

PENALTY PARAMETER SELECTION AND ASYMMETRY CORRECTIONS TO LAPLACE APPROXIMATIONS IN BAYESIAN P-SPLINES MODELS

Philippe Lambert, Oswaldo Gressani

LIDAM Discussion Paper ISBA
2022 / 30

ISBA

Voie du Roman Pays 20 - L1.04.01

B-1348 Louvain-la-Neuve

Email : lidam-library@uclouvain.be

<https://uclouvain.be/en/research-institutes/lidam/isba/publication.html>

Penalty parameter selection and asymmetry corrections to Laplace approximations in Bayesian P-splines models

PHILIPPE LAMBERT^{*,1,2} and OSWALDO GRESSANI³,

¹ Institut de Mathématique, Université de Liège, Belgium

² Institut de Statistique, Biostatistique et Sciences Actuarielles (ISBA),
Université catholique de Louvain, Belgium

³ Interuniversity Institute for Biostatistics and statistical
Bioinformatics (I-BioStat), Data Science Institute,
Hasselt University, Belgium

* Corresponding author: p.lambert@uliege.be

October 5, 2022

Abstract

Laplacian-P-splines (LPS) associate the P-splines smoother and the Laplace approximation in a unifying framework for fast and flexible inference under the Bayesian paradigm. Gaussian Markov field priors imposed on penalized latent variables and the Bernstein-von Mises theorem typically ensure a razor-sharp accuracy of the Laplace approximation to the posterior distribution of these variables. This accuracy can be seriously compromised for some unpenalized parameters, especially when the information synthesized by the prior and the likelihood is sparse. We propose a refined version of the LPS methodology by splitting the latent space in two subsets. The first set involves latent variables for which the joint posterior distribution is approached from a non-Gaussian perspective with an approximation scheme that is particularly well tailored to capture asymmetric patterns, while the posterior distribution for parameters in the complementary latent set undergoes a traditional treatment with Laplace approximations. As such, the dichotomization of the latent space provides the necessary structure for a separate treatment of model parameters, yielding improved estimation accuracy as compared to a setting where posterior quantities are uniformly handled with Laplace. In addition, the proposed enriched version of LPS remains entirely sampling-free, so that it operates at a computing speed that is far from reach to any existing Markov chain Monte Carlo approach. The methodology is illustrated on the additive proportional odds model with an application on ordinal survey data.

Keywords: Additive model ; P-splines ; Laplace approximation ; Skewness.

1 Motivation

By publishing his *Mémoire sur la probabilité des causes par les événements* (Laplace, 1774), the young French polymath Pierre-Simon de Laplace (1749-1827) seeded an idea today known as the Laplace approximation. At that time, Laplace probably could not have imagined that almost two centuries later, his approximation technique would be resurrected (see e.g. Leonard, 1982;

Tierney and Kadane, 1986; Rue *et al.*, 2009) to play a pivotal role in modern Bayesian literature. Essentially, the Laplace approximation is a Gaussian distribution centered around the maximum *a posteriori* (MAP) of the target distribution with a variance-covariance matrix that coincides with the inverse of the negative Hessian of the log-posterior target evaluated at the MAP. Recently, the ingenuity of Laplace’s approximation crossed path with P-splines, the brainchild of Paul Eilers and Brian Marx (Eilers and Marx, 1996) to inaugurate a new approximate Bayesian methodology labelled “Laplacian-P-splines” (LPS) with promising applications in survival analysis (Gressani and Lambert, 2018; Gressani *et al.*, 2022), generalized additive models (Gressani and Lambert, 2021), nonparametric double additive location-scale models for censored data (Lambert, 2021) and infectious disease epidemiology (Gressani *et al.*, 2021). The sampling-free inference scheme delivered by Laplace approximations combined with the possibility of smoothing different model components with P-splines in a flexible fashion paves the way for a robust and much faster alternative to existing simulation-based methods. At the same time, the LPS toolbox helps P-spline users gain access to the full potential of Bayesian methods, without having to endure the long and burdensome CPU- and real-time often required by Markov chain Monte Carlo (MCMC) samplers.

