

Scalable Bayesian computation for crossed and nested hierarchical models

Omiros Papaspiliopoulos

Bocconi University, Milan 20100, Italy
e-mail: omiros@unibocconi.it

Timothée Stumpf-Fétizon

Bocconi Institute for Data Science, Bocconi University, Milan 20100, Italy
e-mail: timothee.stumpffetizon@unibocconi.it

Giacomo Zanella

Department of Decision Sciences, Bocconi University, Milan 20100, Italy
e-mail: giacomo.zanella@unibocconi.it

Abstract: We develop sampling algorithms to fit Bayesian hierarchical models, the computational complexity of which scales linearly with the number of observations and the number of parameters in the model. We focus on crossed random effect and nested multilevel models, which are used ubiquitously in applied sciences. The posterior dependence in both classes is sparse: in crossed random effects models it resembles a random graph, whereas in nested multilevel models it is tree-structured. For each class we identify a framework for scalable computation, building on previous work. Methods for crossed models are based on extensions of appropriately designed collapsed Gibbs samplers, where we introduce the idea of *local centering*; while methods for nested models are based on sparse linear algebra (SLA) and data augmentation. We provide a theoretical analysis of the proposed algorithms in some simplified settings, including a comparison with previously proposed methodologies and an average-case analysis based on random graph theory. Numerical experiments, including two challenging real data analyses on predicting electoral results and real estate prices, compare with off-the-shelf Hamiltonian Monte Carlo, displaying drastic improvement in performance. The code for replicating the experiments in the article and for implementing our methods can be found at <https://github.com/timsf/crossed-effects> and <https://github.com/timsf/nested-effects>.

Keywords and phrases: Gibbs sampler, sparse linear algebra, reparameterisation, random graphs, crossed random effects, multilevel models, Polygamma.

Received December 2022.

1. Motivation and objectives

In this article we explore efficient computational frameworks to perform Bayesian inference in large-scale hierarchical models. We focus on two canonical hierarchical structures, which are arguably the two most common in applied Statistics;

crossed effects models, where the data is organized in a multidimensional contingency table and the within-cell distribution is modeled in terms of a linear expansion of a priori independent effects; and nested multilevel models, where data is organized in hierarchical clusters and parameters in a given level of the hierarchy are shrunk to (fewer) parameters in the deeper level. Crossed effects models are the canonical framework for modeling the dependence of an output variable on a number of categorical input variables. In the literature they appear under various names, e.g. cross-classified data, variance component models or multiway analysis of variance [12, 31, 37]. Nested multilevel models are the basic hierarchical structure for data and parameters that are naturally organized in nested clusters [13]. The detailed specifications of these models are given in Sections 2 and 3, respectively.

The number of parameters in these models will be denoted by p and the number of observations by n . The computational challenge is to sample from the corresponding posterior distributions in high-dimensional scenarios where both n and p are large. The “holy grail” is MCMC algorithms whose total computational cost scales linearly in $\max\{n, p\}$. We carefully define the notion of computational cost later in the article. It involves both the cost per iteration and the number of iterations needed. We call algorithms whose cost is linear in $\max\{n, p\}$, *scalable*.

Both structures, crossed and nested, involve a high-dimensional vector of regression parameters, and a low-dimensional vector of hyperparameters. Priors for the regression parameters are usually Gaussian while the data likelihood is generic, and in this article we consider methods and numerics for a wide variety of such choices. The class of sampling algorithms we consider are (variations or approximations of) block-updating coordinate-wise schemes – a.k.a. *blocked Gibbs samplers* – that alternate updating the regression parameters given the hyperparameters, and vice-versa. The main computational challenge, which is the focus of this paper, is how to update the high-dimensional regression parameters in a scalable way.

The paper is structured as follows. Section 2 is devoted to crossed effect models. After describing the models and the high-level computational approach, in Section 2.2 we propose a methodology for scalable sampling of the high-dimensional regression parameters in crossed models with generic likelihood, using what we call a *local centering* approach. In Section 2.3, we consider the practically important special case of multicategorical logistic likelihood where, beyond applying local centering ideas, we introduce a novel scheme for sampling efficiently without imposing identifiability constraints. These can instead be imposed, if desired, at a post-sampling stage. In Section 2.4, in the specific context of binomial likelihoods, we introduce an alternative scheme that combines the collapsed Gibbs sampler recently proposed in [22] with Poly-Gamma data augmentation [25]. Finally, Section 2.5 provides a theoretical analysis of the complexity of the above schemes, where we establish scalability under certain assumptions. In Section 3 we consider nested multilevel models. We first show, in Section 3.1, how to use sparse linear algebra methods to update the high-dimensional regression coefficients for Gaussian likelihoods in a scalable

way (which is not directly possible for crossed models). Then, in Section 3.2, we combine this methodology with data augmentation approaches to obtain scalable samplers for non-Gaussian likelihoods, such as binomial, Cauchy or Poisson.

The article contains various numerical experiments, including both simulated data and illustrations on two challenging large-scale applications, one on electoral survey data and one on real-estate modeling. In all cases, we compare our proposed methods to state of the art alternatives, including Hamiltonian Monte Carlo, showing superior performance, often by more than an order of magnitude. Our numerical experiments also suggest that the conclusions from our theory extrapolate beyond the specific assumptions required for the theoretical analysis. Throughout the article, we use boldface letters to denote vectors (when lower case) and matrices (when upper case).

2. Crossed effects models

We consider the following class of crossed effects models. There are K categorical input variables, known as factors, each with p_k categories, known as levels, and their numbering is arbitrary. In this context, $p = \sum_{k=1}^K p_k$ corresponds to the sum of levels over all factors. The combination of the levels of the different factors defines a $p_1 \times \dots \times p_k$ contingency table. For each observation $j \in \{1, \dots, n\}$ we record the values of the categorical input variables and an output variable that could be multivariate and will be denoted by \mathbf{y}_j , with L denoting the dimension of the response. Typical examples for $L = 1$ are continuous, binary or count observations y_j , and multicategorical observations \mathbf{y}_j for $L > 1$, with details following below. The distribution of \mathbf{y}_j depends on a linear predictor $\boldsymbol{\eta}_j$:

$$\boldsymbol{\eta}_j = \mathbf{X}_j \boldsymbol{\beta} + \mathbf{a}^{(0)} + \sum_{k=1}^K \mathbf{a}_{i_k[j]}^{(k)}$$

where \mathbf{X}_j are continuous input variables (excluding an intercept), $\boldsymbol{\beta}$ and $\mathbf{a}^{(0)}$ are typically called fixed effects, and associated to each level of every categorical factor there is a random effect. The notation we use, which is quite standard in some of the relevant literature (see, e.g., Section 1.1 and Chapter 11 of [13]), is that $i_k[j]$ refers to the level that the j th observation has on the k th factor. Therefore, the available data per “individual” j (in our election survey example, a survey respondent) are the output data \mathbf{y}_j (e.g., a 1-hot encoded vector for the party they plan to vote for), the strata $i_k[j]$ they belong on each categorical factor k (e.g., reside in Catalunya, unemployed, higher education etc), and other, if any, continuous inputs \mathbf{X}_j . In our notation, $\mathbf{a}_i^{(k)}$ is the random effect associated to the level i of the k th factor, and its dimension is L , just as for the observations. We model all random effects as a priori independent with factor-dependent variances. The main reason why we allow for multivariate random effects in this article is to deal with multi-categorical responses, as in the election survey application, in which case there are precision matrices describing the dependence between the elements of each multivariate random effect.

At this level of generality, with \mathcal{L} denoting laws of random variables, the crossed effects model is as follows:

$$\begin{aligned}\mathcal{L}(\mathbf{a}^{(0)}) &= \mathcal{N}(\boldsymbol{\mu}_{pr}, \mathbf{T}_{pr}^{-1}), \\ \mathcal{L}(\mathbf{T}_k) &= \text{Wishart}(\nu_k, \mathbf{I}/\nu_k), \\ \mathcal{L}(\mathbf{a}^{(k)} | \mathbf{T}_k) &= \mathcal{N}(\mathbf{0}, (\mathbf{T}_k \otimes \mathbf{I})^{-1}), \quad k = 1, \dots, K, \\ \mathcal{L}(\mathbf{y}_j | \cdot) &= \mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j), \\ \boldsymbol{\eta}_j &= \mathbf{X}_j \boldsymbol{\beta} + \mathbf{a}^{(0)} + \sum_{k=1}^K \mathbf{a}_{i_k[j]}^{(k)}, \quad j = 1, \dots, n,\end{aligned}\tag{1}$$

where \mathcal{N} denotes Gaussian distributions, $\mathcal{L}(\mathbf{y}_j | \cdot)$ denotes the law of \mathbf{y}_j given all other variables, \otimes stands for the Kronecker product, \mathbf{T}_k are $L \times L$ precision matrices, $\mathbf{a}^{(k)}$ is the vector of all stacked $\mathbf{a}_i^{(k)}$'s, and the \mathbf{I} 's are identity matrices of appropriate dimension. We understand the limiting case $\boldsymbol{\mu}_{pr} = \mathbf{0}, \mathbf{T}_{pr} = \mathbf{0}\mathbf{0}^T$ as the improper flat prior on $\mathbf{a}^{(0)}$. When $L = 1$ we denote the corresponding precision scalars by τ_k and assign them gamma priors. The ‘‘fixed effects’’ coefficients $\boldsymbol{\beta}$ are typically assigned a Gaussian or flat prior.

In our numerical illustrations we focus on three important special cases: Gaussian models, where $L = 1$ and $\mathcal{L}(y_j | \eta_j, \tau) = \mathcal{N}(\eta_j, 1/\tau)$ (with a gamma prior on τ); binomial models, where $L = 1$ and $\mathcal{L}(y_j | \eta_j) = \text{Binomial}(m_j, (1 + \exp\{-\eta_j\})^{-1})$; and multicategorical logit models, where $L > 1$, \mathbf{y}_j is an L -dimensional vector of zeros and a single one (one-hot encoding), and $\mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j) = \text{Categorical}(\text{softmax}\{\boldsymbol{\eta}_j\})$, with

$$\text{softmax}\{\boldsymbol{\eta}\} = \left(\frac{e^{\eta_1}}{\sum_l e^{\eta_l}}, \dots, \frac{e^{\eta_L}}{\sum_l e^{\eta_l}} \right).$$

The multicategorical logistic model introduces the further complication of being ill-identified due to softmax being invariant under translations of $\boldsymbol{\eta}$. We address this issue in Section 2.3, together with other methodological concerns specific to this model.

2.1. Computational strategy

For simplicity we focus on the case of no fixed effects ($\boldsymbol{\beta} = \mathbf{0}$), but our methodology extends easily to models that include fixed effects, as we discuss below. In this context the regression-type parameters are $\mathbf{a} = (\mathbf{a}^{(0)}, \mathbf{a}^{(1)}, \dots, \mathbf{a}^{(K)})$. The vector $\boldsymbol{\gamma}$ includes the precision matrices \mathbf{T}_k as well as any parameters specific to the likelihood $\mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j)$. The methodology developed in the following sections is about sampling from $\mathcal{L}(\mathbf{a} | \mathbf{y}, \boldsymbol{\gamma})$, where $\mathbf{y} = (\mathbf{y}_j)_{j=1, \dots, n}$ denotes all the observed data. We augment such methods with steps that update $\mathcal{L}(\boldsymbol{\gamma} | \mathbf{y}, \mathbf{a})$, which in our context is easily done by sampling from gamma and (low-dimensional) Wishart distributions to obtain a block coordinate-wise

sampling algorithm for $\mathcal{L}(\mathbf{a}, \boldsymbol{\gamma} \mid \mathbf{y})$. Potential issues related to dependence between \mathbf{a} and $\boldsymbol{\gamma}$ in the joint posterior distribution $\mathcal{L}(\mathbf{a}, \boldsymbol{\gamma} \mid \mathbf{y})$ are discussed in the numerical parts and in the discussion on future work in Section 4.

The methods we discuss are generally based on partitioning the regression parameters into $K + 1$ blocks, $(\mathbf{a}^{(0)}, \mathbf{a}^{(1)}, \dots, \mathbf{a}^{(K)})$. Previous work by [8, 22] has proved that the complexity of the vanilla Gibbs sampler that updates each $\mathbf{a}^{(k)}$ conditionally on the rest ($k = 0, 1, \dots, K$) is typically super-linear in n and p . [22] developed an approach based on a joint update of $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)})$ for each $k = 1, \dots, K$, which under certain conditions is provably scalable. Their methodology, which they call collapsed Gibbs sampling, only applies to Gaussian likelihoods. Below we develop methodologies that extend these ideas to the case of general likelihoods. In some cases one can take advantage of data augmentation techniques, such as Polya-gamma data augmentation for binomial likelihoods, to restore conditional Gaussianity and directly apply the collapsed Gibbs sampler. We explore this in Section 2.4. On the other hand, in Section 2.2 we develop a methodology, called *local centering*, that works with any (sufficiently smooth) likelihood, while inheriting the properties of the collapsed Gibbs sampler both in terms of convergence speed and cost per iteration. Note that, for Gaussian likelihoods, one could also perform a direct block update of the high-dimensional multivariate Gaussian distribution $\mathcal{L}(\mathbf{a} \mid \mathbf{y}, \boldsymbol{\gamma})$. However, the cost of this operation would be super-linear in $\max\{n, p\}$ for crossed effect models and thus not scalable. See also Section 2.6 and related discussions in [8, 19, 22, 9, 14].

Regarding the inclusion of fixed effects $\boldsymbol{\beta}$ in the model, we can adapt the computational approach we develop here in various ways. The simplest strategy would be to update the global coefficients $\boldsymbol{\beta}$ from their full conditional distribution through a Gibbs or Metropolis-within-Gibbs update, following the update of \mathbf{a} detailed below. Regardless of the specific implementation, the key aspect is that the dimensionality of $\boldsymbol{\beta}$ is much lower than the dimensionality of \mathbf{a} in our applications of interest. More robust strategies are possible, such as updating $\boldsymbol{\beta}$ jointly with the intercept $\mathbf{a}^{(0)}$, but we defer the discussion of such strategies, as well as of the discussion of the case of high-dimensional $\boldsymbol{\beta}$, to future work.

2.2. Methodology for generic non-Gaussian likelihoods: local centering

In the same spirit as the collapsed Gibbs sampler of [22] for Gaussian likelihoods, we develop a scheme for the joint update of $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)})$ that is feasible for generic likelihoods. To put our proposal into perspective, we first note that off-the-shelf sampling methods will not be successful. When factor k has a large number of levels, this will be a high-dimensional sampling problem. Additionally, there is strong posterior dependence between $\mathbf{a}^{(0)}$ and $\mathbf{a}^{(k)}$, due to the constraint imposed by the data via the linear predictor. The prior precision does not capture this dependence in crossed effects models, since a priori all variables are independent. Thus one cannot use gradient-based samplers for latent Gaussian models

that precondition using the prior precision, as for example in [33]. Moreover, even with hypothetically good preconditioning, the convergence time of standard gradient-based samplers targeting the high-dimensional joint distribution of $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)})$ will typically grow with p_k (see e.g. [26, 7, 36, 40] and related literature for various analysis of the dimensionality dependence of gradient-based MCMC algorithms), resulting in algorithms with super-linear complexity w.r.t. n and p (see Section 2.5 for proper definitions on convergence).

