic distribution of precipitation provides valuable insights into the water balance and the impacts of climate change (Sarmiento, 2000). Understanding rainfall patterns and making location- and time-specific precipitation predictions are critical challenges in agriculture and ecology. In this section, we present a Bayesian hierarchical spatial model to analyze rainfall data collected between 1968 and 1983 in Guárico, one of the 23 states in north-central Venezuela. The dataset also includes the longitude, latitude, altitude of each collection site. Bordered to the north by the central highlands and to the south by the Orinoco River, Guárico spans 25,091 square miles (64,986 square kilometers) of predominantly plains, making it the leading rice-producing state in Venezuela. Analyzing Guárico's rainfall offers valuable insights for local agricultural practices and aids in developing scalable models to better understand the patterns and predictors of annual precipitation.

Table 3 provides a summary of rainfall data for all years and grouped into four-year intervals, revealing noticeable temporal variation over time. This variation suggests the need to account for temporal factors when constructing the model. Furthermore, Figure 5 illustrates the latitude and mean rainfall across all years by location, showing an increase in elevation along the northern edge of Guárico. The data indicate heavier precipitation in the central southern region and along the northern border compared to other areas of the state.

To identify a suitable trend function for the model, we fitted several linear models with mean rainfall as the response variable and several combinations of predictors as explanatory variables. Figure 6 presents the residuals versus predicted value plots used to compare these models and evaluate their fit. Incorporating

TABLE 3: Summary statistics of rainfall by years.

	Records		Rainfall			
	N	Missing	Mean	SD	2.5%	97.5%
All years	1280	190	1047	330	558	1913
1968 - 1971	320	74	1200	331	742	1927
1972 - 1975	320	34	850	219	468	1323
1976 - 1979	320	27	1042	276	645	1688
1980 - 1983	320	55	1066	318	259	1688

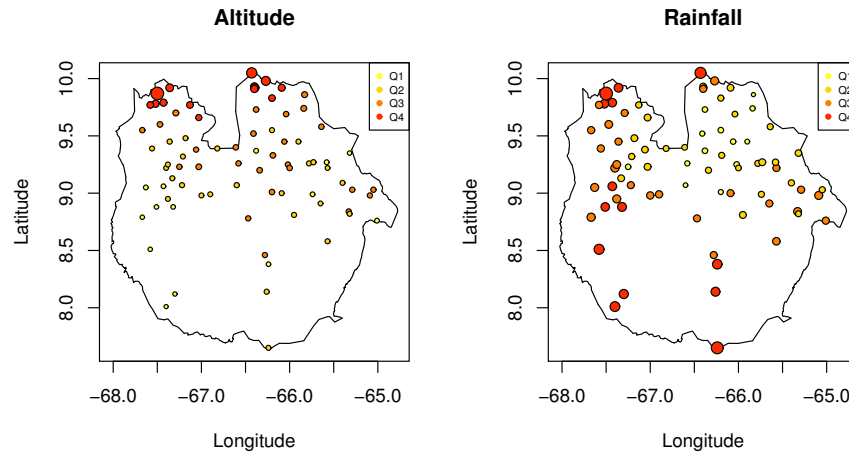


FIGURE 5: Latitude and mean rainfall across all years by location.

rating first-order terms for all variables, along with quadratic terms for year and latitude, substantially improved the model's fit and resulted in well-distributed residuals. In contrast, other second-order terms offered minimal explanatory value and occasionally introduced residual clustering. Based on this analysis, we selected a trend function that includes all first-order terms and quadratic terms for year and latitude to effectively capture the spatial and temporal dynamics.

4.1. Modeling

To analyze the influence of predictors on precipitation and predict rainfall for a given year, we develop a Bayesian hierarchical spatial model. Parameter estimation is performed using an MCMC approach with a block sampler, ensuring efficient computation and reliable inference.

4.1.1. Bayesian Hierarchical Spatial Modeling

Let $\mathbf{X} = (\mathbf{X}_o, \mathbf{X}_*)$ represent the complete dataset, comprising both observed and missing rainfall values, and let $\mathbf{D} = (\mathbf{D}_o, \mathbf{D}_*)$ denote the corresponding trend matrix, which includes the predictors: `longitude`, `latitude`, `altitude`, `year`,

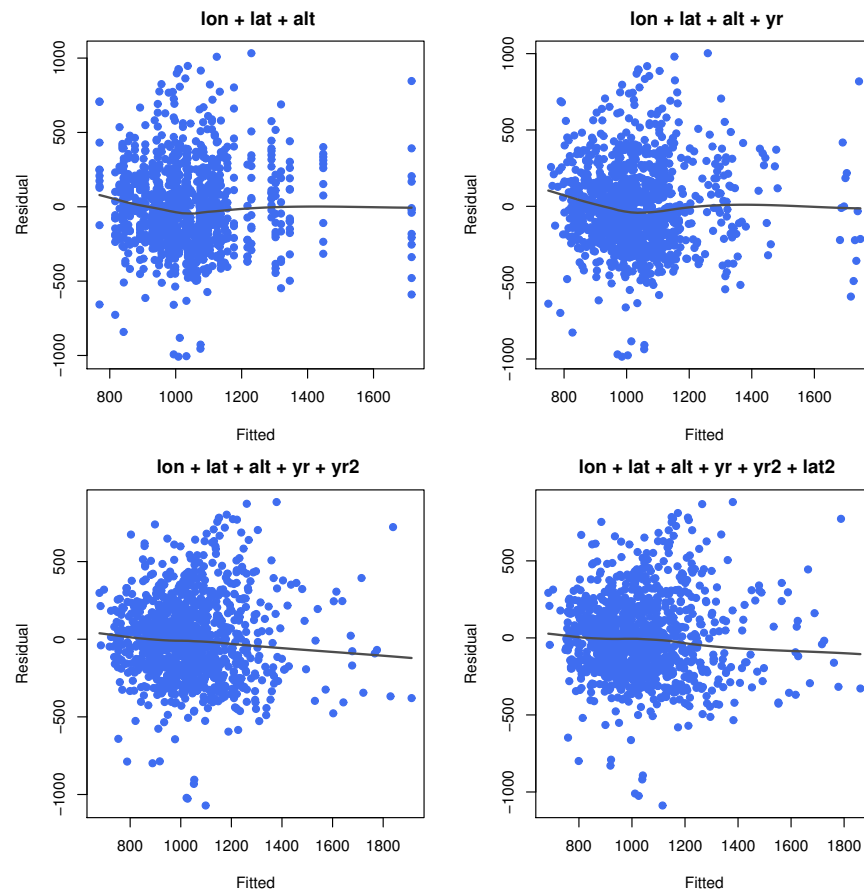


FIGURE 6: Residuals versus predicted values of linear models using mean rainfall as the response variable and several combination of predictors as explanatory variables.

latitude², and year². We assume that \mathbf{X} follows a Gaussian process characterized by the following mean and covariance structure:

$$\mathbf{X} \sim \mathcal{N}(\mathbf{D}\boldsymbol{\beta}, \sigma^2 \mathbf{R}(\phi) + \tau^2 \mathbf{I}),$$

where $\boldsymbol{\beta}$ represents the vector of coefficients of the predictors, σ^2 is the variance of the underlying process, τ^2 is the nugget effect (including observational error), $\mathbf{R}(\phi)$ is a Matérn correlation matrix defined by the range parameter ϕ and a specified order ν , and \mathbf{I} denotes the identity matrix. For a detailed explanation of these spatial modeling fundamentals, refer to [Banerjee et al. \(2015\)](#). For this analysis, ν is fixed at 2.5, selected based on a comparison of the least-squares fit of variograms under Matérn covariance models with varying ν values (Figure 7). A graphical inspection and the corresponding least-squares values indicate that increasing ν beyond 2.5 provides negligible improvement to the fit.

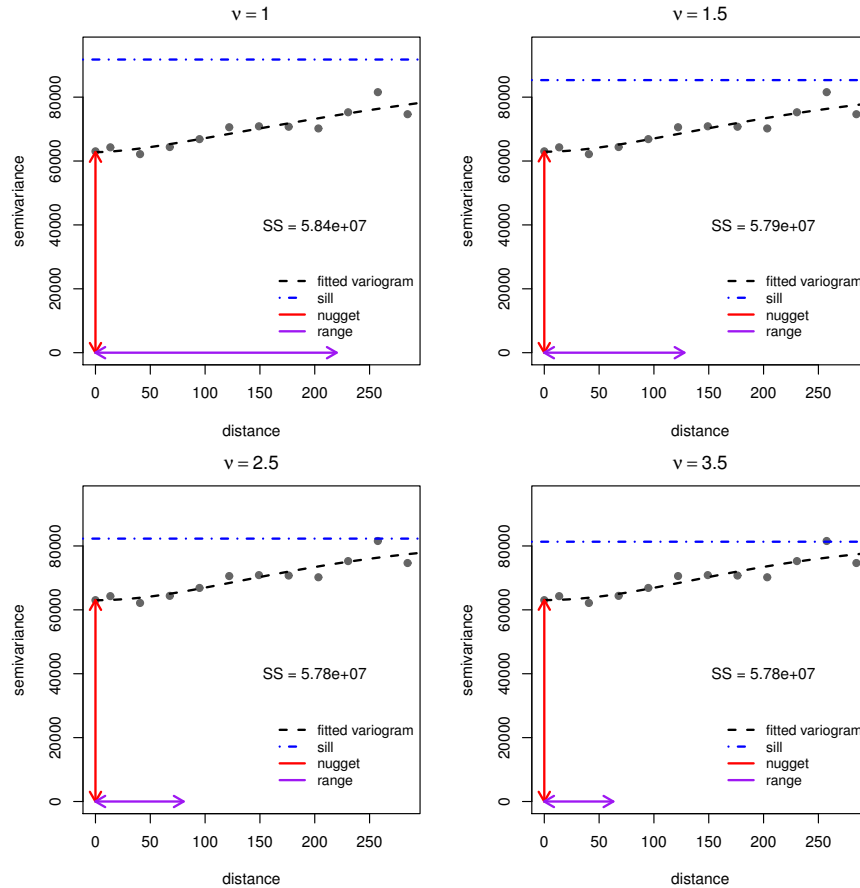


FIGURE 7: Least square fit of variogram using different values of ν .

For convenience in inference and sampling, we make $\gamma^2 = \tau^2/\sigma^2$ and rewrite the variance as $\sigma^2 \mathbf{V}(\phi, \gamma^2)$, where $\mathbf{V}(\phi, \gamma^2) = \mathbf{R}(\phi) + \gamma^2 \mathbf{I}$. Using this reparameterization, we formulate the following Bayesian hierarchical spatial model:

$$\begin{aligned} \mathbf{X} \mid \boldsymbol{\beta}, \sigma^2, \gamma^2, \phi &\sim \mathbf{N}(\mathbf{D}\boldsymbol{\beta}, \sigma^2 \mathbf{R}(\phi) + \tau^2 \mathbf{I}) \\ p(\boldsymbol{\beta}, \sigma^2) &\sim 1/\sigma^2 \\ p(\phi, \gamma^2) &\sim 1/\phi \cdot \mathbf{U}(\gamma^2 \mid 2, 4). \end{aligned}$$

The variogram suggests that the dataset exhibits a significant nugget effect, corresponding to a relatively large γ^2 . As a result, we assign a uniform prior to γ^2 over $(2, 4)$. Noninformative reciprocal priors are chosen for σ^2 and ϕ . To assess the robustness of this prior specification, we performed a simple sensitivity analysis by varying the bounds of the uniform prior on γ^2 and by replacing the reciprocal priors on σ^2 and ϕ with weakly informative Inverse-Gamma distributions whose shape and scale parameters were multiplied and divided by 2. The posterior

summaries of the regression coefficients, the range parameter, and the posterior predictive rainfall surfaces were qualitatively unchanged, indicating that the substantive conclusions are not driven by this particular choice of weakly informative priors.