Although LPS shares some methodological aspects with the popular integrated nested Laplace approximations (INLA) approach (Rue *et al.*, 2009), there are fundamental points of divergence worth mentioning. First, the tools in INLA and its associated R-INLA software are originally built to compute approximate posteriors of univariate latent variables, contrary to LPS that natively delivers approximations to the (multivariate) joint posterior distribution of the latent vector. The key benefit of working with an approximate version of the joint posterior is that pointwise estimators and credible intervals for subsets of the latent vector (and functions thereof) can be straightforwardly constructed. Second, by working with closed-form expressions for the gradients and Hessians involved in the model, LPS is computationally more efficient than the numerical differentiation treatment proposed in INLA. Third, while INLA can be combined with various techniques for smoothing nonlinear model components, LPS is entirely devoted to P-splines smoothers with the key advantage of having a full control over the penalization scheme (as the approximate posterior distribution of the penalty parameter(s) is analytically available) and in that direction, LPS has a closer connection to the work of Wood and Fasiolo (2017), especially in the class of (generalized) additive models (Wood, 2017).

The success of Laplace approximations in Bayesian statistics owes much to a central limit type argument. Under certain regularity conditions, the Bernstein-von Mises theorem (see e.g. Van der Vaart, 2000) ensures that posterior distributions in differentiable models converge to a Gaussian distribution under large samples. In situations involving small to moderate sample sizes, the asymptotic validity of the Laplace approximation can become seriously shattered as it does not account for features involving non-zero skewness (i.e. lack of symmetry) (Ruli *et al.*, 2016). Even under relatively large samples, the Laplace approximation might fail in scenarios involving binary data as the latter are poorly informative for the model parameters and can result in a flat log-likelihood function, thus complicating inference (Ferkingstad and Rue, 2015; Gressani and Lambert, 2021).

Laplacian-P-splines originally belong to the class of latent Gaussian models, where model parameters are dichotomized between a vector of latent variables ξ (including penalized B-spline coefficients, regression coefficients and other parameters of interest) that are assigned a Gaussian prior and another vector of hyperparameters η that involves nuisance parameters, such as the smoothing parameter inherent to P-splines, and for which prior assumptions need not be Gaussian. Combining Bayes’ rule and a simplified Laplace approximation, the conditional posterior distribution of ξ under the LPS framework is approximated by a Gaussian distribution denoted by $\tilde{p}_G(\xi|\hat{\eta}, \mathcal{D})$, where $\hat{\eta}$ is a summary statistic of the posterior hyperparameter vector (e.g. the MAP or posterior mean/median) and $\mathcal{D}$ denotes the observed data. Although the latter approximation is typically accurate for penalized B-spline coefficients, it might be less appropriate for other




Figure 1 – A different approximation scheme is proposed for disjoint subsets of the latent space. As such, the asymmetric patterns for parameters that are suspected to have posterior distributions deviating from Gaussianity can be captured more accurately.

candidates in ξ with large prior variance. In that case, the misfit between the Laplace approximation and a potentially asymmetric (or heavy-tailed) target posterior distribution for a parameter can have a detrimental effect on posterior summary statistics and on any results relying on the generated approximation for the posterior distribution of the model parameters. This motivates us to develop an approach that corrects for potential posterior misfits provided by the Laplace approximation.

A recent technique proposed by [Chiuchiole et al. \(2022\)](#) in the INLA framework consists in using a skew Gaussian copula to correct for skewness when posterior latent variables have a non-negligible deviation from Gaussianity. Our proposal in models involving P-splines consists in splitting the latent parameter space into a set of parameters γ for which the posterior distribution (conditional on the hyperparameters) is approximated in a non-Gaussian fashion with an emphasis on capturing asymmetries, and a set of parameters θ for which the conditional posterior is approached with Laplace approximations. Figure 1 illustrates the separation of the latent space into two subsets, i.e. $\xi = (\gamma^\top, \theta^\top)^\top$. The posterior approximation scheme for γ takes asymmetric patterns into account, while the latent variables in θ (that typically involves penalized B-spline coefficients) are approximated by a Gaussian density. Our refined LPS approach thus allows to obtain an approximate version of the joint posterior distribution for all the components in ξ together with a posterior approximation to the hyperparameter components in η without relying on an MCMC sampling scheme.