Instead, we explicitly leverage the sparse conditional independence structure of $\mathcal{L}(\mathbf{a}^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(-0, -k)})$, where $\mathbf{a}^{(-0, -k)}$ is a short form of $(\mathbf{a}^{(\ell)})_{\ell \neq 0, k}$. We propose to use local hierarchical centering within each block:

$$(\mathbf{a}^{(0)}, \mathbf{a}^{(k)}) \rightarrow (\mathbf{a}^{(0)}, \boldsymbol{\xi}^{(k)}), \quad \boldsymbol{\xi}_i^{(k)} = \mathbf{a}^{(0)} + \mathbf{a}_i^{(k)}. \quad (2)$$

Working locally with a centered parameterisation is motivated by the general principles of parameterisation of hierarchical models: as discussed in [10, 23], the model correlation structure is such that the conditional posterior dependence is weaker between the centered $(\mathbf{a}^{(0)}, \boldsymbol{\xi}^{(k)})$ than the non-centered $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)})$ pair. Now, note that due to the hierarchical structure, we obtain that $\mathcal{L}(\mathbf{a}^{(0)} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(-0, -k)}, \boldsymbol{\xi}^{(k)}) = \mathcal{L}(\mathbf{a}^{(0)} \mid \boldsymbol{\xi}^{(k)})$. Additionally, the $\boldsymbol{\xi}_i^{(k)}$ are independent across i conditionally on $\mathbf{a}^{(0)}$ and the levels of the other factors, hence a high-dimensional sampling problem reduces to p_k independent L -dimensional ones. With these observations made, we propose Algorithm 1.

Algorithm 1 Gibbs sampler with local centering for non-Gaussian likelihoods

for $k = 1 : K$ **do**

 Reparameterise $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)}) \rightarrow (\mathbf{a}^{(0)}, \boldsymbol{\xi}^{(k)})$ according to (2)

 Draw $\mathbf{a}^{(0)}$ from

$$\mathcal{L}(\mathbf{a}^{(0)} \mid \boldsymbol{\xi}^{(k)}) = \mathcal{N} \left(\begin{array}{c} (\mathbf{T}_{pr} + p_k \mathbf{T}_k)^{-1} (\mathbf{T}_{pr} \boldsymbol{\mu}_{pr} + \mathbf{T}_k \sum_{i=1}^{p_k} \boldsymbol{\xi}_i^{(k)}) \\ (\mathbf{T}_{pr} + p_k \mathbf{T}_k)^{-1} \end{array} \right)$$

for $i = 1 : p_k$ **do**

 Draw $\boldsymbol{\xi}_i^{(k)}$ from a kernel that leaves $\mathcal{L}(\boldsymbol{\xi}_i^{(k)} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(-k)})$ invariant

end for

 Reparameterise $(\mathbf{a}^{(0)}, \boldsymbol{\xi}^{(k)}) \rightarrow (\mathbf{a}^{(0)}, \mathbf{a}^{(k)})$ according to (2)

end for

We now detail a gradient-based scheme that leaves $\mathcal{L}(\boldsymbol{\xi}_i^{(k)} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(-k)})$ invariant. The scheme is most concisely described by using the symbol $f_i^{(k)}(\mathbf{a}_i^{(k)})$ to refer to the log-likelihood of \mathbf{y} as a function of $\mathbf{a}_i^{(k)}$, holding other parameters fixed. When $L = 1$, we use a second-order Metropolis-Hastings proposal that requires no tuning parameters and is easy to implement. The approach is based on a second-order expansion of the corresponding density around the current value and amounts to generating a proposal $\tilde{\boldsymbol{\xi}}_i^{(k)} \sim \mathcal{N}(m_i^{(k)}(\boldsymbol{\xi}_i^{(k)}), c_i^{(k)}(\boldsymbol{\xi}_i^{(k)}))$, where

$$\begin{aligned} m_i^{(k)}(\boldsymbol{\xi}_i^{(k)}) &= c_i^{(k)}(\boldsymbol{\xi}_i^{(k)}) (\partial f_i^{(k)}(\boldsymbol{\xi}_i^{(k)} - \mathbf{a}^{(0)}) + \tau_k \mathbf{a}^{(0)} - \partial^2 f_i^{(k)}(\boldsymbol{\xi}_i^{(k)} - \mathbf{a}^{(0)}) \boldsymbol{\xi}_i^{(k)}), \\ c_i^{(k)}(\boldsymbol{\xi}_i^{(k)}) &= \left(\tau_k - \partial^2 f_i^{(k)}(\boldsymbol{\xi}_i^{(k)} - \mathbf{a}^{(0)}) \right)^{-1}, \end{aligned}$$

with τ_k denoting the prior precision of $a_i^{(k)}$ for $L = 1$. We accept $\tilde{\xi}_i^{(k)}$ with probability

$$1 \wedge \frac{f_i^{(k)}(\tilde{\xi}_i^{(k)} - a^{(0)}) \mathcal{N}(\tilde{\xi}_i^{(k)}; a^{(0)}, \tau_k^{-1}) \mathcal{N}(\xi_i^{(k)}; m_i^{(k)}(\tilde{\xi}_i^{(k)}), c_i^{(k)}(\tilde{\xi}_i^{(k)}))}{f_i^{(k)}(\xi_i^{(k)} - a^{(0)}) \mathcal{N}(\xi_i^{(k)}; a^{(0)}, \tau_k^{-1}) \mathcal{N}(\tilde{\xi}_i^{(k)}; m_i^{(k)}(\xi_i^{(k)}), c_i^{(k)}(\xi_i^{(k)})}. \quad (3)$$

When $L > 1$, we use the gradient-based Metropolis-Hastings sampler of [33], which proposes $\tilde{\xi}_i^{(k)} \sim \mathcal{N}(m_i^{(k)}, D_i^{(k)})$, with

$$\begin{aligned} m_i^{(k)} &= C_i^{(k)}(\xi_i^{(k)} / \delta_i^{(k)} + \nabla f_i^{(k)}(\xi_i^{(k)} - a^{(0)}) + T_k a^{(0)}), \\ D_i^{(k)} &= C_i^{(k)} + (C_i^{(k)})^2 / \delta_i^{(k)}, \\ C_i^{(k)} &= (T_k + I / \delta_i^{(k)})^{-1}, \end{aligned} \quad (4)$$

and accepting it with probability

$$1 \wedge \frac{f_i^{(k)}(\tilde{\xi}_i^{(k)} - a^{(0)}) \mathcal{N}(\tilde{\xi}_i^{(k)}; a^{(0)}, T_k^{-1}) \mathcal{N}(\xi_i^{(k)}; m_i^{(k)}(\tilde{\xi}_i^{(k)}), D_i^{(k)})}{f_i^{(k)}(\xi_i^{(k)} - a^{(0)}) \mathcal{N}(\xi_i^{(k)}; a^{(0)}, T_k^{-1}) \mathcal{N}(\tilde{\xi}_i^{(k)}; m_i^{(k)}(\xi_i^{(k)}), D_i^{(k)})}. \quad (5)$$

The scalars $\delta_i^{(k)}$ are level-specific step size parameters that we tune adaptively according to the Robbins-Monro procedure described in, e.g., Algorithm 4 of [1], with target acceptance rate equal to 0.5 and learning rate at iteration t set to $t^{-0.5}$. Notice that the matrices $D_i^{(k)}$ for different i 's share the same eigenspace. Accordingly, an efficient implementation first eigendecomposes $T_k = \mathbf{V}\mathbf{\Lambda}\mathbf{V}^T$, then computes

$$\sqrt{D_i^{(k)}} = \mathbf{V} \sqrt{(\mathbf{\Lambda} + I / \delta_i^{(k)})^{-2} / \delta_i^{(k)} + (\mathbf{\Lambda} + I / \delta_i^{(k)})^{-1}},$$

then simulates the proposal and evaluates its density.

Note that, while the local parameterisation proposed above is effective (as we show below), a global reparameterisation of the model would not be successful. The intuition is that one cannot simultaneously center all K factors. Indeed, [41, Theorem 5] show that for crossed effects models with $K = 2$ factors, all global reparameterisations based on hierarchical centering lead to a complexity that is super-linear in n and p .

2.2.1. Numerical illustration

To demonstrate the scalability of the local centering approach, we generate data of growing dimension, and infer the corresponding models using the local centering and the vanilla Gibbs samplers. Elements of \mathbf{a} are sampled from a standard Gaussian with settings $L = 1$, $K = 2$, $p_1 = p_2$ and increasing values of p_1 . Observations \mathbf{y} are missing completely at random, where each cell in the contingency table is missing with probability 0.9, therefore $n \approx 0.1p_1^2$. For the non-empty

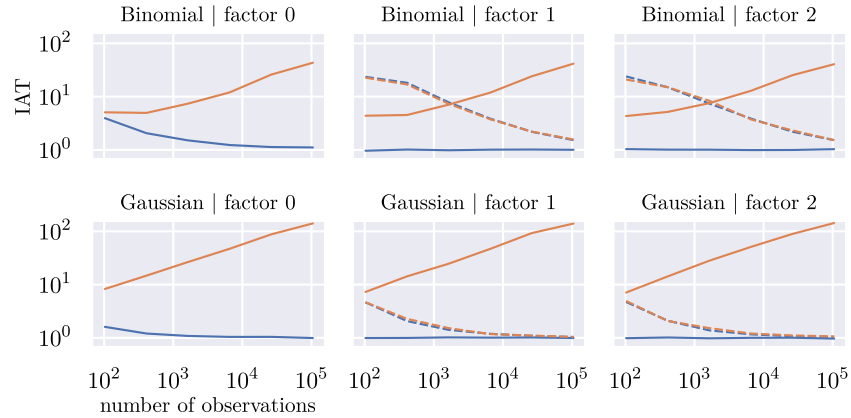


FIG 1. [Crossed effects simulation study, binomial and Gaussian likelihoods] *Integrated autocorrelation time estimates as a function of n for the locally centered (blue) and the vanilla (orange) Metropolis-within-Gibbs sampler. Solid lines refer to $a^{(0)}$ and $p_k^{-1} \sum_{i=1}^{p_k} a_i^{(k)}$, and dashed to τ_k^{-1} , for $k = 1, 2$. Estimates are computed by averaging over 10 MCMC runs with 10000 iterations each. We set the priors $\tau_k \sim \text{Gamma}(1/2, 1/2)$, a flat prior on $a^{(0)}$ and $\tau \sim \text{Gamma}(1/2, 1/2)$ for the Gaussian model. Note that whereas the simulation of data was done with different values of p_k , here we show the results as a function of the implied n .*

cells we simulate observations either from the Gaussian distribution $\mathcal{L}(y_j | \eta_j) = \mathcal{N}(\eta_j, 1)$, or from the binomial law $\mathcal{L}(y_j | \eta_j) = \text{Binomial}(1, (1 + \exp(-\eta_j))^{-1})$, and we use the corresponding likelihood when we learn the models. The precision parameters τ_k and the residual precision τ of the Gaussian model are treated as unknown.

We monitor the empirical estimates of *integrated autocorrelation times* (IAT) for various parameters, and report them in Figure 1. The IAT is a measure of MCMC sampling efficiency (less being better), whose definition and interpretation can be found in Appendix B. For the local centering algorithms this quantity remains upper bounded with respect to n (and thus also of p). This happens for all parameters and suggests that the convergence properties of Algorithm 1 do not deteriorate as n and p increase. Combining this with considerations of cost per iteration, we conjecture that Algorithm 1 is scalable. Section 2.5 contains a more detailed discussion and some theoretical results that support the conjecture.

2.3. Methodology for categorical likelihoods

The local centering approach of Section 2.2 applies to any sufficiently smooth likelihood. A case of special interest, see e.g. the application to electoral survey data below, is the categorical logistic regression model with likelihood

$$\mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j) = \text{Categorical}(\text{softmax}\{\boldsymbol{\eta}_j\}), \quad (6)$$

where softmax is defined as in (2). In this model, $\mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j) = \mathcal{L}(\mathbf{y}_j | \boldsymbol{\eta}_j + c\mathbf{1})$ for any constant c , where $\mathbf{1}$ denotes a vector of 1's of appropriate dimensionality. Thus, $\boldsymbol{\eta}_j$ is only identified by the likelihood up to such translations, which is also the case with the factor levels $\mathbf{a}_i^{(k)}$. Of course, given proper priors, the resulting posterior is proper. Therefore, we set proper priors throughout this section, including on $\mathbf{a}^{(0)}$, albeit weakly informative. Even with a proper prior, the posterior is typically much more informative about the differences in the elements of $\mathbf{a}_i^{(k)}$ than it is about the mean of said vector. As a consequence, the elements have strong posterior dependence, and an MCMC algorithm would need to propose highly correlated moves to fully traverse the posterior. In particular, the Metropolis-Hastings algorithm introduced in Section 2.2 will fail to propose such moves, since it approximates $\mathcal{L}(\boldsymbol{\xi}_i^{(k)} | \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(0)}, \mathbf{a}^{(-k)})$ by way of its gradient and the prior covariance, neither of which capture the extent of the dependence. Therefore, its mixing for categorical likelihoods can be poor, and we provide some numerical evidence to that effect in Figure 3.

Since the differences of the elements of $\mathbf{a}_i^{(k)}$ are usually well-identified, we do expect the Metropolis-Hastings algorithm to mix adequately once the magnitude of the vector has been fixed. Therefore, one way of addressing this issue is to change the prior on the factor levels by imposing an identifiability constraint. One such constraint is to set the L -th element $a_{iL}^{(k)}$ of the vector $\mathbf{a}_i^{(k)}$ to 0, define the remaining $L - 1$ elements as $\dot{\mathbf{a}}_i$, and specify the Gaussian prior

$$\mathcal{L}(\dot{\mathbf{a}}_i^{(k)} | \dot{\mathbf{T}}_k) = \mathcal{N}(\mathbf{0}, \dot{\mathbf{T}}_k^{-1}), \tag{7}$$

with an analogous constraint on $\mathbf{a}^{(0)}$. We may then proceed as in Section 2.2, replacing the symbols $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)}, \mathbf{T}_0, \mathbf{T}_k)$ with $(\dot{\mathbf{a}}^{(0)}, \dot{\mathbf{a}}^{(k)}, \dot{\mathbf{T}}_0, \dot{\mathbf{T}}_k)$. In doing so, we keep in mind that since $a_{iL}^{(k)} = 0$, the linear predictors have a trailing element fixed to 0, i.e.

$$\boldsymbol{\eta}_j = \mathbf{a}^{(0)} + \sum_{k=1}^K \mathbf{a}_{i_k[j]}^{(k)} = \begin{bmatrix} \dot{\mathbf{a}}^{(0)} + \sum_{k=1}^K \dot{\mathbf{a}}_{i_k[j]}^{(k)} \\ 0 \end{bmatrix}, \tag{8}$$

which needs to be accounted for when evaluating the likelihood and its gradient.