4.1.2. Computation

Let n and p be the number of data points and the number of predictors, respectively. Under the model given above, the posterior distribution is given by:

$$p(\mathbf{X}, \boldsymbol{\beta}, \sigma^2, \gamma^2, \phi) \propto (\sigma^2)^{-(n-p)/2} |\mathbf{V}(\phi, \gamma^2)|^{-1/2} \\ \cdot \exp \left\{ -\frac{1}{2\sigma^2} [(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})^\top \mathbf{D}^\top \mathbf{V}^{-1}(\phi, \gamma^2) \mathbf{D}(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}) + S^2(\phi, \gamma^2)] \right\} \\ \cdot (\sigma^2)^{-1} \phi^{-1} I(2 < \gamma^2 < 4),$$

where $\hat{\boldsymbol{\beta}}$ is a solution to

$$\mathbf{D}^\top \mathbf{V}^{-1}(\phi, \gamma^2) \mathbf{D} \hat{\boldsymbol{\beta}} = \mathbf{D}^\top \mathbf{V}^{-1}(\phi, \gamma^2) \mathbf{X},$$

$S^2(\phi, \gamma^2) = (\mathbf{X} - \mathbf{D} \hat{\boldsymbol{\beta}})^\top \mathbf{V}^{-1}(\phi, \gamma^2) (\mathbf{X} - \mathbf{D} \hat{\boldsymbol{\beta}})$, and $I(\cdot)$ is the indicator function. Recall that $\hat{\boldsymbol{\beta}}$ is unique if and only if \mathbf{D} has full column rank.

Note that \mathbf{X} also includes \mathbf{X}_* , which represents the latent variables corresponding to the missing values. The algorithm consists of a block sampler for all model parameters $(\boldsymbol{\beta}, \sigma^2, \gamma^2, \phi)$ and a Gibbs step for \mathbf{X}_* conditioned on all other parameters.

*Block sampler for $(\boldsymbol{\beta}, \sigma^2, \gamma^2, \phi)$ given \mathbf{X}_**

Define $\boldsymbol{\psi} = (\phi, \gamma^2)$. If we consider the following prior $p(\boldsymbol{\beta}, \sigma^2, \boldsymbol{\psi}) \propto p(\boldsymbol{\psi}) |G(\sigma^2 | a, b)|$, it can be shown that the posterior distribution can be rewritten in a product form as follows:

$$p(\boldsymbol{\beta}, \sigma^2, \boldsymbol{\psi} | \mathbf{X}) = p(\boldsymbol{\beta} | \sigma^2, \boldsymbol{\psi}, \mathbf{X}) p(\sigma^2 | \boldsymbol{\psi}, \mathbf{X}) p(\boldsymbol{\psi} | \mathbf{X}),$$

where

$$p(\boldsymbol{\beta} | \sigma^2, \boldsymbol{\psi}, \mathbf{X}) \propto \exp \left\{ -\frac{1}{2\sigma^2} (\boldsymbol{\beta} - \hat{\boldsymbol{\beta}})^\top \mathbf{D}^\top \mathbf{V}^{-1}(\boldsymbol{\psi}) \mathbf{D}(\boldsymbol{\beta} - \hat{\boldsymbol{\beta}}) \right\} \\ p(\sigma^2 | \boldsymbol{\psi}, \mathbf{X}) \propto (\sigma^2)^{-(\frac{n-k}{2} + a + 1)} \exp \left\{ -\frac{1}{2\sigma^2} (S^2(\boldsymbol{\psi}) + b) \right\} \\ p(\boldsymbol{\psi} | \mathbf{X}) \propto |\mathbf{V}(\boldsymbol{\psi})|^{-\frac{1}{2}} |\mathbf{D}^\top \mathbf{V}(\boldsymbol{\psi}) \mathbf{D}|^{-\frac{1}{2}} (S^2(\boldsymbol{\psi}) + 2b)^{-(\frac{n-k}{2} + a)} p(\boldsymbol{\psi}),$$

with $S^2(\boldsymbol{\psi}) = (\mathbf{X} - \mathbf{D} \hat{\boldsymbol{\beta}})^\top \mathbf{V}^{-1}(\boldsymbol{\psi}) (\mathbf{X} - \mathbf{D} \hat{\boldsymbol{\beta}})$. The above also applies to a reciprocal prior for σ^2 by setting $a = 0$ and $b = 0$, such that $p(\sigma^2) \propto 1/\sigma^2$. This factorization suggests that we can apply a blocking strategy to sample from the posterior.

We choose a jumping distribution g_1 for $\boldsymbol{\psi}$ and define the block jumping distribution as:

$$g(\boldsymbol{\beta}, \sigma^2, \boldsymbol{\psi}) = p(\boldsymbol{\beta} | \sigma^2, \boldsymbol{\psi}, \mathbf{X}) p(\sigma^2 | \boldsymbol{\psi}, \mathbf{X}) g_1(\boldsymbol{\psi}),$$

where the acceptance ratio of the current sample $\boldsymbol{\psi}^{(i)}$ and the proposed sample $\boldsymbol{\psi}^*$ is:

$$r = \frac{p(\boldsymbol{\beta}^*, \sigma^{2*}, \boldsymbol{\psi}^* | \mathbf{X})}{p(\boldsymbol{\beta}^{(i)}, \sigma^{2(i)}, \boldsymbol{\psi}^{(i)} | \mathbf{X})} \cdot \frac{g(\boldsymbol{\beta}^{(i)}, \sigma^{2(i)}, \boldsymbol{\psi}^{(i)})}{g(\boldsymbol{\beta}^*, \sigma^{2*}, \boldsymbol{\psi}^*)} = \frac{p(\boldsymbol{\psi}^* | \mathbf{X}) g_1(\boldsymbol{\psi}^{(i)})}{p(\boldsymbol{\psi}^{(i)} | \mathbf{X}) g_1(\boldsymbol{\psi}^*)}.$$

Based on the above result, we only sample $(\boldsymbol{\beta}, \sigma^2)$ if the proposed $\boldsymbol{\psi}^*$ is accepted. Conditioning on $\boldsymbol{\psi}$, we first sample σ^2 , then sample $\boldsymbol{\beta}$ conditioning on both σ^2 and $\boldsymbol{\psi}$. Thus, we sample $(\boldsymbol{\beta}, \sigma^2)$ as follows:

1. Sample $\sigma^2 | \boldsymbol{\psi} \sim \text{IG}((n - k)/2, S^2(\boldsymbol{\psi})/2)$.
2. Sample $\boldsymbol{\beta} | \sigma^2, \boldsymbol{\psi} \sim \text{N}(\hat{\boldsymbol{\beta}}, \sigma^2 \mathbf{V}_\beta)$, where

$$\hat{\boldsymbol{\beta}} = \mathbf{V}_\beta \mathbf{D}^\top \mathbf{V}^{-1}(\boldsymbol{\psi}) \mathbf{X}, \quad \mathbf{V}_\beta = (\mathbf{D}^\top \mathbf{V}^{-1}(\boldsymbol{\psi}) \mathbf{D})^{-1}.$$

In practice, instead of inverting matrices, we use Cholesky decomposition and QR decomposition (e.g., [Golub & Van Loan \(2013\)](#)) to obtain $\hat{\boldsymbol{\beta}}$ and sample from the posterior of $\boldsymbol{\beta}$. Moreover, since the sample size is large, to facilitate sampling, we sample $\boldsymbol{\psi}$ using a grid approach (e.g., [Gelman et al. \(2013, Chap. 2\)](#)), which is analogous to applying independent discrete priors. By doing this, we gain efficiency and stability in computation at the cost of a coarse resolution when exploring the posterior density of $\boldsymbol{\psi}$.

Gibbs step of missing values \mathbf{X}_ given $(\boldsymbol{\beta}, \sigma^2, \gamma^2, \phi)$*

It follows that $(\mathbf{X}_*, \mathbf{X}_o)$ has a joint Normal distribution. In this way, we rearrange the variance matrix into the following form:

$$\tilde{\mathbf{V}} = \begin{pmatrix} \tilde{\mathbf{V}}_{**} & \tilde{\mathbf{V}}_{*o} \\ \tilde{\mathbf{V}}_{o*} & \tilde{\mathbf{V}}_{oo} \end{pmatrix},$$

where $\tilde{\mathbf{V}}_{**}$ and $\tilde{\mathbf{V}}_{oo}$ are the correlation matrices of \mathbf{X}_* and \mathbf{X}_o , respectively, and $\tilde{\mathbf{V}}_{*o}$ is the correlation matrix between \mathbf{X}_* and \mathbf{X}_o . To update the latent variables and impute missing values, we sample $\mathbf{X}_* | \mathbf{X}_o, \boldsymbol{\beta}, \sigma^2, \boldsymbol{\psi} \sim \text{N}(\boldsymbol{\mu}_*, \boldsymbol{\Sigma}_*)$, where:

$$\boldsymbol{\mu}_* = \mathbf{D}_* \boldsymbol{\beta} + \tilde{\mathbf{V}}_{*o} \tilde{\mathbf{V}}_{oo}^{-1} (\mathbf{X}_o - \mathbf{D}_o \boldsymbol{\beta}), \quad \boldsymbol{\Sigma}_* = \sigma^2 \left(\tilde{\mathbf{V}}_{**} - \tilde{\mathbf{V}}_{*o} \tilde{\mathbf{V}}_{oo}^{-1} \tilde{\mathbf{V}}_{o*} \right),$$

where \mathbf{D}_o and \mathbf{D}_* are the corresponding design matrices of \mathbf{X}_o and \mathbf{X}_* , respectively.

4.2. Illustration

We implement a MCMC algorithm to estimate the model parameters. The chain begins with 500 burn-in iterations to allow it to converge to the target distribution, followed by 1,000 effective iterations used for parameter estimation. Standard MCMC convergence diagnostics (including trace plots, potential scale reduction factors, and effective sample sizes for all monitored parameters) indicated good mixing and did not reveal any evidence of convergence problems.

4.2.1. Parameter Estimation

Table 4 presents the posterior distribution of the model parameters. The variables `longitude`, `latitude`, and `altitude` are scaled in kilometers, while the years are recoded from 1 to 16. The results reveal a strong positive trend in precipitation within the trend function as altitude increases. Specifically, for every 100-meter rise in altitude, the annual rainfall increases by an average of 86.5 millimeters. This observation is consistent with the heavy rainfall patterns observed in northern Guárico, located at the edge of the central highlands.

TABLE 4: Posterior statistics of the model parameters.

Parameter	Mean	SD	2.5%	50%	97.5%
<code>longitude</code>	-0.56	0.054	-1.84	0.55	0.39
<code>latitude</code>	-16.7	0.25	-36.2	-16.7	2.3
<code>altitude</code>	865	81.8	748	859	1037
<code>year</code>	-86.3	5.96	-97.5	-86.2	-73.1
<code>latitude</code> ²	0.01	0.01	0.00	0.01	0.02
<code>year</code> ²	4.78	0.34	4.05	4.78	5.44
σ^2	17098	1005	15444	17027	19245
ϕ	49.8	8.63	35.0	47.0	65.0
γ^2	3.25	0.16	3.00	3.25	3.50

The posterior estimates also indicate a second-order trend in latitude and year, aligning with the observation that higher precipitation levels are concentrated in both the southern and far northern regions of Guárico. Geographically, southern Guárico is bordered by the Orinoco River, a natural source of rainfall. Moreover, the central southern region encompasses the Aguaro-Guariquito National Park, where extensive forest coverage promotes water conservation and creates conditions conducive to abundant rainfall.