A simple motivating example inspired from the infectious disease model of [Gressani et al. \(2021\)](#) helps framing the problem. Let $\mathcal{D} = \{y_1, \dots, y_n\}$ be an i.i.d. sample of size n from a negative binomial distribution $\text{NB}(\mu(x), \gamma)$ having a probability mass function following the parameterization of [Piegorisch \(1990\)](#) with mean $E(y|x) = \mu(x)$, variance $V(y|x) = \mu(x) + \mu(x)^2/\gamma$ and overdispersion parameter $\gamma > 0$. We model the mean with P-splines $\log(\mu(x)) = \theta^\top \mathbf{b}(x)$, where $\mathbf{b}(\cdot)$ is a cubic B-spline basis on the interval $[1, n]$ and θ is a vector of B-spline coefficients. Note that $\lim_{\gamma \rightarrow +\infty} V(y|x) = \mu(x)$, i.e. a Poisson is obtained as a limiting distribution when the




Figure 2 – (Left panel) Sample of size $n = 120$ generated with a nonlinear function $\mu(x)$ in the negative binomial model. (Right panel) Laplace approximation misfit to the posterior distribution of the log overdispersion parameter. Histogram for the posterior of $\log(\gamma)$ resulting from a long MCMC chain, Laplace approximation (solid) and skew- t approximation (dashed).

overdispersion parameter tends to a large number. Following [Lang and Brezger \(2004\)](#), a Gaussian prior is imposed on the vector of B-spline parameters $(\boldsymbol{\theta}|\lambda) \sim \mathcal{N}_{\dim(\boldsymbol{\theta})}(0, (\lambda P)^{-1})$, where P is a (full rank) penalty matrix and λ is the penalty parameter to which we assign an uninformative Gamma prior (with mean a/b and variance a/b^2), denoted by $\lambda \sim \mathcal{G}(a, b)$. To close the Bayesian model, a Gamma prior with large variance is also imposed on the overdispersion parameter γ . The left panel of Figure 2 shows a data set of size $n = 120$ simulated from the above negative binomial model with a nonlinear function for $\mu(x)$ and $\gamma = 6$. The histogram for $\log(\gamma)$ on the right panel of Figure 2 is obtained from a long MCMC chain with a Metropolis-within-Gibbs algorithm. There is a serious misfit between the Laplace approximation (solid curve) and the MCMC output, so that quantities like the posterior standard deviation or selected posterior quantiles for $\log \gamma$ will be poorly approximated by Laplace. The dashed curve represents the skewed distribution that we propose as an alternative candidate to the Laplace approximation and that will be thoroughly discussed in the next section within a Bayesian P-splines context. The fitted skewed distribution is able to capture the asymmetry that is apparent in the MCMC sample, improving the precision of posterior estimates for $\log \gamma$ as compared to Laplace and at a much lower computational cost as compared to MCMC. As such, the proposed asymmetry corrections to Laplace approximations will increase the quality of posterior estimates at virtually no additional computational cost. The article is organized as follows. Section 2 presents the Bayesian Laplacian-P-splines model and gives a detailed description of the proposed asymmetric posterior approximation methodology for non-penalized parameters. In Section 3, we illustrate the method in an additive proportional odds model for ordinal data. Finally, Section 4 concludes with a discussion.

2 Laplace approximation and Bayesian P-splines

2.1 Model specification

Consider a regression model describing the conditional distribution of a response y for given covariates $\mathbf{x}$. Denote by $\boldsymbol{\xi}$ the model parameters: it includes the regression and spline parameters, plus possibly the (log of the) scale and (unconstrained transformed) shape parameters. Denote by $p(\boldsymbol{\xi}|\boldsymbol{\eta})$ the joint prior density of $\boldsymbol{\xi}$ conditionally on a vector of hyperparameters $\boldsymbol{\eta}$. In the context of a P-spline model, the latter might include J unknown smooth functionals specified as $f_j(\cdot) = \sum_{\ell=1}^L \theta_{\ell j} b_{j\ell}(\cdot)$ ($j = 1, \dots, J$) where $\mathcal{B}_j = \{b_{j\ell}(\cdot) : \ell = 1, \dots, L\}$ denotes a B-spline basis associated to equidistant knots spanning the argument range. Vector $\boldsymbol{\eta}$ would typically include positive roughness penalty parameters λ_j with prior density $p(\lambda_j)$ for the j th functional. The