While this solution is computationally effective, it is statistically unsatisfactory because different identifiability constraints lead to different posterior inferences for the outcome probabilities of the categorical variable. We think it preferable to improve the algorithm without changing the model, and leave it to the modeler to impose constraints post-hoc if desired. Hence, we propose a novel scheme that, although conceptually close to the identifiability constraint described above, does not change the prior but rather changes the sampling algorithm. Rather than tethering one of the elements to 0, we follow the principle set out in Section 2.2 by updating the well-identified quantities by Metropolis-within-Gibbs, and the ill-identified translation through an additional Gibbs step which only involves the prior. To obtain the well-identified quantity, define $\ddot{\mathbf{a}}_i^{(k)}$ as the first $L - 1$ elements of $\mathbf{a}_i^{(k)} - a_{iL}^{(k)}\mathbf{1}$, and similarly $\ddot{\mathbf{a}}^{(0)}$ as the first $L - 1$

elements of $\mathbf{a}^{(0)} - a_L^{(0)}\mathbf{1}$. Note that the data \mathbf{y} are conditionally independent of $a_{iL}^{(k)}$ and $a_L^{(0)}$ given $(\ddot{\mathbf{a}}_i^{(k)}, \ddot{\mathbf{a}}^{(0)})$. Additionally, if $\mathcal{L}(\mathbf{a}_i^{(k)}|\gamma) = \mathcal{N}(\mathbf{0}, \mathbf{T}_k^{-1})$ with

$$\mathbf{T}_k^{-1} = \begin{bmatrix} \Sigma_{(-L,-L)}^{(k)} & \boldsymbol{\sigma}_{(-L,L)}^{(k)} \\ \boldsymbol{\sigma}_{(L,-L)}^{(k)} & \sigma_{(L,L)}^{(k)} \end{bmatrix}, \tag{9}$$

then the implied prior on $\ddot{\mathbf{a}}_i^{(k)}$ is $\mathcal{L}(\ddot{\mathbf{a}}_i^{(k)}|\gamma) = \mathcal{N}(\mathbf{0}, \ddot{\mathbf{T}}_k^{-1})$, where

$$\ddot{\mathbf{T}}_k^{-1} = \Sigma_{(-L,-L)}^{(k)} + \sigma_{(L,L)}^{(k)}\mathbf{1}\mathbf{1}^\top - \mathbf{1}\boldsymbol{\sigma}_{(L,-L)}^{(k)} - \boldsymbol{\sigma}_{(-L,L)}^{(k)}\mathbf{1}^\top, \tag{10}$$

with corresponding expressions for $\ddot{\mathbf{a}}^{(0)}$. Thus, our proposed scheme is to apply the methodology of Section 2.2 starting from $(\ddot{\mathbf{a}}_i^{(k)}, \ddot{\mathbf{a}}^{(0)})$ and their implied priors, merely replacing $(\mathbf{a}^{(0)}, \mathbf{a}^{(k)}, \mathbf{T}_0, \mathbf{T}_k)$ with $(\ddot{\mathbf{a}}^{(0)}, \ddot{\mathbf{a}}^{(k)}, \ddot{\mathbf{T}}_0, \ddot{\mathbf{T}}_k)$, and again keeping in mind that the linear predictors are as in (8). We then recover $a_{iL}^{(k)}$ by observing that

$$\begin{aligned} \mathcal{L}(a_{iL}^{(k)}|\gamma, \ddot{\mathbf{a}}_i^{(k)}) &= \mathcal{N}(\mathbf{b}_k^\top \ddot{\mathbf{a}}_i^{(k)}, \mathbf{b}_k^\top (\boldsymbol{\sigma}_{(-L,L)}^{(k)} - \sigma_{L,L}^{(k)}\mathbf{1})), \\ \mathbf{b}_k &= \ddot{\mathbf{T}}_k (\boldsymbol{\sigma}_{(-L,L)}^{(k)} - \sigma_{L,L}^{(k)}\mathbf{1}), \end{aligned} \tag{11}$$

with similar expressions for $a_L^{(0)}$. Thereafter, we transform back according to

$$\left(\begin{bmatrix} \ddot{\mathbf{a}}_i^{(k)} + a_{iL}^{(k)}\mathbf{1} \\ a_{iL}^{(k)} \end{bmatrix}, \begin{bmatrix} \ddot{\mathbf{a}}^{(0)} + a_L^{(0)}\mathbf{1} \\ a_L^{(0)} \end{bmatrix} \right) \rightarrow (\mathbf{a}_i^{(k)}, \mathbf{a}^{(0)}).$$

In effect, the resampling of $a_L^{(0)}$ according to simple prior-induced Gaussian computations induces the required positive dependence in the update. This procedure is repeated for all factors. We refer to an algorithm using the above procedure as a *projected* algorithm. The numerical experiments of the following section suggest that the projected algorithm with local centering mixes at least as well as its counterpart based on prior constraints. If required, such constraints can still be imposed post-sampling.

2.3.1. Illustration on electoral survey data

We experimentally assess algorithmic performance on a large-scale electoral survey dataset. The former previously featured in [20], and allows us to assess algorithms under realistic conditions that arise in applied work. Electoral surveys, consisting of the self-reported voting intention and various demographic traits for a sample of potential voters, provide one of the most natural applications for crossed-effects modeling, see for example [20, 16] and references therein. Here the response variable is the multicategorical voting intention, with each category corresponding to a different party. Demographic traits such as location and age serve as input variables; these may be stratified into multicategorical variables. Electoral surveys are useful both for disaggregating the

TABLE 1
Categorical factors included in the crossed effects elections model.

k	factor name	no. of levels (p_k)
1	province id	52
2	activity	4
3	age	3
4	education	3
5	municipality size	3
6	last vote	3
7	gender	2
–	response	3

national voting intention to local districts, which is crucial for predicting parliamentary representation, and for dealing with insurgent political parties with significant electoral success, which is becoming a common phenomenon in European politics. The survey was carried out by the *Centro de Investigaciones Sociológicas* ahead of the November 2019 Spanish general elections. For the experiments in this paper we focus on a simplified data structure with $L = 3$ parties, $K = 7$ input variables, and 210 regression parameters overall (3 per level, and $p = 70$ levels overall). There are $n = 9234$ survey respondents and the input variables define a contingency table with 33696 cells, of which only 4764 are non-empty, hence the table is *sparse*. Table 1 summarises the dataset characteristics.

The response being multi-categorical, we find ourselves in the setting of Section 2.3. In view of that discussion, we compare four algorithms: the vanilla Metropolis-within-Gibbs sampler for the constrained model (const/Van-MwG), the locally centered algorithm for the constrained model (const/LC-MwG), the locally centered algorithm for the unconstrained model (free/LC-MwG), and the projected local centering algorithm for the unconstrained model (free/PLC-MwG). Table 2 in the Appendix provides a list of the algorithms discussed in this article and their acronyms. Notice that the first two sample from a different posterior than the last two.

We carry out two experiments on this elections dataset, one of which is a comparison between const/Van-MwG, const/LC-MwG and the generic Stan implementation of NUTS, showing that const/LC-MwG provides a consistent speedup of more than an order of magnitude compared to Stan and const/Van-MwG, see Figure 2. The const/Van-MwG has inconsistent performance, mixing fairly well in the factors k with many levels, which identifies the mean of $\mathbf{a}^{(k)}$ and breaks dependence with $\mathbf{a}^{(0)}$. Factors with few levels remain strongly dependent with $\mathbf{a}^{(0)}$ a posteriori, and therefore mix slowly. This finding is in accordance with the theoretical results of [22] (e.g., their Theorem 3) that predict this behaviour, albeit under conditions not met in this situation.

The other experiment is a comparison between the various bespoke Gibbs samplers, which illustrates the importance of either constraining or projecting before local centering, see Figure 3. Both the const/LC-MwG and the free/PLC-MwG algorithms perform well, while the mixing of most parameters in the

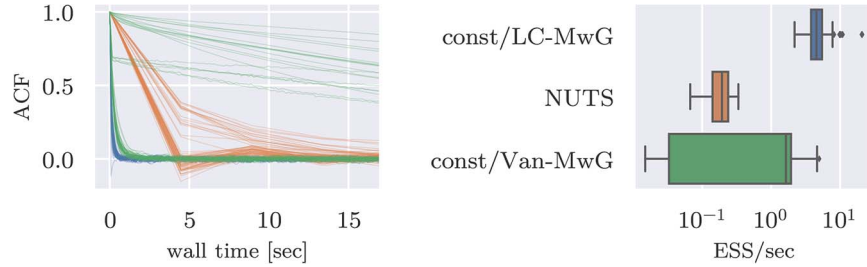


FIG 2. [Crossed effects elections model, constrained categorical likelihood] Comparison of sampling efficiency between the locally centered sampler (*const/LC-MwG*, blue), the vanilla Metropolis-within-Gibbs sampler for the constrained model (*const/Van-MwG*) and Stan/NUTS (orange). The left panel overlays parameter ACFs of \mathbf{a} and $\boldsymbol{\gamma}$ as a function of wall time. The right panel shows the distribution of effective samples per second. Estimates are computed from MCMC runs with 10000 (*const/LC-MwG*, *const/Van-MwG*) and 1000 iterations (NUTS) respectively, with an identical number of iterations used for adaptation. We set the priors $\mathbf{T}_k \sim \text{Wishart}(L-1, \mathbf{I}/(L-1))$, $\hat{\mathbf{a}}^{(0)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.

free/LC-MwG algorithm is two orders of magnitude slower. The extent of the slowdown depends on the strength of the prior, which can restrict the possible range of translations in the softmax; $\mathbf{a}^{(0)}$ and factors k with many levels and hence well-identified hyperparameter \mathbf{T}_k mix better. Since *const/LC-MwG* and *free/PLC-MwG* lead to similar computational efficiency, we recommend the use of the unconstrained model with the *free/PLC-MwG* algorithm and impose, if required, additional constraints in a post-sampling stage.

2.4. Methodology for binomial likelihoods: Polya-gamma augmentation and collapsed Gibbs sampling

While the local centering approach described above has the appeal of applying to general likelihoods, it is not optimal for Gaussian likelihood. In that instance we can directly sample $\mathcal{L}(\mathbf{a}^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{a}^{(-0, -k)})$ by exploiting Gaussian properties, as shown in [22] (and revisited later in this section). For certain non-Gaussian likelihoods we can still benefit from this direct approach when there are latent variables \mathbf{z} such that $\mathcal{L}(\mathbf{a}^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \mathbf{z}, \boldsymbol{\gamma}, \mathbf{a}^{(-0, -k)})$ is Gaussian. We focus here on one such case of significant practical importance, that of binomial likelihood.

[25] observed that when $\mathcal{L}(y_j \mid \eta_j) = \text{Binomial}(m_j, (1 + e^{-\eta_j})^{-1})$, then the corresponding probability mass function admits the following mixture representation:

$$\frac{(e^{\eta_j})^{y_j}}{(1 + e^{\eta_j})^{m_j}} = \int_0^\infty \frac{e^{(y_j - m_j/2)\eta_j - z_j \eta_j^2/2}}{2^{m_j}} \text{PG}(z_j; m_j, 0) dz_j \quad (12)$$

where $\text{PG}(z; \alpha, \beta)$ denotes the density of the so-called *Polya-gamma* distribution with parameters α and β evaluated at z . The *Polya-gamma* distribution is constructed in [25] as a weighted, infinite sum of Gamma random variables.

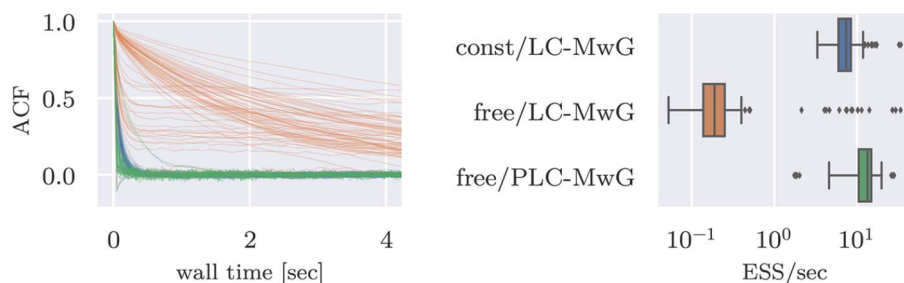


FIG 3. [Crossed effects election model, categorical likelihood] Comparison of sampling efficiency between the locally centered sampler on the constrained model (const/LC-MwG, blue), the local centering sampler on the free model (free/LC-MwG, orange) and the projected local centering sampler on the free model (free/PLC-MwG, green). The left panel overlays parameter ACFs of \mathbf{a} and $\boldsymbol{\gamma}$ as a function of wall time. The right panel shows the distribution of effective samples per iteration. Estimates are computed from MCMC runs with 10000 iterations each, with an identical number of iterations used for adaptation. We set the priors $\mathbf{T}_k \sim \text{Wishart}(L, \mathbf{I}/L)$, $\mathbf{a}^{(0)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in the free likelihood case and $\mathbf{T}_k \sim \text{Wishart}(L - 1, \mathbf{I}/(L - 1))$, $\hat{\mathbf{a}}^{(0)} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ in the constrained likelihood case.

For integer-valued α , a $\text{PG}(\alpha, \beta)$ -variable is obtained as a sum of α $\text{PG}(1, \beta)$ -variables, while simulation of a $\text{PG}(1, \beta)$ -variable is done by exploiting the representation of its density as a bounded Cauchy series whose increments have alternating signs.

This neat construction leads to the following computational strategy. We augment the space with the vector $\mathbf{z} = (z_1, \dots, z_n)$ and design a sampler for $\mathcal{L}(\mathbf{a}, \boldsymbol{\gamma}, \mathbf{z} \mid \mathbf{y})$. It is fairly simple to check (see [25]) that

$$\mathcal{L}(\mathbf{z} \mid \mathbf{y}, \mathbf{a}, \boldsymbol{\gamma}) = \prod_{j=1}^n \text{PG}(m_j, \eta_j), \tag{13}$$

with m_j and η_j defined above. Another direct calculation shows that $\mathcal{L}(\mathbf{a} \mid \mathbf{y}, \mathbf{z}, \boldsymbol{\gamma})$ is Gaussian. Then, we can follow the same paradigm as in Section 2.2, but sample $\mathcal{L}(a^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \mathbf{z}, \boldsymbol{\gamma}, \mathbf{a}^{(-0, -k)})$ directly from the corresponding Gaussian distribution. Algorithm 2 summarizes how this step is carried out.

An obvious advantage of this approach over the one based on local centering, when applicable, is the avoidance of the Metropolis-Hasting step, which typically slows down mixing and requires tuning. However, there are two disadvantages to the Poly-gamma augmentation. One is that the dimension of the distribution to be sampled is increased by the n additional latent variables. Another is that sampling $\text{PG}(m_j, \eta_j)$ has computational cost $\mathcal{O}(m_j)$, hence it will be inefficient for responses with large counts (although see [39] for some mitigation of this problem).

Notice that we may contrast the above algorithm, which we call the Poly-gamma collapsed Gibbs sampler (PG/Col-G), with another sampler on the augmented space that updates effects one at a time, and which we call the Poly-gamma vanilla Gibbs sampler (PG/Van-G). We also compare those samplers

Algorithm 2 The collapsed Gibbs sampler for $\mathcal{L}(\mathbf{a}|\mathbf{y}, \boldsymbol{\gamma}, \mathbf{z})$.

Define the quantities

$$v_j = (y_j - m_j/2)/z_j, \quad z_i^{(k)} = \sum_{j:i_k[j]=i} z_j, \quad s_i^{(k)} = z_i^{(k)}/(\tau_k + z_i^{(k)}). \quad (14)$$

for $k = 1 : K$ **do**
 Draw $a^{(0)}$ from

$$\mathcal{L}(a^{(0)} | \mathbf{y}, \mathbf{a}^{(-0, -k)}) = \mathcal{N} \left(\frac{\frac{1}{\sum_i s_i^{(k)}} \sum_i \frac{s_i^{(k)}}{z_i^{(k)}} \sum_{j:i_k[j]=i} z_j (v_j - \sum_{k' \neq k} a_{i_{k'}[j]}),}{(\tau_k \sum_i s_i^{(k)})^{-1}} \right)$$

for $i = 1 : p_k$ **do**
Draw $a_i^{(k)}$ from

$$\mathcal{L}(a_i^{(k)} | \cdot) = \mathcal{N} \left(\frac{\frac{s_i^{(k)}}{z_i^{(k)}} \sum_{j:i_k[j]=i} z_j (v_j - a^{(0)} - \sum_{k' \neq k} a_{i_{k'}[j]}),}{(z_i^{(k)} + \tau_k)^{-1}} \right)$$

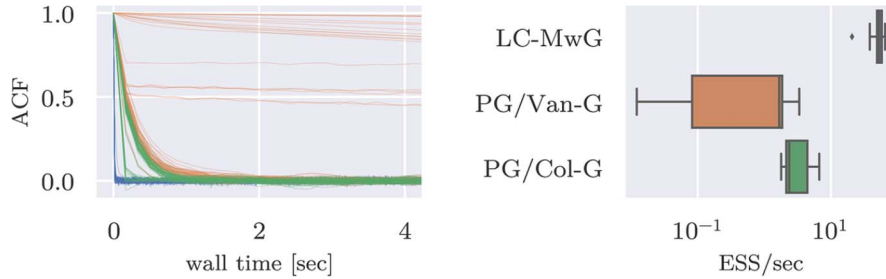
end for**end for**

FIG 4. [Crossed effects election model, binomial likelihood] Comparison of sampling efficiency between the collapsed Gibbs sampler on the augmented space (PG/Col-G, green), the locally centered marginal sampler (LC-MwG, blue) and the vanilla Gibbs sampler on the augmented space (PG/Van-G, orange). The left panel overlays parameter ACFs of \mathbf{a} and $\boldsymbol{\gamma}$ as a function of wall time. The right panel shows the distribution of effective samples per iteration. Estimates are computed from MCMC runs with 10000 iterations each. We set the priors $\tau_k \sim \text{Gamma}(1/2, 1/2)$ and a flat prior on $a^{(0)}$.

to the local centering algorithm (LC-MwG) applied to binomial likelihood.