The posterior mean of γ^2 is 3.25, reflecting a significant nugget effect and indicating substantial observational error, likely influenced in part by inaccuracies in the manual recording of rainfall data. Furthermore, the formulation of the trend function impacts the distribution of variability between the trend and the spatial component, which may contribute to the estimation of γ^2 . The posterior mean of the range parameter ϕ is 49.8, implying a relatively gradual decay in spatial correlation.

4.2.2. Goodness-of-fit

To assess model fit, we perform validation by excluding two randomly selected observations from the dataset. The model is then refitted using the remaining data. Next, we generate samples from the posterior predictive distribution for the excluded observations and compare these predictions with the observed values. If the true values lie close to the center of the predictive distribution, it provides evidence that the model captures the underlying data structure effectively and fits the dataset well.

Figure 8 presents the histograms of the posterior predicted rainfall for the two hold-out observations, with the true values indicated by blue vertical lines. Both posterior predictive distributions exhibit an approximately symmetric and unimodal shape, and in both cases, the true values are positioned near the centers of their respective distributions. This alignment suggests that the model provides a good fit to the data.

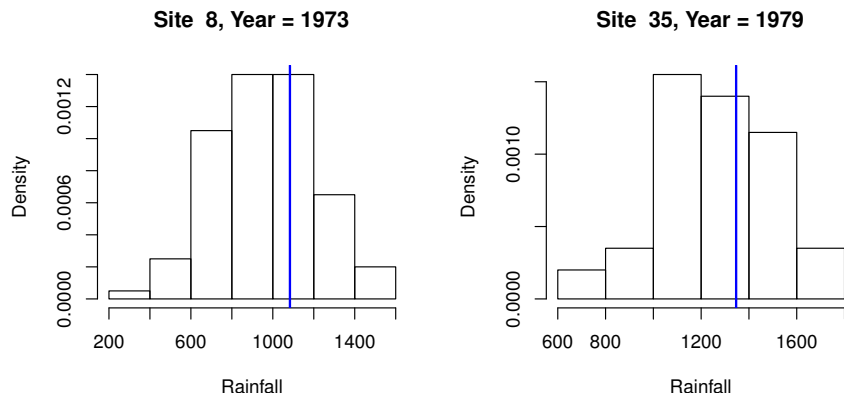


FIGURE 8: Model checking.

4.2.3. Predictions

Predictions at new locations with a covariate matrix $\hat{\mathbf{D}}$ are jointly Normal with the complete data \mathbf{X} . For each sample of \mathbf{X} (including a sample of \mathbf{X}_*) and the parameter vectors, we draw $\hat{\mathbf{X}} \sim N(\boldsymbol{\mu}_{\text{new}}, \boldsymbol{\Sigma}_{\text{new}})$, where

$$\boldsymbol{\mu}_{\text{new}} = \hat{\mathbf{D}}\boldsymbol{\beta} + \mathbf{R}\mathbf{V}^{-1}(\mathbf{X} - \mathbf{D}\boldsymbol{\beta}), \quad \boldsymbol{\Sigma}_{\text{new}} = \sigma^2(\hat{\mathbf{V}} - \mathbf{R}\mathbf{V}^{-1}\mathbf{R}),$$

with $\hat{\mathbf{V}}$ representing the correlation matrix of $\hat{\mathbf{X}}$, and \mathbf{R} denoting the correlation matrix between $(\hat{\mathbf{X}}, \mathbf{X})$. To further analyze predictions, we compute the probability of rainfall exceeding 1200 millimeters, which corresponds to the third quartile of all non-missing observations, using the predictive samples. Since altitude data at the new location is unavailable, we impute the altitudes using those from the nearest collection site.

Figure 9 illustrates the posterior predictions for four selected years: 1968, 1971, 1979, and 1975. These years were chosen intentionally due to their observed mean annual rainfall, based on non-missing data, which decreases sequentially: approximately 1200 mm in 1968, 1100 mm in 1971, 1000 mm in 1979, and 900 mm in 1975. The maps present the posterior mean and variance of the predicted annual rainfall, as well as the probability of rainfall exceeding 1200 millimeters, plotted over a fixed grid. The color scheme uses red for low values (low precipitation, variance, and low probability of heavy rainfall), yellow for medium values, and white for high precipitation and probabilities.

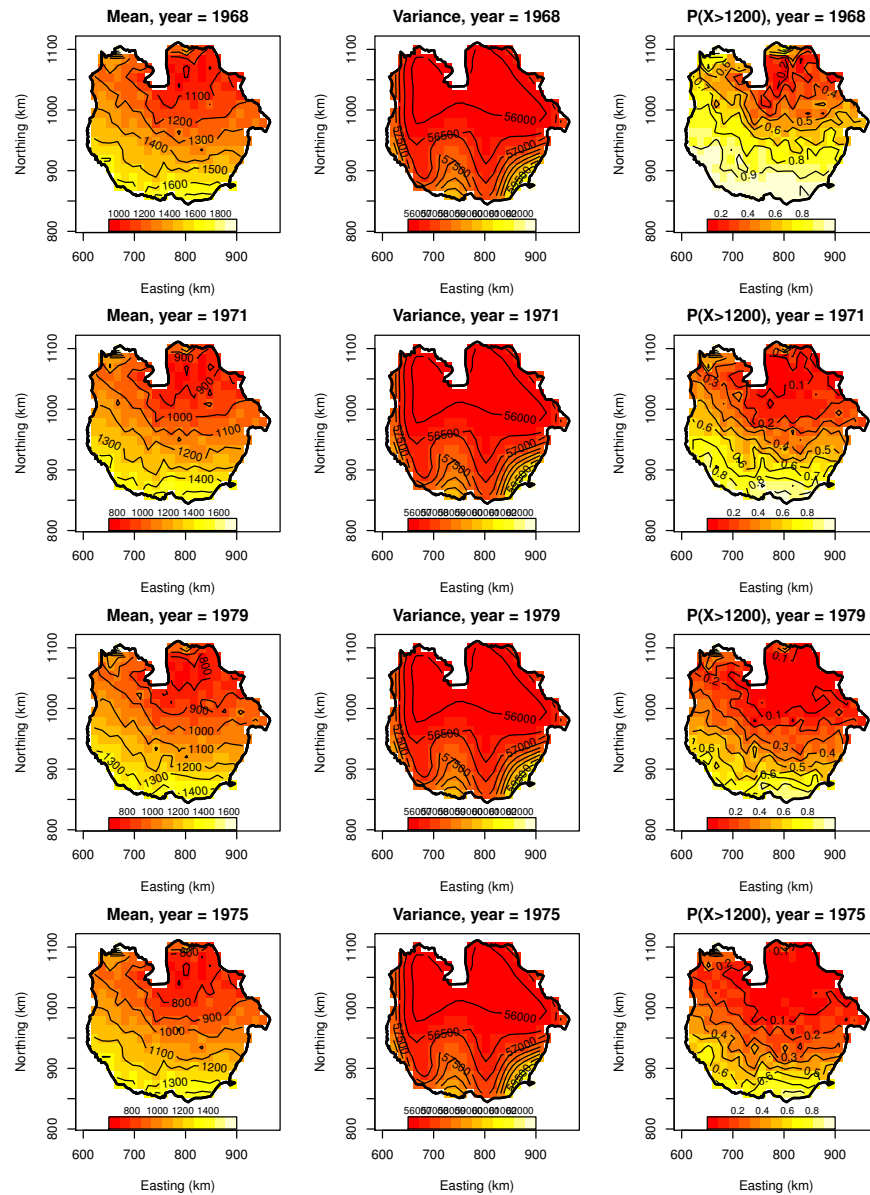


FIGURE 9: Posterior predictions for 1968, 1971, 1979, and 1975.

In all four years, the posterior predictions identify regions of higher precipitation and a greater probability of rainfall exceeding 1200 millimeters in the central southern area and along the northern border of Guárico, where elevation is higher, relative to the rest of the state. While the spatial patterns remain consistent across the years, the overall mean rainfall varies, with 1968 being the wettest and 1979 the driest. The predictive maps also show an increasing prevalence of red areas

over time, indicating a decline in predicted mean precipitation and a reduced probability of rainfall exceeding 1200 millimeters from 1968 to 1971, 1979, and 1975. This pattern aligns closely with the decreasing trend observed in the empirical data for these years.

4.2.4. Remarks

Overall, the coefficients of the predictors are consistent with the spatial and temporal patterns observed in the data. The predictions indicate a strong association between rainfall and the site's topographical features, such as whether it is situated on a plain or in the highlands, its proximity to forests or rivers, or the absence of such features. The model also captures a temporal trend effectively.

However, the model has some limitations. First, some model parameters exhibit large standard deviations, which may result from the grid search approach used for the correlation parameters ϕ and γ^2 . This method limits the range of the posterior distribution that can be explored. Another possible reason is that the likelihood surface for this model is relatively flat over a wide range of values for ϕ and γ^2 , leading to uncertainty in parameter estimation.

Additionally, the inclusion of both first- and second-order terms of the year in the trend component, while fitting the dataset well, may present challenges for long-term predictions. To address this, alternative approaches to modeling the temporal trend, such as incorporating it into the covariance matrix through auto-regression, could be considered for improved long-term forecasts.

Finally, computational constraints, due to the large covariance matrix (1280×1280), resulted in a relatively short MCMC chain and fixing the order ν in the Matérn correlation function. Extending the chain length and placing a prior on ν could enhance the inference, providing more robust results and reducing potential biases.

5. Bayesian Modeling of High-Order Markov Chains

[Raftery \(1985\)](#) introduced the mixture transition distribution (MTD) model to facilitate efficient inference for time-homogeneous Markov chains (MCs) with high-order dependencies. Over the past three decades, frequentist methods have been developed to extend the MTD model and improve computational stability, although some challenges persist ([Raftery, 1985](#)). This section focuses on presenting a Bayesian approach to analyzing the MTD model for time-homogeneous Markov chains.

5.1. High-Order Markov Chains

In high-order Markov chains, the state at time t depends on the previous ℓ states, where $\ell \geq 2$. Let $\{X_t\}$ denote an ℓ -order m -state time-homogeneous

Markov chain with a countable state space S . The transition probabilities are defined as:

$$\Pr(X_t = i_0 \mid X_{t-\ell} = i_\ell, \dots, X_{t-1} = i_1) = q_{i_\ell \dots i_0}, \quad i_t, \dots, i_0 \in S.$$

The full transition matrix of a high-order Markov chain is often sparse, containing many structural zeros. To simplify the presentation, we use a reduced form of the transition matrix. For instance, assuming $\ell = 2$ and $m = 2$, the transition matrix can be represented as:

$$\mathbf{P} = \begin{bmatrix} p_{111} & p_{112} \\ p_{121} & p_{122} \\ p_{211} & p_{212} \\ p_{221} & p_{222} \end{bmatrix}.$$

The maximum likelihood estimators (MLEs) for a fully parameterized ℓ -order m -state Markov chain are given by:

$$\hat{p}_{i_\ell \dots i_0} = \frac{n_{i_\ell \dots i_0}}{n_{i_\ell \dots i_1+}}, \quad \text{where} \quad n_{i_\ell \dots i_1+} = \sum_{i_0=1}^m n_{i_\ell \dots i_0}.$$

Here, $n_{i_\ell \dots i_0}$ represents the number of transitions from $(X_{t-\ell} = i_\ell, \dots, X_{t-1} = i_1)$ to $X_t = i_0$. In a fully parameterized ℓ -order m -state Markov chain model, the number of free parameters is $m^\ell(m-1)$.