frequentist penalty on changes in differences of neighbour spline parameters (Eilers and Marx, 1996) can be translated in a Bayesian context using a conditional prior on $\boldsymbol{\theta}_j = (\theta_{j1}, \dots, \theta_{jL})^\top$ (Lang and Brezger, 2004), $p(\boldsymbol{\theta}_j|\lambda_j) \propto \exp\left(-\frac{1}{2} \boldsymbol{\theta}_j^\top (\lambda_j \mathbf{P}) \boldsymbol{\theta}_j\right)$, with $\mathbf{P} = \mathbf{D}_r^\top \mathbf{D}_r$ denoting a penalty matrix corresponding to a finite difference penalty matrix $\mathbf{D}_r$ of order r . For example, when $r = 2$, one has $\boldsymbol{\theta}_j^\top \mathbf{P} \boldsymbol{\theta}_j = \|\mathbf{D}_r \boldsymbol{\theta}_j\|_2^2 = \sum_{\ell=1}^{L-2} (\theta_{\ell+2,j} - 2\theta_{\ell+1,j} + \theta_{\ell,j})^2$. The penalty parameter $\lambda_j > 0$ is used to tune the smoothness of the associated additive term with, at the limit when $\lambda_j \rightarrow +\infty$, a polynomial of order $r - 1$ for $f_j(\cdot)$. Different prior distributions could be chosen for λ_j with Brezger and Lang (2006) suggesting to take Gamma priors $\lambda_j \sim \mathcal{G}(a_j, b_j)$ (with mean a_j/b_j and variance a_j/b_j^2). A small value for b_j ($= 10^{-4}$, say) combined with $b_j = a_j$ or $a_j = 1$ ensures a large prior variance with some more weight set on small or large values of λ_j , respectively. Mixtures of Gamma densities were also investigated in Jullion and Lambert (2007) with $(\lambda|\delta) \sim \mathcal{G}(\nu/2, \nu\delta/2)$ and $\delta \sim \mathcal{G}(a_\delta, b_\delta)$ yielding, in the special case $a_\delta = b_\delta = .5$ and $\nu = 1$, a Beta prime distribution $\mathcal{B}'(.5, .5)$ for λ or a half-Cauchy prior for $\sqrt{\lambda}$ with $p(\lambda) \propto \lambda^{-.5}(1 + \lambda)^{-1}$ (Lambert and Brehmhorst, 2019).

Penalties can also be combined and sophisticated in multiple ways, see e.g. the book by Eilers and Marx (2021) for inspiring examples. More generally, we assume that the joint conditional prior for the vector $\boldsymbol{\theta}$ stacking all the vectors of penalized B-spline coefficients in the model can be written as

$$p(\boldsymbol{\theta}|\boldsymbol{\lambda}) \propto \exp\left(-\frac{1}{2} \boldsymbol{\theta}^\top \mathcal{P}_\lambda \boldsymbol{\theta}\right), \quad (1)$$

where $\mathcal{P}_\lambda$ is a positive semi-definite matrix. The vector of model parameters can be reorganized as follows, $\boldsymbol{\xi} = (\boldsymbol{\gamma}^\top, \boldsymbol{\theta}^\top)^\top \in \mathbb{R}^{k_1+k_2}$, where $\boldsymbol{\gamma} \in \mathbb{R}^{k_1}$ denotes the vector of non-penalized parameters. If $\mathcal{D}$ generically denotes the available data and if $\boldsymbol{\lambda}$ stands for the vector of hyperparameters $\boldsymbol{\eta}$ in the specific context of P-spline models, then the joint posterior for $\boldsymbol{\xi}$ directly follows from Bayes theorem,

$$p(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D}) \propto \mathcal{L}(\boldsymbol{\xi}|\mathcal{D}) p(\boldsymbol{\gamma}) p(\boldsymbol{\theta}|\boldsymbol{\lambda}) p(\boldsymbol{\lambda}).$$

It is typically explored using Markov chain Monte Carlo methods (MCMC). In this paper, we build up on the methodology described in Gressani and Lambert (2018, 2021) and in Lambert (2021) where Laplace approximations to the conditional posterior of $(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})$ and an additional approximation to the marginal posterior of $(\boldsymbol{\lambda}|\mathcal{D})$ enable to bypass sampling algorithms, see Section 2.2.