Figure 4 shows that the scalable algorithms LC-MwG and PG/Col-G significantly outperform the non-scalable algorithm PG/Van-G on the elections dataset. As in Figure 2, the vanilla algorithm slows down the most for the factors with few levels. Note that while PG/Col-G and LC-MwG have similar ESS per iteration, PG/Col-G has substantially higher iteration time due to the latent variable update.

2.5. Theoretical analysis

We now perform a theoretical analysis of the complexity of the algorithms described above. The computational complexity of an MCMC algorithm depends both on the cost per iteration and on the number of iterations required to obtain each effective sample. In this work, we quantify the number of iterations required through the relaxation time of the associated Markov chain (other options are possible, but we do not discuss them here), leading to the following definition.

Definition 1. We define the complexity of an MCMC algorithm as

$$\text{Cost(MCMC)} = (\text{cost per iteration}) \times (\text{relaxation time}), \quad (15)$$

where the relaxation time refers to the reciprocal of 1 minus the geometric rate of convergence of the associated Markov chain [28, Proposition 1].

We now analyse the cost per iteration and the relaxation time of the proposed schemes. Following the discussion in Sections 1 and 2.1, our theoretical analysis applies to the algorithm for sampling $\mathcal{L}(\mathbf{a} \mid \mathbf{y}, \boldsymbol{\gamma})$, even though all our numerics pertain to sampling $\mathcal{L}(\mathbf{a}, \boldsymbol{\gamma} \mid \mathbf{y})$.

2.5.1. Cost per iteration of the proposed schemes

The cost per iteration of the Gibbs sampler based on local centering proposed in Section 2.2 is of order $\mathcal{O}(nKL + KL^3 + pL^2)$, as we show now. The evaluation of the gradients terms in (4) has a $\mathcal{O}(nL)$ cost for each factor, since

$$\nabla f_i^{(k)}(\mathbf{a}_i^{(k)}) = \sum_{j:i_k[j]=i} \nabla_{\mathbf{a}_i^{(k)}} \log p(\mathbf{y}_j \mid \mathbf{a}),$$

implies that exactly N terms of the form $\nabla_{\mathbf{a}_i^{(k)}} \log p(\mathbf{y}_j \mid \mathbf{a})$ for $j = 1, \dots, n$ need to be computed, at $\mathcal{O}(L)$ cost each. By similar arguments, the likelihood evaluations that are part of the acceptance probability computation also require $\mathcal{O}(nL)$ operations per factor. Given the gradient and likelihood computations, one can eigendecompose \mathbf{T}_k at $\mathcal{O}(L^3)$ cost and then perform the updates of $\boldsymbol{\xi}_i^{(k)}$ for all i at $\mathcal{O}(p_k L^2)$ cost. Summing over factors we obtain the $\mathcal{O}(nKL + KL^3 + pL^2)$ cost. The update from $\mathcal{L}(\mathbf{a}^{(0)} \mid \boldsymbol{\xi}^{(k)})$ requires $\mathcal{O}(L^3)$ operations, which does not influence the overall cost and can further be reduced to $\mathcal{O}(L^2)$ if the prior precision \mathbf{T}_{pr} is diagonal and \mathbf{T}_k has already been eigendecomposed.

Conveniently, with $L = 1$, this formula also holds for the collapsed Gibbs sampler for binomial likelihoods in Section 2.4, i.e. the complexity of updating $\mathcal{L}(\mathbf{a}^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \mathbf{a}^{(-0,-k)}, \boldsymbol{\gamma}, \mathbf{z})$ for all $k = 1, \dots, K$ is $\mathcal{O}(nK + p)$. This is due to needing to sample from p Gaussian distributions, and carrying out $\mathcal{O}(nK)$ summations to calculate their parameters. When sampling $\mathcal{L}(\mathbf{z} \mid \mathbf{y}, \mathbf{a}, \boldsymbol{\gamma})$ we incur the additional cost of $\mathcal{O}(\sum_j m_j)$, for a total cost of $\mathcal{O}(nK + p + \sum_j m_j)$.

For the categorical const/LC and free/pLC algorithms, notice that if the response is $L + 1$ -categorical, then local centering is done in L dimensions. The additional computations in the free/pLC algorithm can also be carried out with complexity $\mathcal{O}(KL^3)$. Conversely, for free/pLC, the update to γ is still $L + 1$ -dimensional, therefore the complexity of the full algorithm is $\mathcal{O}(nKL + K(L + 1)^3 + pL^2)$.

2.5.2. Relaxation time analysis

While the relaxation time of MCMC algorithms is hard to characterize in general, that of the Gibbs sampler for crossed effects models with Gaussian likelihood has become more amenable to explicit computations by combining the classical results of [27] with the so-called multigrid decomposition developed in [41, 22]. In this section we extend those results to analyze and compare the sampling algorithms discussed in previous sections, providing two main results. The first, contained in Proposition 2, Corollary 1 and Theorem 1, provides an analysis of local centering and a comparison of the latter with the collapsed Gibbs Sampler. These results show that local centering always increases the relaxation time compared to exact collapsing, but does so only by a constant factor, thus leading to the same asymptotic complexity. The second contribution, formalized in Theorem 2, is to expand the results in [22] by leveraging connections to random graph theory to show that both collapsing and local centering induce algorithms that are scalable as n and p diverge for average-case observation designs.

In this section we focus on models with Gaussian likelihood. Also, we assume $L = 1$ to keep the notation simple, even if most of the results would naturally extend to the $L > 1$ case. More precisely, we study algorithms targeting the law $\mathcal{L}(\mathbf{a} \mid \mathbf{y}, \gamma)$ induced by the following model

$$\begin{aligned} \mathcal{L}(a^{(0)}) &= \mathcal{N}(\mu_{pr}, \tau_{pr}^{-1}), \\ \mathcal{L}(\mathbf{a}^{(k)} \mid \tau_k) &= \mathcal{N}(\mathbf{0}, (\tau_k \mathbf{I})^{-1}), \quad k = 1, \dots, K, \\ \mathcal{L}(y_j \mid \mathbf{a}, \gamma) &= \mathcal{N}(a^{(0)} + \sum_{k=1}^K a_{i_k[j]}^{(k)}, \tau^{-1}) \quad j = 1, \dots, n, \end{aligned} \tag{16}$$

for fixed hyperparameters $\gamma = (\tau, \tau_1, \dots, \tau_k)$. We compare the following three schemes:

- (GS) the vanilla Gibbs sampler that updates from $\mathcal{L}(\mathbf{a}^{(k)} \mid \mathbf{y}, \gamma, \mathbf{a}^{(-k)})$ for $k = 0, \dots, K$;
- (cGS) the collapsed Gibbs sampler updating from $\mathcal{L}(a^{(0)}, \mathbf{a}^{(k)} \mid \mathbf{y}, \gamma, \mathbf{a}^{(-0, -k)})$ for $k = 1, \dots, K$;
- (Loc) the Gibbs sampler based on local centering that executes Algorithm 1 with the update of $\xi_i^{(k)}$ being the exact update from the conditional $\mathcal{L}(\xi_i^{(k)} \mid \mathbf{y}, \gamma, \mathbf{a}^{(-k)})$, which is feasible in the Gaussian case.

We first define some concepts and quantities related to the data observation pattern.

Definition 2. We denote occurrence and co-occurrence counts associated to levels of different factors by

$$n_i^{(k)} = \sum_{j=1}^n \mathbf{1}(i_k[j] = i),$$

$$n_{ii'}^{(k,k')} = \sum_{j=1}^n \mathbf{1}(i_k[j] = i, i_{k'}[j] = i').$$

We say that a design has balanced levels if $n_i^{(k)} = n/p_k$ for each factor k and level i . We define

$$\bar{n} = Kn/p$$

for $p = \sum_k p_k$, to be interpreted as the average number of observed data points per factor level.

The multigrid decomposition relates the rate of convergence of the Gibbs sampling Markov chain to that of simpler Markov chains. We summarise the main result in Proposition 1 below, which is an amalgamation of Theorem 1, Lemmas 1 and 2 of [22], extended to the case of local centering. The proof is analogous to the ones of the results just mentioned, and is omitted here for brevity.

Proposition 1 (Multigrid decomposition). For a factor k , let $\bar{a}^{(k)} = \sum_i a_i^{(k)}/p_k$ be the factor average, and $\delta a_i^{(k)} = a_i^{(k)} - \bar{a}^{(k)}$ be the factor residuals. Let $\{\mathbf{a}(t)\}$, where t indexes iteration, be the Markov chain generated by either (GS), (cGS) or (Loc). Let $\{\phi(t)\}$ and $\{\psi(t)\}$ be the processes that correspond to the factor averages (including $a^{(0)}$) and factor residuals respectively. Then, under balanced levels, the processes $\{\phi(t)\}$ and $\{\psi(t)\}$ are Markov chains and independent of each other, and the rate of convergence of $\{\mathbf{a}(t)\}$ is the maximum of their convergence rates. Additionally, the convergence rate of $\{\psi(t)\}$ is the same for all three algorithms.

Proposition 1 reduces the computation of the rate of convergence of the algorithm of interest to that of the rates of two different Markov chains, one low-dimensional that operates on factor averages, and another high-dimensional on factor residuals. The latter is common to all algorithms under consideration, therefore the algorithms differ in the speed of the chain of the averages. Although defined on a space of fixed dimension, the chain of the averages can have a mixing time that deteriorates as n and p grow. In fact, Proposition 3 of [22] shows that vanilla Gibbs sampling is not scalable precisely for this reason in regimes where both n and p grow. On the other hand, the same proposition shows that the chain of averages generates independent samples in the case of the collapsed Gibbs sampler, therefore the rate of convergence of this algorithm is completely determined by that of the residuals. While in the case of local centering the

chain of averages does not consist of independent draws, in Proposition 2 we show that it has a relaxation time that is bounded by constants that do not depend on n and p . The remaining results of this section, whose proofs can be found in Appendix A, are derived under the following assumption:

(A1) Assume model (16) with balanced levels and an improper flat prior on $a^{(0)}$, i.e. $\mu_{pr} = \tau_{pr} = 0$.

Proposition 2. *Under (A1), the relaxation time of the Markov chain of averages $\{\psi(t)\}$ induced by (Loc) is upper bounded by $1 + \min_{k=1,\dots,K} \tau_k/\tau$.*

The previous results imply that (cGS) and (Loc) have similar convergence properties, as formally shown below.

Corollary 1. *Under (A1), the relaxation times of (cGS) and (Loc), denoted by T_{cgs} and T_{loc} respectively, satisfy*

$$T_{\text{cgs}} \leq T_{\text{loc}} \leq T_{\text{cgs}} + C,$$

where C is a constant depending only on γ and not on n , p or K .

Obviously, for non-Gaussian likelihoods for which the sampling of the $\xi_i^{(k)}$ cannot be done directly, but where a gradient-based sampler is used instead, one would expect a further increase in relaxation time. However, as shown by numerical results in the above sections, this will typically involve only constants that do not depend on n and p .

Therefore, in view of Corollary 1, for the rest of the section we focus on the collapsed Gibbs sampler, and refer to its computational complexity as $\text{Cost}(\text{cGS})$. We now provide a more refined result on its scalability in the two-factor case, i.e. $K = 2$, building on previous work in [22]. On the basis of the previous discussion, a characterisation of the relaxation time of the chain of residuals is central to the following result. It involves an auxiliary Markov chain on a discrete state-space that we introduce first.

Definition 3. T_{aux} is the relaxation time of the two-component deterministic-scan Gibbs sampler that targets the discrete distribution with state space $\{1, \dots, p_1\} \times \{1, \dots, p_2\}$ and probability mass function proportional to $n_{ii'}^{(1,2)}$ for (i, i') being an element of the state space.

Theorem 1. *Assume (A1) and $K = 2$. Then*

$$\text{Cost}(\text{cGS}) = \mathcal{O}(\max\{n, p\} \min\{\bar{n}, T_{\text{aux}}\}), \quad (17)$$

where the constants depend only on γ and not on n or p .

The factor $\max\{n, p\}$ comes from the cost per iteration, and $\min\{\bar{n}, T_{\text{aux}}\}$ comes from the relaxation time, which in view of the previous discussion is that of the chain of residuals.

To obtain some intuition about the implications of this result, it is convenient to interpret the 2-factor ($K = 2$) model as a recommender system, where one

factor pertains to users, and the other to products. Note that \bar{n} can be interpreted as the average number of products rated by a customer; it is precisely this number when $p_1 = p_2$ and the design has balanced levels. In very sparse designs, increasing values of n will not be associated with increasing values of \bar{n} (more products and customers, but each customer rates a bounded number of products). In such regimes the sampler is scalable. In less sparse designs where \bar{n} increases with n , the contingency table is more populated, and provided it is populated enough for the auxiliary Markov chain to mix well, the sampler will again be scalable.

The above discussion suggest that \bar{n} and T_{aux} typically exhibit complementary behaviour with respect to sparsity of the data co-occurrence matrix. However, the value of T_{aux} does not only depend on the level of sparsity, but also on the specific observation pattern, encoded in the co-occurrence counts $n^{(1,2)} = (n_{ii'})_{i=1,\dots,p_1;i'=1,\dots,p_2}$, and can vary greatly across different designs with the same sparsity level. Indeed, T_{aux} is the only term in the bound (17) that does not only depend on the amount of data and parameters in the model, but also on the observation pattern. We now perform a more refined analysis of $\text{Cost}(cGS)$ in (17) by considering random observation patterns and relating T_{aux} to the rich literature on spectral properties of random walks on random graphs.

We consider random binary designs with balanced levels. Specifically, given non-negative integers d_1, d_2 and n such that n is a multiple of both d_1 and d_2 , we define $\mathcal{D}(n, d_1, d_2)$ as the collection of all $p_1 \times p_2$ binary designs with balanced levels and n observations in total, where $p_1 = n/d_1, p_2 = n/d_2$. In other words $\mathcal{D}(n, d_1, d_2)$ contains all designs $n^{(1,2)}$ such that $n_i^{(1)} = d_1, n_{i'}^{(2)} = d_2$ and $n_{ii'}^{(1,2)} \in \{0, 1\}$ for all $i = 1, \dots, p_1$ and $i' = 1, \dots, p_2$. We will assume that the observed design is sampled uniformly from such collection, i.e. $n^{(1,2)} \sim \text{Unif}(\mathcal{D}(n, d_1, d_2))$. The following theorem provides a characterization of T_{aux} and thus of $\text{Cost}(cGS)$ in such regimes.