5.2. The MTD Model

The MTD model reduces the dimensionality of a fully parameterized model by decomposing the transition matrix \mathbf{P} into a linear combination of contributions from each of the past ℓ states. The transition probability is expressed as:

$$\Pr(X_t = i_0 \mid X_{t-\ell} = i_\ell, \dots, X_{t-1} = i_1) = \sum_{g=1}^{\ell} \lambda_g \Pr(X_t = i_0 \mid X_{t-g} = i_g) = \sum_{g=1}^{\ell} \lambda_g q_{i_g i_0},$$

where $q_{i_g i_0}$ are elements of an $m \times m$ transition matrix $\mathbf{Q} = [q_{i_g i_0}]$, with each row forming a probability distribution. The vector $\boldsymbol{\lambda} = (\lambda_\ell, \dots, \lambda_1)$ contains lag parameters representing the contribution from each lag.

To ensure valid probabilities, the model imposes the following constraints:

$$\sum_{g=1}^{\ell} \lambda_g = 1, \quad \lambda_g \geq 0, \quad g = 1, \dots, \ell. \quad (1)$$

Under the MTD model, the number of free parameters is $m(m-1) + (\ell-1)$, significantly fewer than the $m^\ell(m-1)$ parameters in the full model (see Figure 10). Frequentist estimation of the MTD model involves numerically maximizing the log-likelihood under the linear constraints defined in (1), which can be computationally demanding. Although iterative algorithms have been proposed to enhance computational efficiency, they do not guarantee convergence to the global maximum of the log-likelihood (Berchtold & Raftery, 2002).




FIGURE 10: Number of free parameters under the full model and MTD.

5.3. Modeling

In contrast to frequentist estimation algorithms for MTD models, the Bayesian approach avoids the challenges associated with global maximization of the likelihood. For time-homogeneous MCs with a known order, MCMC algorithms can be applied to a model augmented with latent variables, enabling efficient parameter estimation. When the order is fixed but unknown, Bayesian methods leverage posterior model probabilities to perform order inference, providing a systematic approach to determining the most likely model complexity.

5.3.1. Bayesian MTD Models

Let $\{X_t : t = 1, \dots, N + \ell\}$ represent a time-homogeneous m -state ℓ -order MC with state space $S = \{1, \dots, m\}$. Given the first ℓ observations, the likelihood of $\{X_t\}$ under the MTD model can be expressed as

$$p(x_{\ell+1}, \dots, x_{\ell+N} \mid \boldsymbol{\lambda}, \mathbf{Q}, x_1, \dots, x_\ell) = \prod_{t=\ell+1}^{\ell+N} \left(\sum_{g=1}^{\ell} \lambda_g q_{x_t-g, x_t} \right),$$

where $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_\ell)$ is the vector of weights for the past ℓ lags, and $\mathbf{Q} = [q_{i,j}]$ is the $m \times m$ transition matrix for the MTD model.

From the above likelihood, it is apparent that the model resembles a mixture model. Introducing latent variables $\{w_t : t = \ell + 1, \dots, \ell + N\}$ as indicators for the mixture components allows the summation to be converted into a product. In this context, $\{w_t\}$ is a set of indicators specifying the mixture component from which each data point originates. The augmented likelihood then becomes

$$p(x_{\ell+1}, \dots, x_{\ell+N} \mid \boldsymbol{\lambda}, \mathbf{Q}, w_{\ell+1}, \dots, w_{\ell+N}, x_1, \dots, x_\ell) = \prod_{t=\ell+1}^{\ell+N} \prod_{g=1}^{\ell} (\lambda_g q_{x_t-g, x_t})^{I(w_t=g)},$$

where $I(\cdot)$ is the indicator function. This model can be efficiently estimated using MCMC methods. Finally, we assign independent Dirichlet prior distributions to $\boldsymbol{\lambda}$ and each \mathbf{q}_i :

$$\boldsymbol{\lambda} \sim \text{Dir}(b_1, \dots, b_\ell), \quad \mathbf{q}_i \sim \text{Dir}(a_{i,1}, \dots, a_{i,m}), \quad i = 1, \dots, m.$$

5.3.2. Computation

Under the model given above, the posterior distribution is given by

$$p(x_{\ell+1}, \dots, x_{\ell+N}, \boldsymbol{\lambda}, \mathbf{Q}, w_{\ell+1}, \dots, w_{\ell+N} \mid x_1, \dots, x_\ell) \propto \prod_{t=\ell+1}^{\ell+N} \prod_{g=1}^{\ell} (\lambda_g q_{x_t-g, x_t})^{I(w_t=g)} \\ \cdot \prod_{i=1}^m \prod_{j=1}^m q_{ij}^{a_{ij}-1} \cdot \prod_{g=1}^{\ell} \lambda_g^{b_g-1}.$$

MCMC algorithm for MCs with known order

Based on the joint posterior, we run an MCMC algorithm to estimate the parameters using independent $\text{Dir}(1/2, \dots, 1/2)$ priors for both $\boldsymbol{\lambda}$ and \mathbf{q}_i , for $i = 1, \dots, m$. The chain starts with 1,000 burn-in iterations, followed by 5,000 effective iterations. The algorithm is based on the following update steps:

1. Sample $w_t \sim \text{Mult}(\gamma_1, \dots, \gamma_\ell)$, for $t = \ell + 1, \dots, \ell + N$, where

$$\gamma_g = \frac{\lambda_g q_{x_t-g, t}}{\sum_{r=1}^{\ell} \lambda_r q_{x_t-r, t}}, \quad g = 1, \dots, \ell.$$

2. Sample $\boldsymbol{\lambda} \sim \text{Dir}(b_1^*, \dots, b_\ell^*)$, where

$$b_g^* = b_g + \sum_{t=\ell+1}^{\ell+N} I(w_t = g), \quad g = 1, \dots, \ell.$$

3. Sample $\mathbf{q}_i \sim \text{Dir}(a_{i,1}^*, \dots, a_{i,2}^*)$, where

$$a_{i,j}^* = a_{i,j} + \sum_{t=\ell+1}^{\ell+N} I(x_{t-w_t} = i, x_t = j), \quad i, j = 1, \dots, m.$$

MCs with unknown order

When the order of an MC is unknown, two main approaches can be considered. The first involves implementing a reversible jump MCMC algorithm, which allows for variable dimensions of the parameter vector (Richardson & Green, 1997). The second approach fits models of different orders and evaluates them using model selection criteria, such as the Bayes factor.

Here, we use posterior model probabilities for order selection (Scott, 2002). Assume the chain \mathbf{x} has an order $\ell \in \{1, \dots, L\}$, and let $\boldsymbol{\phi} = (\boldsymbol{\phi}_1, \dots, \boldsymbol{\phi}_L)$ denote

the vector of parameters, where ϕ_ℓ is the parameter vector for the ℓ -th order MTD model. A Monte Carlo estimate of the posterior model probability, based on B draws from the posterior distribution, is given by Scott (2002):

$$p(\ell | \mathbf{x}) = \int p(\ell | \mathbf{x}, \phi) p(\phi | \mathbf{x}) d\phi \approx \frac{1}{B} \sum_{j=1}^B p(\ell | \mathbf{x}, \phi^{(j)}),$$

where $p(\ell | \mathbf{x}, \phi^{(j)}) \propto p(\mathbf{x} | \phi^{(j)}, \ell) p(\ell)$. Thus, the posterior model probability $p(\ell | \mathbf{x})$ can be approximated by the mean of the posterior likelihood of the model, up to a proportionality constant.

5.4. Illustration

We implement the method on the infant death rate data assuming up to 5 orders of dependence and draw posterior inference on the order of the chain and the parameters using the aforementioned Bayesian methods. In this way, we use the provisional monthly data on infant death rate per 1,000 population in the United States between July 2011 and June 2014, as reported in the National Vital Statistics Reports, which consists of 36 consecutive months of data. The tables are publicly available on the CDC website, and the technical details are explained at http://www.cdc.gov/nchs/data/nvsr/nvsr58/nvsr58_25.htm.

Instead of using the raw data of the infant death rate (Figure 11, left panel), we analyze the change in the rate compared to the previous month. By assigning state 1 to a decrease, state 2 to no change, and state 3 to an increase, we transform the absolute rates into a three-state chain with state space $S = \{1, 2, 3\}$. We assume up to 5 orders of dependence and take the first five observations as given throughout the analysis.

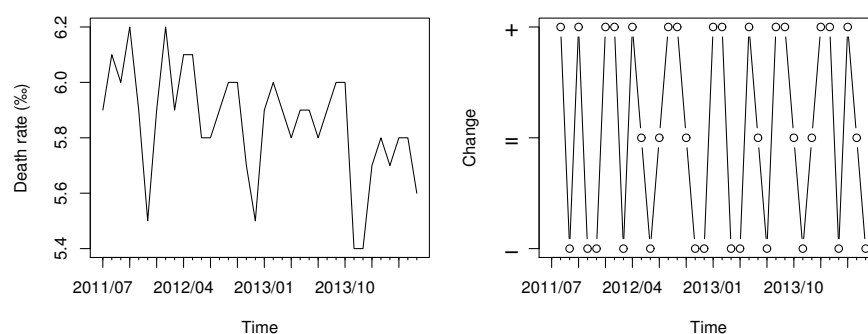


FIGURE 11: Monthly U.S infant death rate (per 1,000 population, left panel) and change in rate (+: increase, =: no change, -: decrease, right panel), from July 2011 to June 2014.

5.4.1. Order of the Chain

We fit five Bayesian MTD models with orders $\ell = 1, \dots, 5$ and calculate the posterior model probability for each, as presented in Table 5. The second- and third-order models exhibit higher posterior model probabilities than the first-order model. Table 5 highlights that the second-order MC ($\ell = 2$) is the most probable within the order space $L = \{1, \dots, 5\}$. Accordingly, we select the second-order model as the most suitable. For subsequent analyses, we adopt $\ell = 2$ as the order of the MTD model.

TABLE 5: Comparison of models based on posterior model probabilities.

Order	Mean Log-Posterior Likelihood	Posterior Model Probability
1	-29.58	0.07
2	-27.17	0.74
3	-28.74	0.15
4	-30.25	0.03
5	-31.76	0.01

5.4.2. Parameter estimation

Table 6 presents the summary statistics of the posterior samples for λ and $q_{i,j}$'s. The substantially larger posterior mean of λ_2 compared to λ_1 suggests that X_{t-2} has a stronger impact than X_{t-1} , which also explains the superior fit of the MTD model with $\ell = 2$ over the first-order model.

TABLE 6: Posterior distribution of λ and $q_{i,j}$.

Parameter	Prior Mean	Posterior			
		Mean	SD	2.5%	97.5%
λ_1	0.500	0.184	0.183	0.000	0.684
λ_2	0.500	0.816	0.183	0.316	1.000
q_{11}	0.330	0.237	0.135	0.009	0.519
q_{12}	0.330	0.716	0.142	0.419	0.965
q_{21}	0.330	0.268	0.195	0.015	0.755
q_{22}	0.330	0.284	0.179	0.007	0.675
q_{31}	0.330	0.645	0.149	0.311	0.891
q_{32}	0.330	0.232	0.127	0.044	0.525

The posterior mean stationary distribution π represents the equilibrium probabilities of the system's state combinations, reflecting the long-term proportion of time the system spends in each state. In this context, the highest probabilities correspond to states (1, 3), (3, 1), and (3, 3), with $\pi_{13} = 0.153$, $\pi_{31} = 0.150$, and $\pi_{33} = 0.155$, indicating these states dominate the system's dynamics. Conversely, states such as (2, 2) and (2, 1), with stationary probabilities of $\pi_{22} = 0.073$ and $\pi_{21} = 0.083$, are visited less frequently, highlighting their lower significance in the system's behavior.

Figure 12 displays box plots corresponding to the posterior samples of π , highlighting a considerable number of samples that fall outside the 1.5 interquartile

range for $\pi_{(2,2)}$. This is largely attributed to the small size of the original dataset (30 data points, excluding the first five observations treated as given) and the lack of observed transitions from state 2 to state 2. Expanding the sample size could mitigate the spread of the distributions and enhance the precision of the estimates.




FIGURE 12: Box plots of posterior samples of π .

5.4.3. Model Assessment

We employ the posterior predictive loss criterion proposed by [Gelfand & Ghosh \(1998\)](#) to evaluate the model's performance. This criterion selects the model \mathcal{M} that minimizes:

$$D_r(\mathcal{M}) = P^{\mathcal{M}} + \frac{r}{r+1} G^{\mathcal{M}},$$

where

$$P^{\mathcal{M}} = \sum_{t=1}^n \text{Var}^{\mathcal{M}}(\hat{y}_t), \quad G^{\mathcal{M}} = \sum_{t=1}^n \left(y_t - \mathbb{E}^{\mathcal{M}}(\hat{y}_t) \right)^2,$$

and $r \geq 0$. Here, $\mathbb{E}^{\mathcal{M}}(\hat{y}_t)$ and $\text{Var}^{\mathcal{M}}(\hat{y}_t)$ denote the mean and variance of the posterior predictive distribution at time t under model \mathcal{M} . The criterion consists of two components: The penalty term $P^{\mathcal{M}}$, which accounts for model complexity, and the goodness-of-fit term $G^{\mathcal{M}}$, weighted by $\frac{r}{r+1}$. This balance ensures a trade-off between model parsimony and fit to the data.

The posterior predictive loss criterion described above is used to compare the estimated ℓ -th order MTD model M_ℓ with the fully parameterized model M_0 . For the fully parameterized model M_0 , the $m^\ell \times m$ transition matrix \mathbf{P} is also estimated within a Bayesian framework. In this approach, each row of transition probabilities in \mathbf{P} is assigned an independent $\text{Dirichlet}(1/2, \dots, 1/2)$ prior, reflecting a weakly informative prior distribution that encourages balance while allowing flexibility in the transition probabilities. This enables a fair comparison between the simpler M_ℓ model and the more complex M_0 , taking into account both model fit and complexity.

We fit the fully parameterized second-order Markov chain model (M_0) and compare its posterior predictive loss to that of the second-order MTD model (M_2), with the results summarized in Table 7. The comparison reveals that M_2 has lower values for both D_1 and D_∞ , indicating a preference for the second-order MTD model over the fully parameterized model in this case. Analyzing the P and G terms shows that M_2 achieves a comparable fit to the data as M_0 , while being more parsimonious. This highlights the advantage of the MTD model in balancing model fit and complexity.

TABLE 7: Posterior predictive loss of fully parameterized second-order Markov chain model (M_0) and second-order MTD model (M_2).

Model	Parameters	P	G	D_1	D_∞
M_0	18	23.21	20.38	33.4	43.59
M_2	7	21.46	20.88	31.9	42.34

5.4.4. Prediction

Given the weight vector λ and the MTD transition matrix \mathbf{Q} , the ℓ -order transition matrix \mathbf{P} can be reconstructed as follows:

$$\mathbf{P}_{i_\ell \dots i_1 i_0} = P(X_t = i_0 \mid X_{t-\ell} = i_\ell, \dots, X_{t-1} = i_1) = \sum_{g=1}^{\ell} \lambda_g q_{i_g i_0}, \quad i_\ell, \dots, i_0 \in S,$$

where $q_{i_g i_0}$ are the elements of the MTD transition matrix $\mathbf{Q} = [q_{i_g i_0}]$. This formulation leverages the structure of the MTD model, which simplifies the high-dimensional ℓ -order Markov chain into a weighted combination of simpler transitions.

To predict future sample paths, new observations can be generated sequentially, conditioned on the observed data. By averaging over the posterior sample draws obtained from the MCMC algorithm, the posterior predictive probability of a specific state at each time point can be calculated. This approach provides a probabilistic framework for forecasting, integrating uncertainty from the posterior distribution into the predictions.

For each posterior sample, we reconstruct the full transition matrix \mathbf{P} and use it to predict the relative change in the infant death rate over the subsequent 12 months. The results are summarized in Figure 13, which illustrates the predicted probabilities for each state in the following year. In this figure, the states represent the relative changes: tan indicates a decrease, coral represents no change, and red signifies an increase compared to the previous month. These probabilities, derived from the posterior samples, provide insights into the expected dynamics of the infant death rate over time, accounting for uncertainty in the model estimates.




FIGURE 13: Bar chart of posterior predicted probability of states for the next 12 months (1: Tan, 2: Coral, 3: Red).

5.4.5. Remarks

We present a Bayesian framework for analyzing Mixture Transition Distribution (MTD) models in high-order Markov chains. By assessing posterior model probabilities, we determine the optimal order of dependence in the data, formulating this task as a Bayesian model selection problem where the model with the highest posterior probability serves as the selection criterion. The chosen MTD model is then compared to the fully parameterized Markov model, demonstrating superior performance in balancing goodness-of-fit and model parsimony. Furthermore, we generate near-future predictions, illustrating the practical utility of the MTD approach. Our results indicate that Bayesian MTD models can be effectively implemented using MCMC methods and provide strong performance relative to the full Markov model, making them a compelling choice for modeling high-order dependencies.

6. A Variational Algorithm for Dirichlet Process Mixtures

The mean field variational inference method provides a framework for approximating complex posterior distributions by reformulating the problem as an optimization task. In Bayesian inference, the posterior distribution $p(\boldsymbol{\theta} \mid \mathbf{y})$ often involves computationally intractable integrals, especially in high-dimensional spaces. Variational methods (Jordan et al., 1999) address this challenge by replacing the exact posterior with an approximate distribution $q_{\boldsymbol{\eta}}(\boldsymbol{\theta})$, parameterized by variational parameters $\boldsymbol{\eta}$, and optimizing the fit between the two distributions. The goal is to minimize the Kullback-Leibler (KL) divergence, which quantifies the difference between $q_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ and $p(\boldsymbol{\theta} \mid \mathbf{y})$. This divergence is defined as

$$KL(p||q) = \int_{\Theta} \log \frac{q_{\boldsymbol{\eta}}(\boldsymbol{\theta})}{p(\boldsymbol{\theta} | \mathbf{y})} q_{\boldsymbol{\eta}}(\boldsymbol{\theta}) d\boldsymbol{\theta},$$

where $\boldsymbol{\theta}$ represents latent variables within a parameter space Θ . By minimizing this divergence, the variational approach ensures that $q_{\boldsymbol{\eta}}(\boldsymbol{\theta})$ closely approximates the true posterior (Bishop, 2006).

The mean field approach, foundational to variational inference methods simplifies the problem by assuming that the posterior distribution can be factorized into independent components. Specifically, the joint posterior $p(\boldsymbol{\theta} | \mathbf{y})$ is approximated as

$$q_{\boldsymbol{\eta}}(\boldsymbol{\theta}) = \prod_{k=1}^K q_{\eta_k}(\theta_k),$$

where each $q_{\eta_k}(\theta_k)$ represents the distribution of the k -th latent variable, governed by its variational parameters. This factorization significantly reduces the computational complexity of working with the full posterior, allowing efficient optimization. Early implementations of this method in probabilistic graphical models demonstrated its scalability and practicality for high-dimensional inference (Koller & Friedman, 2009).

The algorithm iteratively updates the variational parameters η_k to refine the approximation. For each latent variable, the full conditional posterior $p(\theta_k | \boldsymbol{\theta}_{-k}, \mathbf{y})$ is used as a reference. In many practical settings, this conditional posterior belongs to the exponential family, and its form can be written as

$$p(\theta_k | \boldsymbol{\theta}_{-k}, \mathbf{y}) \propto \exp \left\{ \sum_{\ell=1}^L g_{k,\ell}(\boldsymbol{\theta}_{-k}, \mathbf{y}) t_{\ell}(\theta_k) - h(\theta_k) \right\}.$$

Here, $t_{\ell}(\theta_k)$ represents the sufficient statistics of the conditional distribution, $g_{k,\ell}(\boldsymbol{\theta}_{-k}, \mathbf{y})$ are functions capturing the dependencies on other latent variables $\boldsymbol{\theta}_{-k}$ and the observed data \mathbf{y} , and $h(\theta_k)$ is a normalizing term. This approach follows from the exponential family framework provided in Blei et al. (2017).

The optimization step involves deriving the variational distribution $q_{\eta_k}(\theta_k)$ that best approximates $p(\theta_k | \boldsymbol{\theta}_{-k}, \mathbf{y})$. This is achieved by setting the parameters $\eta_{k,\ell}$ of $q_{\eta_k}(\theta_k)$ to the expected value of the corresponding terms in the conditional posterior: $\eta_{k,\ell} = \mathbb{E}_{q_{-k}}[g_{k,\ell}(\boldsymbol{\theta}_{-k}, \mathbf{y})]$. Here, the expectation is taken with respect to q_{-k} , the variational distribution over all latent variables except θ_k . This expectation ensures that the parameters of $q_{\eta_k}(\theta_k)$ incorporate information from the rest of the model (Hoffman et al., 2013).

The iterative process of updating η_k for each latent variable continues until the KL divergence converges to a minimum, providing an optimal approximation for the posterior. This approach, by leveraging factorization and conditional independence, allows for scalable inference in high-dimensional models. Mean field variational inference thus offers a powerful and computationally efficient alternative to exact Bayesian methods, striking a balance between accuracy and feasibility in complex problems.

6.1. Dirichlet Process Mixture Model

In a Dirichlet Process Mixture Model (DPMM; e.g., [Rodriguez & Müller \(2013\)](#), [Müller & Mitra \(2013\)](#), [Müller et al. \(2018\)](#)), the Dirichlet Process (DP; [Ferguson \(1973\)](#) and [Ferguson \(1974\)](#)) serves as a prior for the mixing distribution G , enabling an unknown and potentially infinite number of mixture components. This approach is particularly effective for analyzing data with complex structures or when the number of underlying subgroups (clusters) is unknown. The model is formulated as:

$$F(\cdot | G) = \int K(\cdot | \theta) dG(\theta),$$

where $G \sim DP(\alpha, G_0)$ indicates that G is drawn from a DP with concentration parameter α and base measure G_0 , and $K(\cdot | \theta)$ is a parametric distribution defined by parameters θ . Each θ specifies the parameters for an individual component of the mixture. In this way, the corresponding mixture density or probability mass function in the DPMM can be expressed as:

$$f(\cdot | G) = \int k(\cdot | \theta) dG(\theta),$$

where $k(\cdot | \theta)$ denotes the density (or probability mass function) of $K(\cdot | \theta)$. Since G is random, the mixture density $f(\cdot | G)$ and the cumulative distribution $F(\cdot | G)$ are also random.

In a DPMM, a DP prior is placed on G , the distribution of parameters θ . The DP ensures that G is a discrete distribution, potentially consisting of infinitely many "atoms" or components. Thus, data points y_i are generated in two steps: First, a parameter θ_i is drawn from G , and second, y_i is sampled from $K(\cdot | \theta_i)$, a distribution governed by θ_i . Since G is discrete, many of the θ_i values will coincide, naturally forming groupings or clusters among the y_i values. This characteristic enables the DPMM to infer clusters in the data adaptively, even without prior knowledge of the number of components.

6.1.1. Alternative Formulation

Using the constructive definition of the DP ([Sethuraman, 1994](#)), a random distribution G drawn from a $DP(\alpha, G_0)$ can be represented as:

$$G = \sum_{\ell=1}^{\infty} p_{\ell} \delta_{Z_{\ell}}$$

where δ_a denotes a Dirac delta function centered at a . This representation, known as the stick-breaking process, provides an intuitive way to interpret G as a countable mixture of point masses, each located at parameters Z_{ℓ} with corresponding weights p_{ℓ} .