Theorem 2. *Assume (A1) and $K = 2$. Let d_1 and d_2 be fixed integers larger than 4, and assume $n^{(1,2)} \sim \text{Unif}(\mathcal{D}(n, d_1, d_2))$. Then, for any $\epsilon > 0$, we have*

$$T_{\text{aux}} \leq 1 + \frac{2}{\sqrt{\min\{d_1, d_2\}} - 2} + \epsilon \tag{18}$$

and

$$\text{Cost}(cGS) = \mathcal{O}(\max\{n, p\}), \tag{19}$$

asymptotically almost surely as $n \rightarrow \infty$, with constants in (19) depending only on γ and not on n, p, d_1 or d_2 .

Theorem 2 suggests that the collapsed Gibbs Sampler for 2-factor models is scalable in the average case, where average refers to the data co-occurrence pattern $n^{(1,2)}$. In particular, the almost sure statements (18)–(19) imply that for fixed d_1 and d_2 , the fractions of balanced levels designs in $\mathcal{D}(n, d_1, d_2)$ for which the collapsed Gibbs sampler is scalable goes to 1 as $n \rightarrow \infty$.

As shown in the proof of Theorem 2, which can be found in Appendix A, (18) follows from relating T_{aux} to the spectrum of random bipartite bi-regular graphs. This allows to then apply an extension for bipartite graphs, which was recently developed in [5], of the so-called Friedman's second largest eigenvalue theorem. The statement in (19) then follows by combining (18) and Theorem 1. Thus, at a high level, the positive result in (19) is related to the fact that random graphs have near-optimal connectivity properties (i.e. they are expanders), as shown by a large body of random graph theory literature, see e.g. references in [5].

Interestingly, the upper bound in (18) goes to 1 as $d_1, d_2 \rightarrow \infty$. This suggests that in regimes where d_1 and d_2 grow with n , which is a reasonable assumption in many contexts, then T_{aux} will get closer and closer to 1, which in turn by Theorem 1 implies that the relaxation time of the collapsed Gibbs Sampler will converge to 1. While this cannot be deduced from Theorem 2, since the regime studied there keeps d_1 and d_2 fixed and does not allow for d_1 and d_2 to grow n , the above results still point to a *blessing of dimensionality*, where the sampling algorithm actually converges faster and faster as n and p grow. This conjecture is coherent with the empirical results in, e.g., Figure 1 as well as results in [22, 14].

The theory developed in this section holds for Gaussian likelihoods. We see this more as a necessity of our strategy in proving the result, as opposed to the nature of the result itself, which we believe relates to the conditional independence structure in the model. This is, again, corroborated by the findings in Figure 1, which suggest that both the change from Gibbs to Metropolis-within-Gibbs, from known to unknown γ and from Gaussian to non-Gaussian likelihood slow down convergence only by a factor that is constant w.r.t. n and p , and do not affect the scalability of the resulting MCMC algorithm.

2.6. Related literature and alternative approaches

Early work on Bayesian computation for crossed effects models includes [34] and [11, Section 6], who discuss, respectively, the use of identifiability constraints and reparameterisation for computational purposes. [8, 9] provide a systematic discussion of the super-linear cost of standard frequentist fitting procedures when $K = 2$, and develop a scalable method of moments. [22] provide a detailed analysis of the Gibbs sampler's complexity, covering also sparse balanced designs and $K \geq 2$ factors, and propose the collapsed Gibbs sampler for Gaussian likelihood. [14] propose a coordinate-wise descent method for computing the MAP estimator (phrased in their article as a backfitting procedure to compute generalized least squares estimates) that has close connections to collapsed Gibbs sampling. They prove $\mathcal{O}(\max\{n, p\})$ complexity results also for somewhat unbalanced designs, their arguments are based on simple but effective concentration inequalities and linear algebra techniques.

A common practice among practitioners is the imposition of constraints within each factor, for example $\sum_i \mathbf{a}_i^{(k)} = \mathbf{0}$. Of course, such constraints are not needed for the inference to make sense in a Bayesian formulation, if desired

they can be imposed a posteriori, with different constraints leading to different inferences. To avoid confusion, this is an entirely different issue to the identifiability issues that arise in categorical logistic models, and which are discussed in Section 2.3. From the point of view of this article it is worth considering their impact on algorithmic performance. Theorem 6 of [41] derives the relaxation times of the vanilla Gibbs sampler for crossed effects models with Gaussian likelihood under identifiability constraints. Some constraints can make the vanilla Gibbs sampler scalable, whereas others cannot. Unfortunately, those found to be effective lead to difficult sampling problems when considering non-Gaussian likelihoods.

One could combine collapsing with techniques that seek to reduce posterior dependence between \mathbf{a} and $\boldsymbol{\gamma}$, such as parameter expansion [17, 18]. While useful and applicable, in the simulation studies of [22] for crossed effects models with Gaussian likelihoods, this strategy was found to have relatively limited impact on the resulting MCMC efficiency and it did not appear to change the overall complexity of the parent algorithm. Nevertheless, this is an important aspect of designing useful algorithms for large scale problems.

[19] and [16] propose mean field variational Bayes procedures for crossed models with Gaussian and logistic likelihoods, respectively, motivated by the overly high computational cost of Laplace approximations and HMC sampling for such models. There is an interesting interplay and scope for cross-fertilization between the methodology we develop in this article and variational inference, and we will report our progress in that direction in future work.

Many popular software for models with Gaussian latent variables, such as *lme4* [3] or *INLA* [30], perform approximate integration of \mathbf{a} using variants of the Laplace approximation. However, these approaches rely on sparse linear algebra techniques, which in general incurs super-linear complexity in $\max\{n, p\}$ for crossed effects models, see e.g. numerical results and related discussions in [8, 19, 22, 9, 14, 15]. On the contrary, such methods are entirely effective for nested multilevel models, as discussed in the following section.

3. Nested multilevel models

We consider the following broad class of nested multilevel hierarchical models. There is an underlying hierarchical organization in the data that is structured as a tree. To give a concrete example, in the large scale applications we consider later in this Section, the spatial domain in Spain is organized according to 5-digit postal codes where any longer sequence of digits (from 2 to 5) corresponds to an administrative unit (e.g., municipality) that is nested geographically within those that omit the last digits (e.g., city, province etc). A different example is classes, nested within schools, nested within municipalities, etc. We denote the root of the tree as i_0 , its offsprings as $i_0 i_1$, for $i_1 = 1, \dots, n_0$, for each such offspring its offsprings are $i_0 i_1 i_2$, for $i_2 = 1, \dots, n_{i_0 i_1}$, etc. Overall there are K levels apart from the root (hence $K = 4$ in the postal code example). In full generality we observe a matrix of covariates $\mathbf{X}_{i_0 \dots i_k}$ and a vector of responses

$\mathbf{y}_{i_0 \dots i_k}$ and at every node in the tree, although it is more common that such observations are available only at level K . For example, in one of the real estate models we consider in the sequel for each leaf at $K = 4$ the covariate is a national house pricing index and the response local average house sale prices, with replications given over different quarters. The model relates the response to the covariates via a linear predictor:

$$\boldsymbol{\eta}_{i_0 \dots i_k} = \mathbf{X}_{i_0 \dots i_k} \boldsymbol{\beta}_{i_0 \dots i_k}.$$

Inference borrows strength across units by shrinking the regression parameters to those at the deeper level $\boldsymbol{\beta}_{i_0 \dots i_{k-1}}$ in terms of a Gaussian shrinkage prior:

$$\mathcal{L}(\boldsymbol{\beta}_{i_0 \dots i_k} \mid \boldsymbol{\beta}_{i_0 \dots i_{k-1}}) = \mathcal{N}(\mathbf{A}_{i_0 \dots i_k} \boldsymbol{\beta}_{i_0 \dots i_{k-1}}, \mathbf{T}_{i_0 \dots i_{k-1}}^{-1}),$$

where $\mathbf{A}_{i_0 \dots i_k}$'s are design matrices, often simply the identity matrices. Putting these ingredients together, the full model is as follows:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\beta}_{i_0}) &= \mathcal{N}(\boldsymbol{\mu}_{pr}, \mathbf{T}_{pr}^{-1}), \\ \mathcal{L}(\mathbf{T}_{i_0 \dots i_k}) &= \text{Wishart}(\nu_{i_0 \dots i_k}, \mathbf{I} / \nu_{i_0 \dots i_k}), \\ \mathcal{L}(\boldsymbol{\beta}_{i_0 \dots i_k} \mid \boldsymbol{\beta}_{i_0 \dots i_{k-1}}, \mathbf{T}_{i_0 \dots i_{k-1}}) &= \mathcal{N}(\mathbf{A}_{i_0 \dots i_k} \boldsymbol{\beta}_{i_0 \dots i_{k-1}}, \mathbf{T}_{i_0 \dots i_{k-1}}^{-1}), \\ \boldsymbol{\eta}_{i_0 \dots i_k} &= \mathbf{X}_{i_0 \dots i_k} \boldsymbol{\beta}_{i_0 \dots i_k}, \\ \mathcal{L}(\mathbf{y}_{i_0 \dots i_k} \mid \cdot) &= \mathcal{L}(\mathbf{y}_{i_0 \dots i_k} \mid \boldsymbol{\eta}_{i_0 \dots i_k}), \\ &k = 0, \dots, K. \end{aligned} \quad (20)$$

When the likelihood is Gaussian, the posterior $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \boldsymbol{\gamma})$ is a multivariate Gaussian that can be sampled directly in a provably scalable way through sparse linear algebra. The result is given in Proposition 3. For non-Gaussian likelihoods this result can be leveraged in different ways. One, which we explore in this work, is in situations where latent variables \mathbf{z} can be augmented such that $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \mathbf{z}, \boldsymbol{\gamma})$ is Gaussian. This is in the same spirit as collapsed Gibbs sampling with Polya-gamma augmentation, which we developed in Section 2.4 for crossed effects models. In Section 3.2 we show results for binomial and Cauchy likelihoods. A different strategy appropriate for light-tailed likelihoods, such as the Poisson, is to combine the efficient Gaussian sampling with distributed gradient-based sampling of the regression parameters on the leaves of the tree; this is also discussed in Section 3.2.

3.1. Methodology for Gaussian likelihoods: sparse linear algebra

For nested multilevel models with Gaussian likelihoods, we can generate samples from the posterior $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \boldsymbol{\gamma})$ directly at $\mathcal{O}(\max\{n, p\})$ cost, as shown in Proposition 3. Specifically, we consider the model (20) with likelihood

$$\mathcal{L}(\mathbf{y}_{i_0 \dots i_k} \mid \boldsymbol{\eta}_{i_0 \dots i_k}) = \mathcal{N}(\boldsymbol{\eta}_{i_0 \dots i_k}, \tau_{i_0 \dots i_k}^{-1} \mathbf{I}). \quad (21)$$

As a consequence of (20)–(21), the conditional distribution $\mathcal{L}(\beta \mid \mathbf{y}, \gamma)$ is multivariate Gaussian. Sampling a multivariate Gaussian distribution can be reduced to solving linear systems; for details on this perspective and associated tools we refer to [29], who provides results that can be used to prove the following claims. When, given appropriate matrices \mathbf{V} , \mathbf{U} and \mathbf{R} , the model can be expressed as

$$\begin{aligned} \mathcal{L}(\beta \mid \gamma) &= \mathcal{N}(\mathbf{m}, \mathbf{V}^{-1}), \\ -2 \log p(\mathbf{y} \mid \beta, \gamma) &= \beta^\top \mathbf{U} \beta - 2\mathbf{y}^\top \mathbf{R}^\top \beta, \end{aligned} \tag{22}$$

(the second equation up to constants in β), then $\mathcal{L}(\beta \mid \mathbf{y}, \gamma)$ is also a Gaussian distribution with precision $\mathbf{Q} = \mathbf{U} + \mathbf{V}$. Let \mathbf{L} be the lower-triangular Cholesky factor of \mathbf{Q} , and \mathbf{z} be a standard Gaussian vector. Then, a sample from the posterior $\mathcal{L}(\beta \mid \mathbf{y}, \gamma)$ is obtained by solving the following two systems:

$$\begin{aligned} \mathbf{L}\mathbf{w} &= \mathbf{R}\mathbf{y} + \mathbf{V}\mathbf{m}, \\ \mathbf{L}^\top \beta &= \mathbf{w} + \mathbf{z}. \end{aligned} \tag{23}$$

It is well-known that \mathbf{L} can be computed column-wise from first to last column. Hence, the computation of \mathbf{L} , and the simultaneous solution of the top equation in (23) by forward substitution, can be thought of as a *forward pass*. Having completed the forward pass, we can then solve the second equation in (23) by backward substitution, in what we can think of as a *backward pass*. The cost of solving these systems is dominated by the cost of computing \mathbf{L} , which is in general cubic in the dimension of β .

A direct calculation verifies the following claims. Let \mathbf{Q} denote the posterior precision associated with the Gaussian $\mathcal{L}(\beta \mid \mathbf{y}, \gamma)$. We refer to the elements of \mathbf{Q} by indicating what blocks of regression coefficients they pertain to, rather than labeling the blocks with ordinal indices. For instance, we write $\mathbf{Q}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_{k-1}}]$ to refer to the sub-matrix of \mathbf{Q} that corresponds to $\beta_{i_0 \dots i_k}$ and $\beta_{i_0 \dots i_{k-1}}$ in the nested multilevel model. In the nested specification (20), $\beta_{i_0 \dots i_k}$ and $\beta_{i_0 \dots i_{k'}}$ are conditionally independent unless $k' \in \{k - 1, k, k + 1\}$, hence $\mathbf{Q}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_{k'}}] = \mathbf{0}\mathbf{0}^\top$. Accordingly, \mathbf{Q} is sparse in the sense that it only has $\mathcal{O}(p)$ non-zero elements. These are located in the following positions (with the obvious additional non-zero blocks due to symmetry):

$$\begin{aligned} \mathbf{Q}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_{k-1}}] &= -\mathbf{T}_{i_0 \dots i_{k-1}} \mathbf{A}_{i_0 \dots i_k}, \\ \mathbf{Q}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_k}] &= \begin{cases} \mathbf{T}_{i_0 \dots i_{k-1}} + \tau_{i_0 \dots i_k} \mathbf{X}_{i_0 \dots i_k}^\top \mathbf{X}_{i_0 \dots i_k} & (k < K) \\ + n_{i_0 \dots i_k} \mathbf{A}_{i_0 \dots i_{k+1}}^\top \mathbf{T}_{i_0 \dots i_k} \mathbf{A}_{i_0 \dots i_{k+1}} & \\ \mathbf{T}_{i_0 \dots i_{k-1}} + \tau_{i_0 \dots i_k} \mathbf{X}_{i_0 \dots i_k}^\top \mathbf{X}_{i_0 \dots i_k} & (k = K) \end{cases}, \end{aligned} \tag{24}$$

where we define $\mathbf{T}_{i_0 i_{-1}} = \mathbf{T}_{pr}$ for $k = 0$, and we recall that $n_{i_0 \dots i_k}$ is the number of offsprings of node $\beta_{i_0 \dots i_k}$ in the graphical model. With respect to (20)–(22), for nested multilevel models $\mathbf{V}\mathbf{m}$ is a vector of zeros except for the location that corresponds to β_0 , in which case it takes the value $\mathbf{T}_{pr} \boldsymbol{\mu}_{pr}$. Thus, the top equation in (23) is simplified accordingly. Finally, as defined in (22), $\mathbf{R}\mathbf{y}$ can trivially be computed at $\mathcal{O}(\max\{n, p\})$ cost.

The Cholesky factor \mathbf{L} does not necessarily inherit the sparsity of \mathbf{Q} , and it is known that the ordering of the indices plays a role in the sparsity of the Cholesky factor. This role can be understood through the concept of future sets, see for example in Theorem 2.8 of [29]. Therefore, the computational performance of linear algebra methods for sampling from Gaussian distributions depends strongly on the way β and \mathbf{Q} are organized. We now define a specific ordering that turns out to be optimal for nested multilevel models.