Under this framework, the probability model $f(\cdot \mid G)$ can be expressed as a countable mixture of parametric densities:

$$f(\cdot \mid G) = \sum_{\ell=1}^{\infty} p_{\ell} k(\cdot \mid Z_{\ell}),$$

where $k(\cdot \mid Z_{\ell})$ is a parametric density function parameterized by Z_{ℓ} , and p_{ℓ} are the weights generated by the stick-breaking process. The construction begins with the first weight $p_1 = v_1$, and subsequent weights are defined recursively as $p_{\ell} = v_{\ell} \prod_{r=1}^{\ell-1} (1 - v_r)$, where $v_r \sim \text{Beta}(1, \alpha)$ for $r = 1, 2, \dots$. The stick-breaking process guarantees that the weights p_1, p_2, \dots sum to 1, ensuring a valid probability distribution over the infinite components.

The locations Z_{ℓ} are independently sampled from the base measure G_0 , which serves as the prior over the parameter space and encodes prior knowledge about the data-generating process. Importantly, the sequences $\{v_r : r = 1, 2, \dots\}$ and $\{Z_{\ell} : \ell = 1, 2, \dots\}$ are independent, ensuring that the generation of weights and the selection of component locations are decoupled. This separation simplifies the modeling and enhances the interpretability of the DP-based mixture.

6.1.2. Hierarchical Formulation

Typically, DPMs are employed to model data as follows:

$$y_i \mid G, \phi \stackrel{\text{iid}}{\sim} f(\cdot \mid G, \phi) = \int k(\cdot \mid \theta, \phi) dG(\theta), \quad i = 1, \dots, n,$$

where $G \sim \text{DP}(\alpha, G_0)$ is a random distribution drawn from a DP with concentration parameter α and base measure G_0 . Then, a parametric prior $p(\phi)$ is placed on the global parameter ϕ , and hyperpriors may be specified for α or the parameters $\boldsymbol{\vartheta}$ governing $G_0 = G_0(\cdot \mid \boldsymbol{\vartheta})$, as follows:

1. Observation model: Each y_i is conditionally independent, given its associated latent parameter θ_i and the global parameter ϕ :

$$y_i \mid \theta_i, \phi \sim k(y_i \mid \theta_i, \phi), \quad i = 1, \dots, n.$$

2. Latent mixing parameters: The latent variables θ_i are independently drawn from the random distribution G :

$$\theta_i \mid G \sim G, \quad i = 1, \dots, n.$$

3. DP prior: The mixing distribution G is itself drawn from a Dirichlet Process:

$$G \mid \alpha, \boldsymbol{\vartheta} \sim \text{DP}(\alpha, G_0(\cdot \mid \boldsymbol{\vartheta})),$$

where α controls the concentration of G , and G_0 acts as the base measure, parameterized by $\boldsymbol{\vartheta}$.

4. Priors on model parameters: Independent priors are placed on the global parameter ϕ , the concentration parameter α , and the base measure parameter $\boldsymbol{\vartheta}$:

$$\phi, \alpha, \boldsymbol{\vartheta} \sim p(\phi) p(\alpha) p(\boldsymbol{\vartheta}).$$

In the first step of the previous formulation, the mixing parameters θ_i can be replaced with configuration variables L_1, L_2, \dots . Each L_i is defined such that $L_i = \ell$ if and only if $\theta_i = Z_\ell$, for $i = 1, 2, \dots$ and $\ell = 1, 2, \dots$. This representation simplifies the model by associating each observation i with one of the distinct components indexed by ℓ , offering a more concise framework for managing clustering assignments. Therefore, observation model can be rewritten as:

$$\begin{aligned} y_i \mid L_i, Z_{L_i}, \phi &\stackrel{\text{ind}}{\sim} k(y_i \mid Z_{L_i}, \phi) \\ L_i \mid \boldsymbol{p} &\stackrel{\text{iid}}{\sim} \sum_{\ell=1}^N p_\ell \delta_\ell(L_i) \\ Z_\ell \mid \boldsymbol{\vartheta} &\stackrel{\text{iid}}{\sim} G_0(\cdot \mid \boldsymbol{\vartheta}), \end{aligned}$$

for $i = 1, 2, \dots$ and $\ell = 1, 2, \dots$

6.2. Illustration

We propose a mean field variational algorithm to perform inference on a DPMM, where $y_i \mid L_i, Z_{L_i}, \phi \stackrel{\text{ind}}{\sim} \mathcal{N}(y_i \mid Z_{L_i}, 1/\phi)$, for $i = 1, \dots, n$. The model assumes a parametric prior $\phi \sim \mathcal{G}(a_\phi, b_\phi)$ on the precision parameter, a fixed concentration parameter $\alpha = 1$, and a baseline distribution $G_0 = \mathcal{N}(\psi, 1/\nu)$.

Here, a_ϕ , b_ϕ , ψ , and ν are treated as hyperparameters. While this formulation considers ψ , ν , and α as fixed, they can also be assigned additional priors to incorporate further uncertainty into the model. This flexibility allows for a fully Bayesian treatment when desired, enabling richer inference while maintaining computational tractability. In this case, the posterior distribution is given by:

$$\begin{aligned} p(\mathbf{Z}, \mathbf{L}, \mathbf{v}, \phi \mid \mathbf{y}) &\propto \prod_{\ell=1}^{\infty} v_\ell^{1-1} (1 - v_\ell)^{\alpha-1} \prod_{\ell=1}^{\infty} \exp \left\{ -\frac{1}{2} \nu (Z_\ell - \psi)^2 \right\} \\ &\cdot \prod_{\ell=1}^{\infty} \phi^{\frac{1}{2} \sum_{i=1}^n I(L_i=\ell)} \exp \left\{ -\frac{1}{2} \sum_{i=1}^n I(L_i = \ell) \phi (y_i - Z_\ell)^2 \right\} \\ &\cdot \prod_{\ell=1}^{\infty} \prod_{i=1}^n v_\ell^{I(\ell=L_i)} (1 - v_\ell)^{I(\ell < L_i)} \cdot \phi^{a_\phi-1} \exp \{-b_\phi \phi\}, \quad (2) \end{aligned}$$

where $I(\cdot)$ is the indicator function.

6.2.1. Variational Algorithm

The variational distribution for this model is designed with a truncation level N , which approximates the infinite-dimensional stick-breaking representation of the DP. To simplify the representation, boundary conditions are imposed such that $q(v_N = 1) = 1$ and $q(v_\ell = 0) = 1$ for $\ell > N$, effectively truncating the stick-breaking process beyond the N -th component.

The variational distribution is formulated as a factorized family of distributions, given by:

$$q(\phi, \mathbf{v}, \mathbf{Z}, \mathbf{L}) = q_\xi(\phi) \cdot \prod_{\ell=1}^{N-1} q_{\gamma_\ell}(v_\ell) \cdot \prod_{\ell=1}^N q_{\eta_\ell}(Z_\ell) \cdot \prod_{i=1}^n q_{\varpi_i}(L_i).$$

Under this approximation, the precision parameter ϕ follows a Gamma distribution, $q_\xi(\phi) = \mathbf{G}(\xi_1, \xi_2)$, the stick-breaking weights v_ℓ follow a Beta distribution, $q_{\gamma_\ell}(v_\ell) = \mathbf{Beta}(\gamma_{\ell,1}, \gamma_{\ell,2})$, the component-specific parameters Z_ℓ follow a Normal distribution, $q_{\eta_\ell}(Z_\ell) = \mathbf{N}(\eta_{\ell,1}, \eta_{\ell,2})$, and finally, the latent cluster assignments L_i follow a Multinomial distribution, $q_{\varpi_i}(L_i) = \mathbf{Mult}(\varpi_i)$. Here, ξ_1 , ξ_2 , $\gamma_{\ell,1}$, $\gamma_{\ell,2}$, $\eta_{\ell,1}$, $\eta_{\ell,2}$, and ϖ_i denote the variational parameters.

This variational distribution leverages conjugate forms for the posterior approximations, enabling efficient optimization of the variational parameters during inference. The independence assumption and the choice of conjugate distributions make this approach computationally feasible, even for large datasets and high truncation levels N .

The variational update rules for the parameters are outlined below:

1. Precision Parameter: The posterior conditional density is given by:

$$p(\phi \mid \dots) \propto \phi^{a_\phi + \frac{1}{2} \sum_{\ell=1}^N \sum_{i=1}^n I(L_i = \ell) - 1} \cdot \exp \left\{ -\phi \left(b_\phi + \frac{1}{2} \sum_{\ell=1}^N \sum_{i=1}^n I(L_i = \ell) (y_i - Z_\ell)^2 \right) \right\}.$$

The variational parameters are updated as:

$$\begin{aligned} \xi_1 &= a_\phi + \frac{1}{2} \sum_{i=1}^n \sum_{\ell=1}^N \mathbb{E}[I(L_i = \ell)] = a_\phi + \frac{n}{2}, \\ \xi_2 &= b_\phi + \frac{1}{2} \sum_{\ell=1}^N \sum_{i=1}^n \mathbb{E}[I(L_i = \ell)] \cdot \left[y_i^2 - 2y_i \mathbb{E}(Z_\ell) + \mathbb{E}(Z_\ell)^2 + \text{Var}(Z_\ell) \right]. \end{aligned}$$

Substituting $\mathbb{E}[I(L_i = \ell)] = \varpi_{i,\ell}$, $\mathbb{E}(Z_\ell) = \eta_{\ell,1}$, and $\text{Var}(Z_\ell) = \eta_{\ell,2}$, we have:

$$\xi_2 = b_\phi + \frac{1}{2} \sum_{\ell=1}^N \sum_{i=1}^n \varpi_{i,\ell} \left(y_i^2 - 2y_i \eta_{\ell,1} + \eta_{\ell,1}^2 + \eta_{\ell,2} \right).$$

2. Stick-breaking weights: The posterior conditional density is given by:

$$p(v_\ell | \dots) \propto v_\ell^{1+\sum_{i=1}^n I(L_i=\ell)-1} (1-v_\ell)^{\alpha+\sum_{i=1}^n I(L_i>\ell)-1}.$$

The variational parameters are updated as:

$$\begin{aligned}\gamma_{\ell,1} &= 1 + \sum_{i=1}^n \mathbb{E}[I(L_i = \ell)] = 1 + \sum_{i=1}^n \varpi_{i,\ell}, \\ \gamma_{\ell,2} &= \alpha + \sum_{i=1}^n \mathbb{E}[I(L_i > \ell)] = \alpha + \sum_{i=1}^n \sum_{j=\ell+1}^N \varpi_{i,j}.\end{aligned}$$

3. Component Parameters: The posterior conditional density is:

$$p(Z_\ell | \dots) \propto \exp \left(-\frac{1}{2} \left\{ \left(\nu + \sum_{i=1}^n I(L_i = \ell) \phi \right) Z_\ell^2 - 2 \left(\nu \psi + \sum_{i=1}^n I(L_i = \ell) \phi y_i \right) Z_\ell \right\} \right).$$

The variational parameters are updated as:

$$\begin{aligned}\eta_{\ell,1} &= \eta_{\ell,2} \cdot \left(\nu \psi + \sum_{i=1}^n \mathbb{E}[I(L_i = \ell)] \cdot \mathbb{E}(\phi) \cdot y_i \right), \\ \eta_{\ell,2} &= \left(\nu + \sum_{i=1}^n \mathbb{E}[I(L_i = \ell)] \cdot \mathbb{E}(\phi) \right)^{-1}.\end{aligned}$$

Substituting $\mathbb{E}[I(L_i = \ell)] = \varpi_{i,\ell}$ and $\mathbb{E}(\phi) = \xi_1/\xi_2$, we have:

$$\begin{aligned}\eta_{\ell,1} &= \left(\nu + \frac{\xi_1}{\xi_2} \sum_{i=1}^n \varpi_{i,\ell} \right)^{-1} \left(\nu \psi + \frac{\xi_1}{\xi_2} \sum_{i=1}^n \varpi_{i,\ell} y_i \right), \\ \eta_{\ell,2} &= \left(\nu + \frac{\xi_1}{\xi_2} \sum_{i=1}^n \varpi_{i,\ell} \right)^{-1}.\end{aligned}$$