Definition 4. For nested multilevel models as defined in (20), we define the “depth-last” ordering, according to which β is organized by placing β_0 last, then the $\beta_{i_0 i_1}$ ’s (the order among them is irrelevant), then the $\beta_{i_0 i_1 i_2}$ ’s, and so on, until we reach the level furthest from the root.

For nested multilevel models, Proposition 3 shows that under the “depth-last” ordering the Cholesky factor \mathbf{L} has the same sparsity as \mathbf{Q} , i.e. it has $\mathcal{O}(p)$ non-zero elements, and that it may be computed at $\mathcal{O}(p)$ cost. Therefore, (23) can be used at a $\mathcal{O}(\max\{n, p\})$ cost to draw samples from $p(\beta \mid \mathbf{y}, \gamma)$. The proposition shows how the non-zero elements can be computed sequentially. However, the main consequence of the Proposition is to establish the sparsity of \mathbf{L} under the given ordering. Any sparse linear algebra software (e.g. CHOLMOD/SuiteSparse) will efficiently compute \mathbf{L} from a representation of \mathbf{Q} as a sparse matrix with “depth-last” ordering, hence the user would not need to hard-code the iterations described in the proposition.

Proposition 3. For the nested multilevel model as defined in (20) and (21), and under the “depth-last” ordering defined in Definition 4, let \mathbf{Q} be the corresponding posterior precision matrix and \mathbf{L} its associated Cholesky factor. Then, \mathbf{L} has the same sparsity structure as \mathbf{Q} and its only non-zero blocks are computed according to

$$\begin{aligned} \mathbf{L}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_k}]^\top \mathbf{L}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_{k-1}}] &= -\mathbf{T}_{i_0 \dots i_{k-1}} \mathbf{A}_{i_0 \dots i_k}, \\ \mathbf{L}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_k}] \mathbf{L}[\beta_{i_0 \dots i_k}, \beta_{i_0 \dots i_k}]^\top &= \mathbf{C}_{i_0 \dots i_k} + \mathbf{T}_{i_0 \dots i_{k-1}}, \end{aligned} \tag{25}$$

where as before $\mathbf{T}_{i_0 \dots i_{k-1}} = \mathbf{T}_{pr}$ for $k = 0$. The matrices $\mathbf{C}_{i_0 \dots i_k}$ are defined according to the recursion

$$\begin{aligned} \mathbf{C}_{i_0 \dots i_k} &= \begin{cases} \tau_{i_0 \dots i_K} \mathbf{X}_{i_0 \dots i_K}^\top \mathbf{X}_{i_0 \dots i_K} & (k = K) \\ \sum_{j_{k+1}} \tilde{\mathbf{C}}_{i_0 \dots i_k j_{k+1}} + \tau_{i_0 \dots i_k} \mathbf{X}_{i_0 \dots i_k}^\top \mathbf{X}_{i_0 \dots i_k} & (k < K) \end{cases}, \\ \tilde{\mathbf{C}}_{i_0 \dots i_k} &= \mathbf{A}_{i_0 \dots i_k}^\top \mathbf{B}_{i_0 \dots i_k}^\top (\mathbf{B}_{i_0 \dots i_k} \mathbf{T}_{i_0 \dots i_{k-1}}^{-1} \mathbf{B}_{i_0 \dots i_k}^\top + \mathbf{I})^{-1} \mathbf{B}_{i_0 \dots i_k} \mathbf{A}_{i_0 \dots i_k}, \end{aligned} \tag{26}$$

where $\mathbf{B}_{i_0 \dots i_k}$ follows from the decomposition $\mathbf{C}_{i_0 \dots i_k} = \mathbf{B}_{i_0 \dots i_k}^\top \mathbf{B}_{i_0 \dots i_k}$, and the notation $\sum_{j_{k+1}} \tilde{\mathbf{C}}_{i_0 \dots i_k j_{k+1}}$ is to be understood as summing over the children of node $i_0 \dots i_k$.

The cost of carrying out the recursion in Proposition 3 is dominated by $\mathcal{O}(p)$ decompositions of $L \times L$ -dimensional matrices, where we recall that L is the dimension of the regression parameters at each node. Therefore, the non-zero

elements are obtained at $\mathcal{O}(pL^3)$ cost. The Proposition can be proved in different ways, one using a future set argument, and another using *belief propagation*. In fact, the sequence of computational steps induced by sparse linear algebra based on “depth-last” ordering is almost identical to the one induced by belief propagation algorithms applied to the joint distributions defined by the posterior of the nested multilevel model. Belief propagation, as well as the general junction tree algorithm, is a framework for computations in graphical models, see Chapter 6 of [6] and Chapter 2 of [38], but for nested multilevel models the two paradigms are essentially equivalent. For an introduction to belief propagation see Section 8.4 in [4]. For details on how belief propagation can be used for nested multilevel models see the technical report [24], which also provides a historical perspective on this and links to various references. For other references discussing scalable sparse linear algebra methods to perform matrix factorisation for sparse matrices with non-zero elements corresponding to a tree see e.g. [35, Section 13.2] and references therein. Given that the proof can be worked out on the basis of one of these established approaches, we omit it from the present article.




FIG 5. [Nested effects real estate model, Gaussian likelihood] Comparison of sampling efficiency between MCMC based on Gibbs with SLA (blue) and Stan/NUTS (orange). The left panel overlays ACFs of β and γ as a function of wall time. The right panel shows the distribution of effective samples per second. Due to the large number of parameters, we only plot ACFs for a subset. For each algorithm, we include ACFs for the 20 slowest parameters and a representative subsample of the rest. Estimates are computed from MCMC runs with 10000 (SLA) and 1000 iterations (NUTS) respectively, with a warm-up period of 1000 iterations for NUTS. We set the hyperparameters $\nu_k = L$ for T_k , a flat prior on β_0 , and $\tau \sim \text{Gamma}(1/2, 1/2)$.

3.2. Methodology for non-Gaussian likelihoods: data augmentation and blocked Gibbs sampling

The direct sampling approach based on belief propagation or sparse linear algebra can still be useful to build scalable sampling algorithms even when the likelihoods are non-Gaussian. One such setting arises when conditionally on latent variables $z_{i_0 \dots i_k}$ associated to each observation $y_{i_0 \dots i_k}$, $\mathcal{L}(\beta \mid y, z, \gamma)$ is Gaussian. We have already discussed this approach with the binomial likelihood in Section 2.4, and we can transfer the result to the nested setting. Moreover,

heavier-tailed likelihoods, e.g. the Laplace or the Student-t, can be represented as scale mixtures of Gaussians, i.e.

$$\begin{aligned} \mathcal{L}(\mathbf{y}_{i_0 \dots i_k} \mid \boldsymbol{\eta}_{i_0 \dots i_k}) \\ = \int_0^\infty \mathcal{N}(\mathbf{y}_{i_0 \dots i_k}; \boldsymbol{\eta}_{i_0 \dots i_k}, (z_{i_0 \dots i_k} / \tau_{i_0 \dots i_k}) \mathbf{I}) p(z_{i_0 \dots i_k}) dz_{i_0 \dots i_k}. \end{aligned} \quad (27)$$

Letting \mathbf{z} denote the vector of all such latent variables, our proposed scheme is to work with the data augmented law $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma}, \mathbf{z} \mid \mathbf{y})$, and target it with a Gibbs sampler, where we sample from $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \boldsymbol{\gamma}, \mathbf{z})$ using the methodology of Section 3.1. $\mathcal{L}(\mathbf{z} \mid \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\gamma})$ decomposes according to

$$\mathcal{L}(\mathbf{z} \mid \mathbf{y}, \boldsymbol{\beta}, \boldsymbol{\gamma}) = \prod_{k=0}^K \prod_{i_0 \dots i_k} \mathcal{L}(z_{i_0 \dots i_k} \mid \mathbf{y}_{i_0 \dots i_k}, \boldsymbol{\beta}, \boldsymbol{\gamma}), \quad (28)$$

with each $z_{i_0 \dots i_k}$ easily updated.

Various important specifications fall outside of this class, such as the Poisson regression model $\mathcal{L}(y_{i_0 \dots i_k} \mid \eta_{i_0 \dots i_k}) = \text{Pois}(e^{\eta_{i_0 \dots i_k}})$. This instance is most easily described when the data \mathbf{y} is only observed at the leaf level K , which in any case is the most common. Defining $\boldsymbol{\beta}^{(K)}$ as the set of all regression parameters $\beta_{i_0 \dots i_K}$ at level K , we suggest the strategy of alternating updates to $\mathcal{L}(\boldsymbol{\beta}^{(-K)} \mid \boldsymbol{\beta}^{(K)}, \boldsymbol{\gamma})$ and $\mathcal{L}(\boldsymbol{\beta}^{(K)} \mid \mathbf{y}, \boldsymbol{\beta}^{(K-1)}, \boldsymbol{\gamma})$, noting that $\boldsymbol{\beta}^{(-K)}$ is conditionally independent of \mathbf{y} . Indeed, $\boldsymbol{\beta}^{(-K)}$ is now structured as a tree with $K-1$ levels and Gaussian data $\boldsymbol{\beta}^{(K)}$ at the new leaves. Thus, we sample from the full conditional distribution law of $\boldsymbol{\beta}^{(-K)}$ according to the methodology of Section 3.1. Conversely, while $\mathcal{L}(\boldsymbol{\beta}^{(K)} \mid \mathbf{y}, \boldsymbol{\beta}^{(K-1)}, \boldsymbol{\gamma})$ is not Gaussian, it decomposes according to

$$\mathcal{L}(\boldsymbol{\beta}^{(K)} \mid \mathbf{y}, \boldsymbol{\beta}^{(K-1)}, \boldsymbol{\gamma}) = \prod_{i_0 \dots i_K} \mathcal{L}(\beta_{i_0 \dots i_K} \mid \mathbf{y}_{i_0 \dots i_K}, \boldsymbol{\beta}^{(K-1)}, \boldsymbol{\gamma}), \quad (29)$$

with each $\beta_{i_0 \dots i_K}$ updated independently by a Metropolis-within-Gibbs step akin to the one presented in Section 2.2, at a total cost of order $\mathcal{O}(p)$ for the full update from $\mathcal{L}(\boldsymbol{\beta}^{(K)} \mid \mathbf{y}, \boldsymbol{\beta}^{(K-1)}, \boldsymbol{\gamma})$. While the cost of the combined update remains $\mathcal{O}(p)$, the main potential weakness of that strategy lies in the induced dependence between $\boldsymbol{\beta}^{(K)}$ and $\boldsymbol{\beta}^{(-K)}$, which in some settings increases the relaxation time of the algorithm as p increases. The work in [23] and the associated theory in [21] suggests that for likelihood with tails lighter than the Gaussian, such as the Poisson, this centered parameterisation (in the terminology of [23]) is expected to perform well, whereas for heavy-tailed likelihoods it would not. Hence, we view the two strategies proposed in this Section as complementary.

3.3. Numerical experiments with a real estate model

We assess algorithmic performance on a semi-synthetic dataset that originates from a broader task of temporal prediction of real estate prices at high spatial




FIG 6. [Nested effects real estate model, binomial likelihood] *Sampling efficiency of the PG/SLA algorithm for binomial likelihoods. The plot overlays ACFs of β and γ as a function of MCMC lag. Due to the large number of parameters, we only plot ACFs for a subset: we include ACFs for the 20 slowest parameters and a representative subsample of the rest. Estimates are computed from an MCMC run with 10000 iterations. We set the hyperparameters $\nu_k = L$ for T_k and $(\mu_{pr} = \mathbf{0}, T_{pr} = \mathbf{I})$ for β_0 , corresponding to a standard Gaussian prior.*

resolution. Our approach combines a type of multilevel capital asset pricing model, which relates the local average square meter price of different type of real estate properties to national economic indices (e.g. GDP), with a multivariate time series model for these indices. In this article we focus on the multilevel regression component of this model. The spatial domain is Spain, organized according to 5-digit postal codes which define $K = 4$ levels, plus the root. The first two digits define a province in the country and the later codes correspond to higher spatial resolutions. There are 44 observations at each postal code, each of which is the difference in consecutive quarters of the logarithm of local average sales price. The quarters span the period 2007–2018. The predictors, an intercept and a national housing price index (hence $L = 2$), are only available at the leaf-level. In this data structure, there are $p \approx 8 \times 10^3$ regression parameters and $n \approx 4 \times 10^6$ observations. Due to confidentiality constraints, we work with a dataset that has been simulated from a misspecified version of (20), using the part of the dataset that is publicly available and parameter values β that are consistent with the real data.

In two further experiments, we adopt the latent variable augmentation setting of Section 3.2, but keep the same prior structure as in the Gaussian experiment. In one instance, we transform the response into a binary variable according to its sign, set $\mathcal{L}(y_{i_0 \dots i_K t} | \eta_{i_0 \dots i_K t}) = \text{Binomial}(1, (1 + \exp\{-\eta_{i_0 \dots i_K t}\})^{-1})$, and use Polya-gamma augmentation to obtain conditionally Gaussian observations, as in Section 2.4. We show results in Figure 6. Remarkably, the elements of β appear to mix instantly, while elements of γ have significantly larger IATs. These are ordered by the level that the parameter pertains to: the higher the level, the slower the mixing. This is due to the rising number of coefficients per hyperparameter, and therefore increasing strength of the prior. This could be remedied, e.g. by way of parameter expansion. In the other instance, we postulate a heterogeneous observation model where $\mathcal{L}(y_{i_0 \dots i_K t} | \eta_{i_0 \dots i_K t}) = \text{Cauchy}(\eta_{i_0 \dots i_K t}, \tau^{-1/2})$ for a proportion α of the terminal leaves, and $\mathcal{N}(\eta_{i_0 \dots i_K t}, \tau^{-1})$ otherwise. Here,




FIG 7. [Nested effects real estate model, Cauchy likelihood] Comparison of sampling efficiency of the GMix/SLAwG algorithm for different proportions $\alpha = 0.1, 0.2, \dots, 1$ of heavy-tailed observation densities. The panel shows ACFs of the model likelihood for proportions ranging from $\alpha = 0$ (lightest) to $\alpha = 1$ (darkest). Estimates are computed from an MCMC run with 10000 iterations. We set the hyperparameters $\nu_k = L$ for T_k , a flat prior on β_0 , and $\tau \sim \text{Gamma}(1/2, 1/2)$.

the t subscript denotes replications over times for $t = 1, \dots, 44$. We vary $\alpha = 0.1, 0.2, \dots, 1$, with $\alpha = 0$ corresponding to the model and associated algorithm that is analyzed in Figures 5. For observations $y_{i_0 \dots i_K t}$ with non-Gaussian observation density we augment with latent variables $z_{i_0 \dots i_K t}$, given which the density is Gaussian, and proceed with the blocked Gibbs sampler outlined above for each of the resulting 10 models. In Figure 7 we plot the estimated ACF for the likelihood of the data (marginalizing over the latent variables \mathbf{z} , where present). The experiment shows that as the proportion of latent variables increases the mixing of the algorithm slows down, but the deterioration is bounded. In particular, the autocorrelation is almost negligible after 10 iterations in all settings. Considering the large value of n and p of this dataset, the results are coherent with the conjecture that the algorithm's relaxation time remains small as n and p diverge.

In a final experiment, the response $y_{i_0 \dots i_K t}$ is the number of transactions at the location $i_0 \dots i_K$ and time t , and we model it as $\text{Pois}(\exp\{\eta_{i_0 \dots i_K t}\})$. There is a trend in the transaction numbers, which we account for by introducing the time index t as an additional covariate, and therefore $L = 3$. We apply our strategy for light-tailed likelihoods as laid out in Section 3.2, with updates to $\beta_{i_0 \dots i_K}$ proposed according to a second order expansion of $\mathcal{L}(\beta_{i_0 \dots i_K} \mid \mathbf{y}_{i_0 \dots i_K}, \beta^{(K-1)}, \gamma)$. We show results in Figure 8. The algorithm achieves good mixing throughout the tree, with the exception of a few nodes located on the penultimate level $L = 3$. Our investigation suggests that this is due to the fact that the model we use for the data is purposely mis-specified and that this creates the slower, although by no means slow, mixing of a few internal nodes. In fact, the mis-specification creates a non-trivial dependence between precision hyperparameters and a few regression parameters.




FIG 8. [Nested effects real estate model, Poisson likelihood] *Sampling efficiency of the LT/SLAwG algorithm for Poisson likelihoods. The plot overlays ACFs of β and γ as a function of MCMC lag. Due to the large number of parameters, we only plot ACFs for a subset: we include ACFs for the 20 slowest parameters and a representative subsample of the rest. Estimates are computed from an MCMC run with 10000 iterations. We set the hyperparameters $\nu_k = L$ for \mathbf{T}_k and $(\mu_{pr} = \mathbf{0}, \mathbf{T}_{pr} = \mathbf{I})$ for β_0 , corresponding to a standard Gaussian prior.*

4. Discussion

We have developed efficient sampling algorithms for large scale Bayesian inference with two of the most popular families of models for applied research, crossed effects and nested multilevel models with Gaussian priors and non-Gaussian likelihoods, and we have obtained rigorous results on their scalability properties. Our results support the claim that appropriately designed blocked Gibbs samplers can effectively exploit the sparse conditional dependence structure inherent to these models in order to achieve an overall computational complexity that is $\mathcal{O}(\max\{n, p\})$. From a practical point of view, these algorithms are observed to provide dramatic computational speed-ups compared to off-the-shelf HMC, both for simulated data and in challenging applications.

There are various open questions and directions for future research that would complement and extend the present work, two of which we elaborate on. Firstly, we worked under the assumption that scalable sampling from the conditional distribution of the high-dimensional regression parameters, i.e. $\mathcal{L}(\mathbf{a} \mid \mathbf{y}, \boldsymbol{\gamma})$ and $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \boldsymbol{\gamma})$ for crossed and nested models respectively, is sufficient to scalably sample from the joint distributions $\mathcal{L}(\mathbf{a}, \boldsymbol{\gamma} \mid \mathbf{y})$ and $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma} \mid \mathbf{y})$, at least in commonly encountered scenarios. While this is strongly supported by simulations, theoretical results on when the assumption is met, and when it isn't, would further strengthen our conclusions. A first step in this direction is made in [2], which analyses the two-block Gibbs sampler targeting the joint distribution $\mathcal{L}(\boldsymbol{\beta}, \boldsymbol{\gamma} \mid \mathbf{y})$ for two-level nested hierarchical models. The theory developed therein is asymptotic and supports the claim that under appropriate assumptions, the convergence time of the above two block sampler remains bounded as $p, n \rightarrow \infty$. Secondly, while we demonstrated that SLA methods are ideally suited to nested models, it turns out that they are not effective for crossed effect models. In ongoing work we are performing a systematic exploration of SLA for

crossed models, both to substantiate this claim, as well as to assess the effectiveness of *approximate* SLA methods, which sample from an approximation to $\mathcal{L}(\boldsymbol{\beta} \mid \mathbf{y}, \boldsymbol{\gamma})$.

Appendix A: Proofs

Proof of Theorem 1. The cost per iteration has been established to be $\mathcal{O}(2n + p) = \mathcal{O}(\max\{n, p\})$, hence the proof focuses on quantifying the relaxation time. Denote the relaxation time of (cGS) by T_{cgs} . We now prove that

$$c \min\{\bar{n}, T_{\text{aux}}\} \leq T_{\text{cgs}} \leq C \min\{\bar{n}, T_{\text{aux}}\} \quad (30)$$

for any balanced level design and for some positive and finite constants c and C , which depend only on $\boldsymbol{\gamma}$. The thesis follows then directly from (30).

By [22, Theorem 4] we have $T_{\text{cgs}} = (1 - \rho_1 \rho_2 \rho_{\text{aux}})^{-1}$ where $\rho_k = (n\tau + p_k \tau_k)^{-1} n\tau$ for $k = 1, 2$; ρ_{aux} is the L^2 -rate of convergence of the auxiliary Gibbs sampler satisfying $T_{\text{aux}} = (1 - \rho_{\text{aux}})^{-1}$, τ is the likelihood precision and τ_k the prior precision of the k -th factor. Since $\rho \mapsto (1 - \rho)^{-1}$ is increasing on $(0, 1)$ and $\rho_1 \rho_2 \rho_{\text{aux}} \leq \min\{\rho_1, \rho_2, \rho_{\text{aux}}\}$, it follows that

$$T_{\text{cgs}} \leq (1 - \min\{\rho_1, \rho_2, \rho_{\text{aux}}\})^{-1} = \min\{(1 - \min\{\rho_1, \rho_2\})^{-1}, (1 - \rho_{\text{aux}})^{-1}\}.$$

By definition of ρ_1 and ρ_2 ,

$$\begin{aligned} (1 - \min\{\rho_1, \rho_2\})^{-1} &= 1 + \frac{n\tau}{\max\{p_1 \tau_1, p_2 \tau_2\}} \\ &\leq 1 + C_1 \frac{n}{\max\{p_1, p_2\}} \\ &\leq 1 + C_1 \bar{n}, \end{aligned} \quad (31)$$

with $C_1 = \frac{\tau}{\min\{\tau_1, \tau_2\}}$ and exploiting $\max\{p_1, p_2\} \geq 2^{-1}(p_1 + p_2)$. Combining the above inequalities with $T_{\text{aux}} = (1 - \rho_{\text{aux}})^{-1}$, we obtain

$$T_{\text{cgs}} \leq \min\{1 + C_1 \bar{n}, T_{\text{aux}}\} \leq C \min\{\bar{n}, T_{\text{aux}}\},$$

where we can take $C = 1 + C_1$ since $\bar{n} \geq 1$ by construction.

For the lower bound, $\rho_1 \rho_2 \rho_{\text{aux}} \geq (\min\{\rho_1, \rho_2, \rho_{\text{aux}}\})^3$ implies

$$T_{\text{cgs}} \geq (1 - \min\{\rho_1, \rho_2, \rho_{\text{aux}}\}^3)^{-1} \geq 3^{-1} (1 - \min\{\rho_1, \rho_2, \rho_{\text{aux}}\})^{-1},$$

where the second inequality follows from $(1 - \rho^t)^{-1} \geq t^{-1}(1 - \rho)^{-1}$ for every $t \geq 1$ and $\rho \in (0, 1)$. To prove the latter claim note that $(1 - \rho^t)^{-1} \geq t^{-1}(1 - \rho)^{-1}$ if and only if $\rho^t - t\rho + (t - 1) \geq 0$, and that the left-hand side is a decreasing function of ρ on $(0, 1)$ that vanishes when $\rho = 0$. It follows that

$$T_{\text{cgs}} \geq 3^{-1} \min \left\{ (1 - \min\{\rho_1, \rho_2\})^{-1}, (1 - \rho_{\text{aux}})^{-1} \right\}.$$

Then, analogously to (31),

$$\begin{aligned} (1 - \min\{\rho_1, \rho_2\})^{-1} &\geq \frac{n\tau}{\max\{p_1\tau_1, p_2\tau_2\}} \\ &\geq c_1 \frac{n}{\max\{p_1, p_2\}} \\ &\geq \frac{c_1}{2} \bar{n}, \end{aligned}$$

with $c_1 = \frac{\tau}{\max\{\tau_1, \tau_2\}}$ and exploiting $\max\{p_1, p_2\} \leq p_1 + p_2$. Thus, we obtain

$$T_{\text{cgs}} \geq \min \left\{ \frac{c_1}{2} \bar{n}, T_{\text{aux}} \right\} \geq c \min \{ \bar{n}, T_{\text{aux}} \},$$

where $c = \min\{1, c_1/2\}$ as desired. □

Proof of Proposition 2. Throughout the proof, $\{\bar{a}(t)\}$ denotes the chain of averages. We highlight that $a^{(0)}$ takes $K + 1$ intermediate values, denoted by $(a^{(0)}(t, k))_{k=0}^K$. The chain evolves according to the following scheme.

Given $\bar{a}(t) = (a^{(0)}(t), \bar{a}^{(1)}(t), \dots, \bar{a}^{(K)}(t))$ do:

Set $a^{(0)}(t, 0) \leftarrow a^{(0)}(t)$

for $k = 1 : K$ **do**

Draw $a^{(0)}(t, k)$ from

$$\mathcal{N} \left(a^{(0)}(t, k-1) + \bar{a}^{(k)}(t), (p_k \tau_k)^{-1} \right)$$

Draw $\bar{a}^{(k)}(t+1)$ from

$$\mathcal{N} \left(\rho_k \left(\bar{y} - a^{(0)}(t, k) - \sum_{\ell=1}^{k-1} \bar{a}^{(\ell)}(t+1) - \sum_{\ell=k+1}^K \bar{a}^{(\ell)}(t) \right), (n\tau + p_k \tau_k)^{-1} \right)$$

end for

Set $a^{(0)}(t+1) \leftarrow a^{(0)}(t, K)$

By construction

$$\begin{aligned} \mathbb{E}[a^{(0)}(t, k) | \bar{a}(t)] &= \mathbb{E}[\mathbb{E}[a^{(0)}(t, k) | a^{(0)}(t, k-1), \bar{a}(t)] | \bar{a}(t)] \\ &= \mathbb{E}[a^{(0)}(t, k-1) | \bar{a}(t)] + \bar{a}^{(k)}(t), \end{aligned}$$

which by induction over k implies that $\mathbb{E}[a^{(0)}(t, k) | \bar{a}(t)] = a^{(0)}(t) + \sum_{\ell=1}^k \bar{a}^{(\ell)}(t)$ for $k = 1, \dots, K$ and thus $\mathbb{E}[a^{(0)}(t+1) | \bar{a}(t)] = a^{(0)}(t) + \sum_{\ell=1}^K \bar{a}^{(\ell)}(t)$. Moreover,

for $k = 1, \dots, K$

$$\begin{aligned}
 \mathbb{E}[\bar{a}^{(k)}(t)|\bar{\mathbf{a}}(t)] &= \rho_k(\bar{y} - \mathbb{E}[a^{(0)}(t, k)|\bar{\mathbf{a}}(t)]) + \sum_{\ell=k+1}^K \bar{a}^{(\ell)}(t) \\
 &\quad + \sum_{\ell=1}^{k-1} \mathbb{E}[\bar{a}^{(\ell)}(t+1)|\bar{\mathbf{a}}(t)] \\
 &= -\rho_k(a^{(0)}(t) + \sum_{\ell=1}^K \bar{a}^{(\ell)}(t)) \\
 &\quad + \sum_{\ell=1}^{k-1} \mathbb{E}[\bar{a}^{(\ell)}(t+1)|\bar{\mathbf{a}}(t)] + \text{const},
 \end{aligned} \tag{32}$$

where we use “const” to indicate fixed values that do not depend on $\bar{\mathbf{a}}(t)$. It follows from (32) that $\mathbb{E}[\bar{a}^{(1)}(t)|\bar{\mathbf{a}}(t)] = -\rho_1 \left(a^{(0)}(t) + \sum_{\ell=1}^K \bar{a}^{(\ell)}(t) \right) + \text{const}$ and, by induction over $k = 1, \dots, K$,

$$\mathbb{E}[\bar{a}^{(k)}(t)|\bar{\mathbf{a}}(t)] = c_k \left(a^{(0)}(t) + \sum_{\ell=1}^K \bar{a}^{(\ell)}(t) \right) + \text{const},$$

with $(c_k)_{k=1}^K$ satisfying $c_k = -\rho_k(1 + \sum_{\ell=1}^{k-1} c_\ell)$. Therefore, the transition from $\bar{\mathbf{a}}(t)$ to $\bar{\mathbf{a}}(t+1)$ corresponds to a linear autoregressive process with a $(K+1) \times (K+1)$ autoregressive matrix \mathbf{B} . By well-known theory for Gaussian autoregression processes, the rate of convergence of such a Markov chain equals the largest modulus eigenvalue of the matrix \mathbf{B} , which we denote by $\rho(\mathbf{B})$. Defining $c_0 = 1$, the above results imply that $B_{k\ell} = c_k$ for all $k, \ell = 0, \dots, K$, meaning that the $K+1$ columns of \mathbf{B} are all equal to $(c_0, \dots, c_K)^T$. It follows that \mathbf{B} is a rank 1 matrix with K eigenvalues equal to 0 and one non-zero eigenvalue equal to $\sum_{k=0}^K c_k$. Hence, to bound $\rho(\mathbf{B})$ we need to upper bound the absolute value of $\sum_{k=0}^K c_k$. Define $d_k = \sum_{\ell=1}^k c_\ell$ for $k = 1, \dots, K$. We now prove by induction that: $d_k \in (-1, 0)$, $d_k \leq -\rho_k$ and d_k is monotonously decreasing in $k = 1, \dots, K$. The hypothesis is true for $k = 1$ since $d_1 = -\rho_1$ and $\rho_k \in (0, 1)$ for all $k = 1, \dots, K$. For $k \geq 2$, we have

$$d_k = \sum_{\ell=1}^k c_\ell = d_{k-1} + c_k = d_{k-1} - \rho_k(1 + d_{k-1}) = (1 - \rho_k)d_{k-1} - \rho_k.$$

Since $(1 - \rho_k) \in (0, 1)$ and $d_{k-1} \in (-1, 0)$, it follows that $(1 - \rho_k)d_{k-1} \in (-1, 0)$ and $d_k \leq -\rho_k$. Also, being a convex combination of $d_{k-1} \in (-1, 0)$ and of -1 , it follows $d_k \in (-1, 0)$ and $d_k < d_{k-1}$. This proves the induction hypothesis. Since $d_K \in (-1, 0)$ we have $\sum_{k=0}^K c_k = 1 + d_K > 0$. By monotonicity of d_k and $d_k \leq -\rho_k$ for $k = 1, \dots, K$, it follows that $d_k \leq \min_{k=1, \dots, K} -\rho_k$. Thus

$\sum_{k=0}^K c_k = 1 + d_K \leq \min_{k=1, \dots, K} (1 - \rho_k)$. Accordingly,

$$\rho(\mathbf{B}) = \sum_{k=0}^K c_k \leq \min_{k=1, \dots, K} 1 - \rho_k = \min_{k=1, \dots, K} \frac{p_k \tau_k}{n\tau + p_k \tau_k}.$$

It follows that the relaxation time of the $(\bar{a}(t))_{t \geq 1}$ chain equals

$$\begin{aligned} (1 - \rho(\mathbf{B}))^{-1} &\leq \left(1 - \min_{k=1, \dots, K} \frac{p_k \tau_k}{n\tau + p_k \tau_k} \right)^{-1} \\ &= 1 + \min_{k=1, \dots, K} \frac{p_k \tau_k}{n\tau} \\ &\leq 1 + \frac{\tau_{\min}}{\tau} = \mathcal{O}(1), \end{aligned}$$

where $\tau_{\min} = \min_{k=1, \dots, K} \tau_k$ is a fixed constant in our analysis. Above we used $p_k \leq n$ for all $k = 1, \dots, K$, which holds by construction as every level of every factor gets at least one observation. \square

Proof of Theorem 2. By the argument in the proof of [22, Prop. 4], we obtain that

$$T_{\text{aux}} = (1 - \lambda_2^*(P_1 P_2))^{-1}, \tag{33}$$

where $P_1 = n^{(1,2)}/d_1$ and $P_2 = (n^{(1,2)})^T/d_2$ are, respectively, $p_1 \times p_2$ and $p_2 \times p_1$ stochastic matrices; and $\lambda_2^*(M)$ denotes the second largest modulus eigenvalue of a matrix M . The spectrum of the $p_1 \times p_1$ transition matrix $P_1 P_2$ is related to the one of the $(p_1 + p_2) \times (p_1 + p_2)$ matrix

$$A = \left(\begin{array}{c|c} \mathbf{0}_{p_1 \times p_2} & n^{(1,2)} \\ \hline (n^{(1,2)})^T & \mathbf{0}_{p_2 \times p_1} \end{array} \right), \tag{34}$$

where $\mathbf{0}_{m \times r}$ denotes a zero matrix of dimension $m \times r$. The matrix A can be interpreted as the adjacency matrix of a bipartite graph with $p_1 + p_2$ nodes, one for each parameter in $\mathbf{a}^{(1)}$ and $\mathbf{a}^{(2)}$, and edges between $a_i^{(1)}$ and $a_{i'}^{(2)}$ if and only if $n_{ii'}^{(1,2)} = 1$. The definitions of A and $P_1 P_2$ imply that

$$\lambda_2^*(P_1 P_2) = \lambda_2(A) / \sqrt{d_1 d_2}, \tag{35}$$

where $\lambda_2(M)$ denotes the second largest eigenvalue of a matrix M ¹.