4. Cluster Assignments: The cluster probabilities are given by:

$$p(L_i = \ell | \dots) \propto v_\ell \prod_{j=1}^{\ell-1} (1-v_j) \cdot \phi^{\frac{1}{2}} \exp \left(-\frac{1}{2} \phi (y_i - Z_\ell)^2 \right).$$

Taking expectations over the variational distributions, $\varpi_{i,\ell}$ is updated as:

$$\begin{aligned}\varpi_{i,\ell} &= \mathbb{E}[\log(v_\ell)] + \sum_{j=1}^{\ell-1} \mathbb{E}[\log(1-v_j)] + \frac{1}{2} \mathbb{E}[\log(\phi)] \\ &\quad - \frac{1}{2} \mathbb{E}(\phi) [y_i^2 - 2y_i \mathbb{E}(Z_\ell) + \mathbb{E}(Z_\ell)^2 + \text{Var}(Z_\ell)].\end{aligned}$$

Using $\mathbb{E}[\log(v_\ell)] = \Gamma'(\gamma_{\ell,1}) - \Gamma'(\gamma_{\ell,1} + \gamma_{\ell,2})$, and similar expansions:

$$\begin{aligned}\varpi_{i,\ell} &= \Gamma'(\gamma_{\ell,1}) - \Gamma'(\gamma_{\ell,1} + \gamma_{\ell,2}) + \sum_{j=1}^{\ell-1} \left[\Gamma'(\gamma_{j,2}) - \Gamma'(\gamma_{j,1} + \gamma_{j,2}) \right] \\ &\quad + \frac{1}{2} \left[\Gamma'(\xi_1) - \log(\xi_2) \right] - \frac{\xi_1}{2\xi_2} \left[y_i^2 - 2y_i\eta_{\ell,1} + \eta_{\ell,1}^2 + \eta_{\ell,2} \right].\end{aligned}$$

The convergence of the variational algorithm is monitored using a threshold for the infinity norm of the variational parameter differences, which is defined as the maximum absolute difference between corresponding elements of the variational parameter vector across successive iterations of the algorithm. This criterion ensures that the updates to the parameters stabilize, allowing the algorithm to reach a reliable approximation. A threshold of 10^{-5} is employed to signify a high level of precision in the iterative optimization process, providing a clear indication of when the algorithm has effectively converged.

The predictive density quantifies the likelihood of a new observation y_{n+1} given the observed data \mathbf{y} . It is approximated by summing over all possible components of the model, weighted by their probabilities under the variational posterior:

$$p(y_{n+1} | \mathbf{y}) \approx \sum_{\ell=1}^N \mathbb{E}_q\{p_\ell\} \mathbb{E}_q(N(y_{n+1} | Z_\ell, \phi)).$$

Here, $\mathbb{E}_q\{p_\ell\}$ represents the expected contribution of each component, derived from the stick-breaking construction:

$$\mathbb{E}_q\{p_\ell\} = \frac{\gamma_{\ell,1}}{\gamma_{\ell,1} + \gamma_{\ell,2}} \left(\prod_{j=1}^{\ell-1} \frac{\gamma_{j,2}}{\gamma_{j,1} + \gamma_{j,2}} \right).$$

This formulation incorporates the uncertainty in the component assignments, providing a probabilistic forecast for new data while accounting for the underlying structure of the model.

The variational bound of the log marginal likelihood serves as a measure of how well the variational distribution approximates the true posterior. This bound is expressed as:

$$\begin{aligned}\log p(\mathbf{y} | \boldsymbol{\theta}, \boldsymbol{\eta}) &\geq \mathbb{E}_q[\log p(\phi | a_\phi, b_\phi)] + \mathbb{E}_q[\log p(\mathbf{v} | \alpha)] \\ &\quad + \sum_{\ell=1}^N \mathbb{E}_q[\log p(Z_\ell | \psi, \nu)] \\ &\quad + \sum_{i=1}^n \left(\mathbb{E}_q[\log p(L_i | \mathbf{v})] + \mathbb{E}_q[\log p(y_i | Z_{L_i}, \phi)] \right) \\ &\quad - \mathbb{E}_q[\log q(\phi, \mathbf{v}, \mathbf{L}, \mathbf{Z})].\end{aligned}$$

Maximizing this bound minimizes the KL divergence between the variational approximation and the true posterior, ensuring an optimal fit. Each term in the

bound captures a specific contribution: Prior distributions for the latent variables and parameters, assignments of latent variables to observed data, the likelihood of the observed data, and the complexity of the variational posterior. This decomposition provides valuable insight into the trade-offs involved in the approximation process, guiding the optimization to a balance between computational feasibility and the accuracy of the posterior estimate.

6.2.2. Data, Priors, and Performance Assessment

The dataset used in this analysis is the **galaxies** dataset from the **MASS** library (Venables & Ripley, 2002) in **R**, consisting of measurements of the velocities of 82 galaxies. These velocities are recorded in kilometers per second (km/s) and are corrected for the expansion of the universe (recessional velocities). Pre-processing involves standardizing the data to ensure a mean of zero ($\bar{y} = 0$) and a standard deviation of one ($s_y = 1$).

For prior specification, a non-informative prior is used for the location parameter, with $\psi = 0$. The scale parameter is determined based on the decomposition of the sample variance. Specifically, the expected variance of the observations is given by:

$$\text{Var}(y) = \text{E}[\text{Var}(y \mid \theta)] + \text{Var}[\text{E}(y \mid \theta)],$$

where $y \mid \theta \sim \text{N}(\theta, 1/\phi)$. These terms are given by $\text{E}(1/\phi) + \text{Var}(\theta)$, respectively, with $\phi \sim \text{G}(a_\phi, b_\phi)$ and $\theta \sim \text{N}(\psi, 1/\nu)$. Substituting these distributions, the decomposition becomes:

$$\text{E}[\text{Var}(y)] = \frac{b_\phi}{a_\phi - 1} + \nu^{-1}.$$

To encourage smoothness in the estimates, small values of $1/\phi$ are prioritized. An allocation ratio of 1 : 7 is used between observation-level and system-level variances. The hyperparameters are set as follows:

$$a_\phi = 1.5, \quad b_\phi = 0.5 \cdot \frac{1}{8}, \quad \nu = \left(\frac{7}{8}\right)^{-1}.$$

We also examined the sensitivity of the results to this weakly informative prior on the scale. Specifically, we considered alternative specifications obtained by halving and doubling a_ϕ and b_ϕ , as well as modestly increasing and decreasing ν so that the prior contribution of the system-level variance ranged between roughly 20% and 50% of the sample variance. Across these configurations, the posterior number of effectively occupied mixture components, the mixture weights, and the posterior predictive density were very similar to those reported below (results not shown), suggesting that the inferences are driven primarily by the likelihood rather than by a particular choice of hyperparameters. For benchmarking, a collapsed Gibbs sampler (e.g., Rodriguez & Müller, 2013) is implemented. The sampler is configured with 1 000 burn-in iterations, followed by 10 000 posterior samples to ensure reliable inference.

6.2.3. Results

Different initial settings were tested to evaluate their impact on the variational algorithm's performance (Table 8 and Figure 14). Among the scenarios considered, assigning data into three initial clusters resulted in the optimal variational bound, indicating that this configuration provides the best initialization for the algorithm.

TABLE 8: Variational approximation with different initial settings.

Scenario	Initial Setting	# of Iterations to Convergence	Running Time (sec)	Lower Bound Log-Prob.
1	20 initial clusters	101	6.88	-268.1
2	20 initial clusters	216	14.85	-275.6
3	4 initial clusters	65	4.81	-274.0
4	3 initial clusters	11	0.84	-245.7




FIGURE 14: Variational approximations with different initial settings.

The running time comparison highlights the efficiency of the variational algorithm (VA) compared to the collapsed Gibbs sampler (CGS). The variational algorithm completed in 0.84 seconds, whereas the collapsed Gibbs sampler required 637.8 seconds, demonstrating a significant improvement in computational speed, particularly on a 10^3 -scale.

For the variational algorithm, posterior samples were obtained using initial setting 4. The results include the predictive density and the heatmap of the incidence matrix (a matrix of co-clustering probabilities indicating the probability

that two data points belong to the same cluster), as shown in 15. Subfigure (a) presents the predictive density, while Subfigure (b) displays the heatmap of the incidence matrix, revealing the clustering structure inferred by the algorithm.

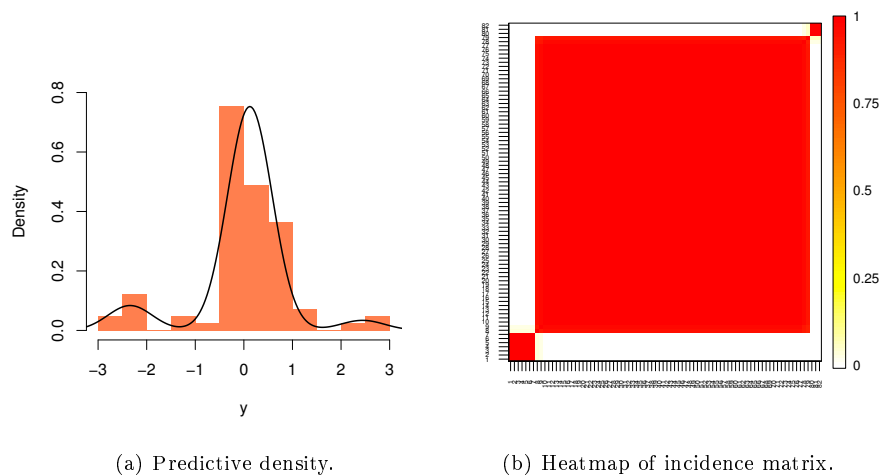


FIGURE 15: Posterior samples from the variational algorithm (initial setting 4).

For the collapsed Gibbs sampler, posterior samples were similarly obtained. Figure 16 presents the corresponding results: Subfigure (a) shows the predictive density, and Subfigure (b) displays the heatmap of the incidence matrix. These visualizations demonstrate the clustering structure identified by the Gibbs sampler. Standard MCMC convergence diagnostics (including trace plots, potential scale reduction factors, and effective sample sizes for all monitored parameters) indicated good mixing and did not reveal any evidence of convergence problems.

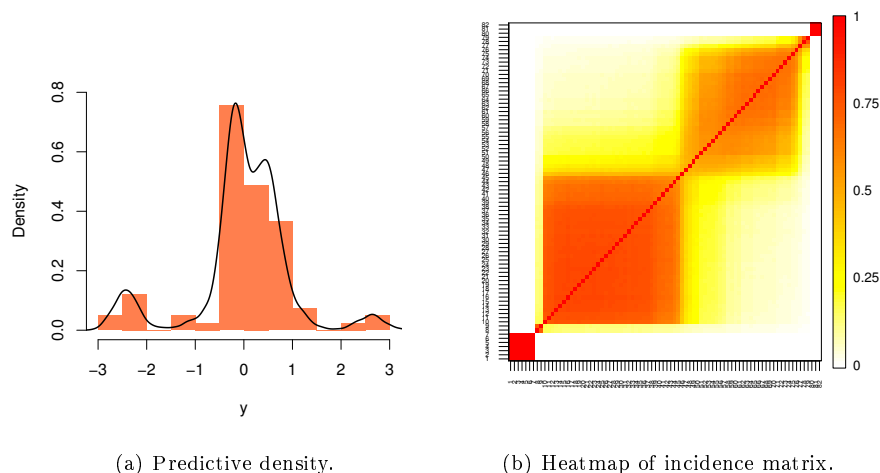


FIGURE 16: Posterior samples from the collapsed Gibbs sampler.