The bipartite graph associated to A , denoted by $G(A)$, is also bi-regular, or more specifically (d_1, d_2) -regular, meaning that the degree of all nodes in the k -th component is d_k , for $k \in \{1, 2\}$. Moreover, it can be seen that the map from $n^{(1,2)}$ to $G(A)$, with A as in (34), defines a bijection from $\mathcal{D}(n, d_1, d_2)$ to $\mathcal{G}(n, d_1, d_2)$, where the latter denotes the collection of all (d_1, d_2) -regular bipartite graphs with a total of n edges. It follows that, when $n^{(1,2)} \sim \text{Unif}(\mathcal{D}(n, d_1, d_2))$ and A

¹More generally the whole spectrum of $P_1 P_2$ and A are related: assuming $p_1 < p_2$ without loss of generality, if λ is an eigenvalue of $P_1 P_2$, then $\sqrt{d_1 d_2} \lambda$ and $-\sqrt{d_1 d_2} \lambda$ are eigenvalues of A , and the remaining $p_2 - p_1$ eigenvalues of A are zero. See e.g. [5] for details.

is as in (34), then $G(A) \sim \text{Unif}(\mathcal{G}(n, d_1, d_2))$. Thus, by (35), we need to study the second largest eigenvalue of the adjacency matrix of a random bipartite bi-regular graph. This has been studied in [5]. Theorem 3.2 therein implies that $\lambda_2(A) \leq \sqrt{d_1 - 1} + \sqrt{d_2 - 1} + \epsilon'$ almost surely as $n \rightarrow \infty$ for any $\epsilon' > 0$, for fixed d_1 and d_2 larger or equal than 3.² Combining the latter with (33) and (35) and performing some basic manipulations we obtain that, for every $\epsilon > 0$, it holds almost surely as $n \rightarrow \infty$ that

$$\begin{aligned} T_{\text{aux}} &\leq \frac{\sqrt{d_1 d_2}}{\sqrt{d_1 d_2} - \sqrt{d_1 - 1} - \sqrt{d_2 - 1}} + \epsilon \\ &\leq 1 + \frac{\sqrt{d_1 - 1} + \sqrt{d_2 - 1}}{\sqrt{d_1 d_2} - \sqrt{d_1 - 1} - \sqrt{d_2 - 1}} + \epsilon \\ &\leq 1 + \frac{2\sqrt{\max\{d_1, d_2\}}}{\sqrt{d_1 d_2} - 2\sqrt{\max\{d_1, d_2\}}} + \epsilon = 1 + \frac{2}{\sqrt{\min\{d_1, d_2\}} - 2} + \epsilon, \end{aligned}$$

where we used $d_1, d_2 > 4$ to ensure that $\sqrt{\min\{d_1, d_2\}} - 2 > 0$. The condition $d_1, d_2 > 4$ is purely assumed for simplicity in order to allow for a more concise and interpretable bound on T_{aux} . This proves (18). Equation (19) then follows directly combining (17) with $T_{\text{aux}} \leq 1 + \frac{2}{\sqrt{\min\{d_1, d_2\}} - 2} \leq 1 + \frac{2}{\sqrt{5} - 2} < \infty$. \square

Appendix B: MCMC performance metrics

The *integrated autocorrelation time* (IAT), is defined as

$$\text{iat } X = 1 + 2 \sum_{t=1}^{\infty} \text{acf}_t X, \quad (36)$$

where X is a stochastic process and $\text{acf}_t X$ is the autocorrelation function (ACF) of X evaluated at t -th lag. In practice, we estimate ACF by empirical averages from a single trajectory of X . We then use these estimates to estimate the IAT following Sokal's truncation heuristic [32]. The heuristic consists of computing a running sum over ACF lags and stopping as soon as the running sum exceeds a multiple of the next lag. We modify this heuristic slightly and only truncate at even lags to account for NUTS' tendency to produce estimates of ACF of different sign at successive lags. We define the *effective sample size* (ESS) as the run length divided by the IAT estimate. Since we are comparing methods with vastly different computing time per iteration, we monitor the run time of the algorithm for both algorithms and use *effective sample size per second* as our performance metric. Notice that we include transience time in the ESS/sec estimate. Moreover, when overlaying ACFs for both methods, we plot them as a function of computing time, as opposed to iteration lag.

²The condition $d_1, d_2 \geq 3$ is assumed at the beginning of [5] and not explicitly recalled in Theorem 3.2.

Appendix C: Table of Algorithms

TABLE 2
Algorithms and corresponding acronyms.

Acronym	Structure	Summary	Appears In
LC-MwG	Crossed	Local centering	Figures 1, 4
Van-MwG	Crossed	Vanilla Metropolis-within-Gibbs	Figure 1
const/LC-MwG	Crossed	Local centering w/ ident. constr.	Figures 2, 3
const/Van-MwG	Crossed	Vanilla Gibbs w/ ident. constr.	Figure 2
free/LC-MwG	Crossed	Local centering w/o ident. constr.	Figure 3
free/PLC-MwG	Crossed	Proj. local centering w/o ident. constr.	Figure 3
PG/Van-G	Crossed	Polya-gamma augm. vanilla Gibbs	Figure 4
PG/Col-G	Crossed	Polya-gamma augm. collapsed Gibbs	Figure 4
SLA	Nested	Gaussian sparse lin. alg.	Figure 5
PG/SLAwG	Nested	Polya-gamma augm. sparse lin. alg.	Figure 6
GMix/SLAwG	Nested	Gaussian mixture sparse lin. alg.	Figure 7
LT/SLAwG	Nested	Light-tailed likelihood sparse lin. alg.	Figure 8

Acknowledgments

The authors are most grateful to Jie Hao Kwa, whose master’s dissertation in 2018 yielded useful insights into sparse linear algebra methods for nested multilevel models, and Maximilian Müller, whose master’s dissertation in 2020 investigated sparse linear methods for crossed effect models. The first two authors would like to thank Jose Garcia-Montalvo for the collaboration on the applied projects that have motivated this work. The article has benefited from comments by Darren Wilkinson.

Funding

The third author acknowledges support from the European Research Council (ERC), through StG “PrSc-HDBayLe” grant ID 101076564.

References

- [1] ANDRIEU, C. and THOMS, J. (2008). A tutorial on adaptive MCMC. *Statistics and Computing* **18** 343–373. [MR2461882](#)
- [2] ASCOLANI, F. and ZANELLA, G. (2023). Complexity of Gibbs samplers through Bayesian asymptotics. arXiv preprint [arXiv:2304.06993](#).
- [3] BATES, D., MÄCHLER, M., BOLKER, B. and WALKER, S. (2015). Fitting linear mixed-effects models using LME4. *Journal of Statistical Software* **67** 1–48. <https://doi.org/10.18637/jss.v067.i01>.
- [4] BISHOP, C. M. (2006). *Pattern Recognition and Machine Learning. Information Science and Statistics*. Springer, New York. <https://doi.org/10.1007/978-0-387-45528-0>. [MR2247587](#)

- [5] BRITO, G., DUMITRIU, I. and HARRIS, K. D. (2022). Spectral gap in random bipartite biregular graphs and applications. *Combinatorics, Probability and Computing* **31** 229–267. [MR4390294](#)
- [6] COWELL, R. G., DAWID, A. P., LAURITZEN, S. L. and SPIEGELHALTER, D. J. (1999). *Probabilistic Networks and Expert Systems. Statistics for Engineering and Information Science*. Springer-Verlag, New York. [MR1697175](#)
- [7] DWIVEDI, R., CHEN, Y., WAINWRIGHT, M. J. and YU, B. (2019). Log-concave sampling: Metropolis-Hastings algorithms are fast. *Journal of Machine Learning Research* **20** 1–42. [MR4048994](#)
- [8] GAO, K. and OWEN, A. (2017). Efficient moment calculations for variance components in large unbalanced crossed random effects models. *Electronic Journal of Statistics* **11** 1235–1296. [MR3635913](#)
- [9] GAO, K. and OWEN, A. B. (2020). Estimation and inference for very large linear mixed effects models. *Statistica Sinica* **30** 1741–1771. [MR4260743](#)
- [10] GELFAND, A. E., SAHU, S. K. and CARLIN, B. P. (1995). Efficient parametrisations for normal linear mixed models. *Biometrika* **82** 479–488. [MR1366275](#)
- [11] GELFAND, A. E., SAHU, S. K. and CARLIN, B. P. (1996). Efficient parametrizations for generalized linear mixed models. In *Bayesian Statistics, 5 (Alicante, 1994)*. *Oxford Sci. Publ.* 165–180. Oxford Univ. Press, New York. [MR1425405](#)
- [12] GELMAN, A. (2005). Analysis of variance: why it is more important than ever. *The Annals of Statistics* **33** 1–53. [MR2157795](#)
- [13] GELMAN, A. and HILL, J. (2007). *Data Analysis Using Regression and Multilevel/Hierarchical Models* **3**. Cambridge University Press New York, New York, USA.
- [14] GHOSH, S., HASTIE, T. and OWEN, A. B. (2020). Backfitting for large scale crossed random effects regressions. arXiv preprint [arXiv:2007.10612](#). [MR4382028](#)
- [15] GHOSH, S., HASTIE, T. and OWEN, A. B. (2022). Scalable logistic regression with crossed random effects. *Electronic Journal of Statistics* **16** 4604–4635. [MR4489236](#)
- [16] GOPLERUD, M. (2021). Fast and accurate estimation of non-nested binomial hierarchical models using variational inference. *Bayesian Analysis* 1–28. <https://doi.org/10.1214/21-BA1266>. [MR4483233](#)
- [17] LIU, J. S. and WU, Y. N. (1999). Parameter expansion for data augmentation. *Journal of the American Statistical Association* **94** 1264–1274. [MR1731488](#)
- [18] MENG, X.-L. and VAN DYK, D. A. (1999). Seeking efficient data augmentation schemes via conditional and marginal augmentation. *Biometrika* **86** 301–320. [MR1705351](#)
- [19] MENICTAS, M., DI CREDICO, G. and WAND, M. P. (2019). Streamlined variational inference for linear mixed models with crossed random effects. arXiv preprint [arXiv:1910.01799](#). [MR4552940](#)
- [20] MONTALVO, J. G., PAPASPILIOPOULOS, O. and STUMPF-FÉTIZON, T.

- (2019). Bayesian forecasting of electoral outcomes with new parties' competition. *European Journal of Political Economy* **59** 52–70.
- [21] PAPASPILIOPOULOS, O. and ROBERTS, G. (2008). Stability of the Gibbs sampler for Bayesian hierarchical models. *The Annals of Statistics* **36** 95–117. <https://doi.org/10.1214/009053607000000749>. MR2387965
- [22] PAPASPILIOPOULOS, O., ROBERTS, G. O. and ZANELLA, G. (2020). Scalable inference for crossed random effect models. *Biometrika* **107** 24–40. MR4064138
- [23] PAPASPILIOPOULOS, O., ROBERTS, G. O. and SKÖLD, M. (2007). A general framework for the parametrization of hierarchical models. *Statistical Science* 59–73. MR2408661
- [24] PAPASPILIOPOULOS, O. and ZANELLA, G. (2017). A note on MCMC for nested multilevel regression models via belief propagation.
- [25] POLSON, N. G., SCOTT, J. G. and WINDLE, J. (2013). Bayesian inference for logistic models using Pólya–Gamma latent variables. *Journal of the American Statistical Association* **108** 1339–1349. MR3174712
- [26] ROBERTS, G. O. and ROSENTHAL, J. S. (1998). Optimal scaling of discrete approximations to Langevin diffusions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **60** 255–268. MR1625691
- [27] ROBERTS, G. O. and SAHU, S. K. (1997). Updating schemes, correlation structure, blocking and parameterization for the Gibbs sampler. *Journal of the Royal Statistical Society Series B: Statistical Methodology* **59** 291–317. MR1440584
- [28] ROSENTHAL, J. S. (2003). Asymptotic variance and convergence rates of nearly-periodic Markov chain Monte Carlo algorithms. *Journal of the American Statistical Association* **98** 169–177. MR1965683
- [29] RUE, H. and HELD, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Chapman & Hall. MR2130347
- [30] RUE, H., MARTINO, S. and CHOPIN, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **71** 319–392. MR2649602
- [31] SEARLE, S. R., CASELLA, G. and MCCULLOCH, C. E. (2009). *Variance Components* **391**. John Wiley & Sons. MR2298115
- [32] SOKAL, A. (1997). Monte Carlo methods in statistical mechanics: foundations and new algorithms. In *Functional Integration* 131–192. Springer. MR1477456
- [33] TITSIAS, M. K. and PAPASPILIOPOULOS, O. (2018). Auxiliary gradient-based sampling algorithms. *Journal of the Royal Statistical Society Series B* **80** 749–767. MR3849342
- [34] VINES, S., GILKS, W. and WILD, P. (1996). Fitting Bayesian multiple random effects models. *Statistics and Computing* **6** 337–346.
- [35] VISHNOI, N. K. et al. (2013). $Lx = b$. *Foundations and Trends® in Theoretical Computer Science* **8** 1–141. MR3071502
- [36] VOGRINC, J., LIVINGSTONE, S. and ZANELLA, G. (2022). Optimal design of the Barker proposal and other locally-balanced Metropolis-Hastings

- algorithms. *Biometrika*. [MR4627772](#)
- [37] VOLFOVSKY, A. and HOFF, P. D. (2014). Hierarchical array priors for ANOVA decompositions of cross-classified data. *Ann. Appl. Stat.* **8** 19–47. [MR3191981](#)
- [38] WAINWRIGHT, M. J. and JORDAN, M. I. (2008). Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* **1** 1–305. <https://doi.org/10.1561/22000000001>.
- [39] WINDLE, J., POLSON, N. G. and SCOTT, J. G. (2014). Sampling Pólya-gamma random variates: alternate and approximate techniques. arXiv preprint [arXiv:1405.0506](#).
- [40] WU, K., SCHMIDLER, S. and CHEN, Y. (2022). Minimax mixing time of the Metropolis-adjusted Langevin algorithm for log-concave sampling. *The Journal of Machine Learning Research* **23** 12348–12410. [MR4577709](#)
- [41] ZANELLA, G. and ROBERTS, G. (2020). Multilevel linear models, Gibbs samplers and multigrid decompositions. *Bayesian Analysis*. [MR4381136](#)