6.2.4. Remarks

We compare the VA and the CGS based on their performance in fitting the data. Both methods demonstrated a reasonable fit, but their behavior differed in several key aspects. In terms of computational efficiency, the VA outperformed the CGS, with a running time difference on the order of 10^3 , highlighting the VA's much faster performance. However, the covariance pattern between parameters, which was captured effectively by the CGS, was lost in the VA. Additionally, the CGS exhibited sensitivity to prior specification, while the VA was relatively insensitive, making it more robust to prior choices.

The VA offers several advantages that make it a compelling choice for large-scale or complex problems. It is highly computationally efficient, with truncation occurring at the approximation level (variational distribution) rather than at the system level (prior specification). This property contributes to its robustness against prior choices and its rapid and quantifiable convergence. Moreover, the performance of the VA is quantifiable using metrics such as the KL divergence, providing insights into the quality of the approximation. Its iterative approximation approach is easier to implement compared to sampling-based methods like MCMC.

However, the VA has notable drawbacks. It loses covariance information between parameters, which may be critical in certain applications. It is highly sensitive to initial settings, with a strong tendency toward local convergence, which can affect the quality of the inference. Additionally, its inference accuracy is generally lower than that of sampling-based methods. Furthermore, the methodology for handling parameters outside the exponential family is unclear, limiting its generalizability in some contexts. The most challenging aspect of implementing the VA is specifying the initial settings. Since the algorithm's performance and convergence depend heavily on the initialization, careful consideration is required to achieve reliable results.

7. Discussion

The results presented in this paper highlight the broad utility of Bayesian inference across diverse domains, ranging from hierarchical and spatial modeling to non-parametric approaches such as Dirichlet process mixtures and mixtures of Brownian motion. By leveraging both Markov chain Monte Carlo (MCMC) methods and variational approximations, this work demonstrates the dual strengths of Bayesian inference: its capacity for accurate, detailed posterior analysis and its adaptability to large-scale problems requiring efficient computation.

MCMC methods remain the gold standard for Bayesian computation due to their ability to accurately sample from complex posterior distributions. Applications in spatial modeling and higher-order Markov chains particularly benefit from the flexibility of MCMC to incorporate intricate dependencies and hierarchical structures. However, their high computational cost and sensitivity to tuning parameters pose challenges, especially for high-dimensional or large-scale datasets.

On the other hand, variational approximations provide a complementary approach, offering rapid convergence and scalability by reformulating posterior inference as an optimization problem. Although this efficiency comes at the cost of some approximation accuracy, it makes variational methods indispensable for applications where computational constraints are paramount.

The comparative analysis of MCMC and variational methods underscores their trade-offs. While MCMC excels in capturing complex covariance structures and delivering high-fidelity posterior estimates, it often requires significant computational resources and careful tuning. In contrast, variational methods are less sensitive to priors and achieve rapid convergence, but they may oversimplify posterior distributions by neglecting parameter correlations. These differences highlight the importance of method selection based on the specific requirements of the problem, including the dataset size, model complexity, and computational resources available.

From a practical standpoint, we view MCMC and variational inference (VI) as complementary tools rather than direct competitors. VI is particularly attractive in settings with very large sample sizes and/or high-dimensional parameter spaces, or when approximate posterior summaries are primarily required for exploratory analysis or as an initialization strategy for subsequent MCMC runs. In such cases, the ability of VI to deliver solutions in a fraction of the computational time can outweigh the loss in accuracy. By contrast, when the scientific goal demands accurate uncertainty quantification for complex posteriors (for example, in the presence of strong parameter dependence, multimodality, or intricate hierarchical structures) MCMC remains the preferred choice despite its higher computational cost and tuning requirements. It is also important to note common caveats: standard ELBO-based mean-field VI often underestimates posterior variances because the optimization penalizes over-dispersion more strongly than under-dispersion, and the method can be sensitive to local optima and initialization. In practice, these issues motivate combining VI with targeted MCMC checks on lower-dimensional submodels or subsets of the data when accurate assessment of uncertainty is critical.

The applications explored in this paper emphasize the versatility of Bayesian methods across scientific fields. For example, spatial modeling in climatology and epidemiology underscores the power of Bayesian inference to capture geographic dependencies and uncertainty, while hierarchical models in oceanography and finance demonstrate its ability to handle nested and multi-level data structures. Non-parametric methods such as Dirichlet process mixtures reveal Bayesian inference's flexibility in modeling infinite-dimensional parameter spaces, making it suitable for adaptive clustering and density estimation.

Despite their strengths, Bayesian methods face ongoing challenges. The sensitivity of MCMC methods to initialization and tuning parameters can limit their accessibility to non-expert users. Similarly, the reliance of variational approximations on strong assumptions about the posterior distribution can lead to biased inference in cases involving highly complex models. Future advancements should focus on developing hybrid methods that combine the strengths of MCMC and variational inference, such as stochastic variational techniques, to improve both accuracy and scalability.

It is important to acknowledge several well-known limitations of the methods discussed here and to point to standard remedies. On the MCMC side, Markov chain algorithms may suffer from slow convergence, poor mixing, and high computational cost in high-dimensional or strongly correlated posteriors, and their performance can be sensitive to tuning choices and parameterization. In practice, these issues are mitigated by reparameterization (e.g., non-centered parameterizations in hierarchical models), gradient-based or adaptive samplers, tempering strategies, and systematic use of convergence diagnostics and effective sample size calculations. Variational inference, in turn, trades exactness for speed and can introduce approximation bias, particularly under mean-field assumptions, often underestimating posterior uncertainty and sometimes being sensitive to initialization. Standard remedies include using richer variational families, tempering or importance-sampling-based corrections, and targeted comparison with MCMC on lower-dimensional submodels or subsets of the data when accurate uncertainty quantification is critical. More broadly, both paradigms remain subject to the usual limitations of model misspecification, for which posterior predictive checks, prior-sensitivity analyses, and the consideration of alternative model structures provide established tools to assess robustness.

Looking ahead, Bayesian inference is poised to benefit from integration with modern computational tools, including parallel processing, machine learning, and GPU-based computation, which promise to enhance scalability and reduce runtime. Additionally, developing methodologies to extend Bayesian inference to more complex, non-standard problems will further expand its applicability. This work aims to bridge the gap between theory and practice, providing a foundation for researchers to explore Bayesian methods across a diverse range of applications and challenges.

Statements and Declarations

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this article.

During the preparation of this work the authors used ChatGPT-4-turbo in order to improve language and readability. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

[Received: September 2025 — Accepted: November 2025]

References

- Albert, J. H. & Chib, S. (1993), ‘Bayesian analysis of binary and polychotomous response data’, *Journal of the American Statistical Association* **88**(422), 669–679.

- Banerjee, S., Carlin, B. P. & Gelfand, A. E. (2015), *Hierarchical Modeling and Analysis for Spatial Data*, 2nd edn, CRC Press, Boca Raton, FL.
- Berchtold, A. & Raftery, A. E. (2002), 'The mixture transition distribution model for high-order Markov chains and non-Gaussian time series', *Statistical Science* pp. 328–356.
- Bernardo, J. M. & Smith, A. F. (2009), *Bayesian theory*, Vol. 405, John Wiley & Sons.
- Bishop, C. M. (2006), *Pattern Recognition and Machine Learning*, Springer, New York, NY.
- Blei, D. M., Kucukelbir, A. & McAuliffe, J. D. (2017), 'Variational inference: A review for statisticians', *Journal of the American Statistical Association* **112**(518), 859–877.
- Cox, R. T. (1946), 'Probability, frequency and reasonable expectation', *American journal of physics* **14**(1), 1–13.
- De Finetti, B. (1970), 'Logical foundations and measurement of subjective probability', *Acta Psychologica* **34**, 129–145.
- Ferguson, T. S. (1973), 'A Bayesian analysis of some nonparametric problems', *The annals of statistics* pp. 209–230.
- Ferguson, T. S. (1974), 'Prior distributions on spaces of probability measures', *The annals of statistics* pp. 615–629.
- Gamerman, D. & Lopes, H. F. (2006), *Markov chain Monte Carlo: stochastic simulation for Bayesian inference*, Chapman and Hall/CRC.
- Gelfand, A. E. & Ghosh, S. K. (1998), 'Model choice: a minimum posterior predictive loss approach', *Biometrika* **85**(1), 1–11.
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A. & Rubin, D. B. (2013), *Bayesian Data Analysis*, 3rd edn, Chapman and Hall/CRC, Boca Raton, FL.
- Ghosh, J. K., Delampady, M. & Samanta, T. (2007), *An introduction to Bayesian analysis: theory and methods*, Springer Science & Business Media.
- Ghosh, M. (2011), 'Objective priors: An introduction for frequentists', *Statistical Science* **26**(2), 187–202.
- Golub, G. H. & Van Loan, C. F. (2013), *Matrix computations*, JHU press.
- Hoffman, M. D., Blei, D. M., Wang, C. & Paisley, J. (2013), 'Stochastic variational inference', *Journal of Machine Learning Research* **14**, 1303–1347.
- Joos, F., Plattner, G.-K., Stocker, T. F., Marchal, O. & Schmittner, A. (1999), 'Global warming and marine carbon cycle feedbacks on future atmospheric co2', *Science* **284**(5413), 464–467.

- Jordan, M. I., Ghahramani, Z., Jaakkola, T. S. & Saul, L. K. (1999), ‘An introduction to variational methods for graphical models’, *Machine Learning* **37**(2), 183–233.
- Koller, D. & Friedman, N. (2009), *Probabilistic Graphical Models: Principles and Techniques*, MIT Press, Cambridge, MA.
- Müeller, P., Quintana, F. A. & Page, G. (2018), ‘Nonparametric Bayesian inference in applications’, *Statistical Methods & Applications* **27**, 175–206.
- Müller, P. & Mitra, R. (2013), ‘Bayesian nonparametric inference—why and how’, *Bayesian analysis (Online)* **8**(2).
- Nykjaer, L. (2009), ‘Mediterranean sea surface warming 1985–2006’, *Climate Research* **39**(1), 11–17.
- Raftery, A. E. (1985), ‘A model for high-order Markov chains’, *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 528–539.
- Richardson, S. & Green, P. J. (1997), ‘On Bayesian analysis of mixtures with an unknown number of components (with discussion)’, *Journal of the Royal Statistical Society: series B (statistical methodology)* **59**(4), 731–792.
- Rodriguez, A. & Müller, P. (2013), Nonparametric Bayesian inference, in ‘NSF-CBMS Regional Conference Series in Probability and Statistics’, Vol. 9, JSTOR, pp. i–110.
- Sarmiento, L. (2000), ‘Water balance and soil loss under long fallow agriculture in the Venezuelan Andes’, *Mountain Research and Development* **20**(3), 246–253.
- Savage, L. J. (1972), *The foundations of statistics*, Courier Corporation.
- Scott, S. L. (2002), ‘Bayesian methods for hidden Markov models’, *Journal of the American Statistical Association* **97**(457).
- Sethuraman, J. (1994), ‘A constructive definition of dirichlet priors’, *Statistica sinica* pp. 639–650.
- Venables, W. N. & Ripley, B. D. (2002), *Modern Applied Statistics with S*, fourth edn, Springer, New York.

Appendix. Notation

The cardinality of a set A is denoted by $|A|$. If P is a logical proposition, then $I(P) = 1$ if P is true and $1_P = 0$ if P is false. The Gamma function is defined as $\Gamma(x) = \int_0^\infty u^{x-1} e^{-u} du$. Matrices and vectors whose entries consist of subscripted variables are represented using bold notation. For example, $\mathbf{x} = (x_1, \dots, x_n)$ denotes an $n \times 1$ column vector with elements x_1, \dots, x_n . We use $\mathbf{0}$ and $\mathbf{1}$ to denote column vectors with all entries equal to 0 and 1, respectively, and \mathbf{I} to represent the

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