2.2 Laplace approximation and penalty parameter selection

Assume that closed form expressions can be obtained for the gradient and Hessian of $\log p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})$,

$$\mathbf{U}_\lambda = \mathbf{U}_\lambda(\boldsymbol{\xi}) = \frac{\partial \log p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})}{\partial \boldsymbol{\xi}} \quad ; \quad \mathbf{H}_\lambda = \mathbf{H}_\lambda(\boldsymbol{\xi}) = \frac{\partial^2 \log p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})}{\partial \boldsymbol{\xi} \partial \boldsymbol{\xi}^\top}.$$

The conditional posterior mode $\hat{\boldsymbol{\xi}}_\lambda$ of $\boldsymbol{\xi}$ can be quickly obtained using the Newton-Raphson (NR) algorithm with the substitution, $\boldsymbol{\xi} \leftarrow \boldsymbol{\xi} - \mathbf{H}_\lambda^{-1} \mathbf{U}_\lambda$, repeated till convergence. The Levenberg-Marquardt algorithm (Marquardt, 1963; Commenges *et al.*, 2006) could be preferred if good initial conditions are not easily found to ensure convergence. A Laplace approximation to the conditional posterior distribution of $\boldsymbol{\xi}$ directly follows: $(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D}) \sim \mathcal{N}_{k_1+k_2}(\hat{\boldsymbol{\xi}}_\lambda, \Sigma_\lambda)$ where $\Sigma_\lambda = -\mathbf{H}_\lambda^{-1}$. Thanks to the Gaussian Markov field (GMRF) prior (Rue and Held, 2005), $p(\boldsymbol{\theta}|\boldsymbol{\lambda})$, assumed in (1) for the penalized parameters, the Normal approximation to the conditional posterior of $(\boldsymbol{\theta}|\boldsymbol{\lambda}, \mathcal{D})$ is usually excellent, see Rue *et al.* (2009) for the same argument in latent Gaussian models. However this might not be true for some non penalized parameters in $\boldsymbol{\xi}$, especially when the combined

information coming from their prior and the likelihood is sparse, see Section 2.3 for a specific handling.

The preceding Laplace approximation can be used to approximate the marginal posterior distribution of the penalty parameters $\boldsymbol{\lambda}$ with the Normal approximation substituted in the denominator of the following identity, $p_\lambda(\boldsymbol{\lambda}|\mathcal{D}) = p(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D})/p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})$, yielding

$$\begin{aligned}\tilde{p}_\lambda(\boldsymbol{\lambda}|\mathcal{D}) &= \frac{p(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D})}{\tilde{p}_G(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})} \propto p(\hat{\boldsymbol{\xi}}_\lambda, \boldsymbol{\lambda}|\mathcal{D}) \left| \hat{\Sigma}_\lambda \right|^{\frac{1}{2}} \\ &\propto \underbrace{\mathcal{L}(\hat{\boldsymbol{\xi}}_\lambda|\mathcal{D}) p(\hat{\boldsymbol{\xi}}_\lambda|\boldsymbol{\lambda}, \mathcal{D})}_{\text{Marginal likelihood}} \left| \hat{\Sigma}_\lambda \right|^{\frac{1}{2}} \times p(\boldsymbol{\lambda}),\end{aligned}\quad (2)$$

see Tierney and Kadane (1986) for the same strategy in the approximation of a marginal distribution. One might prefer to work with $\mathbf{v} = \log \boldsymbol{\lambda}$ and its approximate marginal posterior,

$$\tilde{p}_v(\mathbf{v}|\mathcal{D}) = \tilde{p}_\lambda(e^{\mathbf{v}}|\mathcal{D}) \prod_j e^{v_j}. \quad (3)$$

The maximization of (2) or of the marginal likelihood (as with ‘empirical Bayes’ methods) can be used to select a specific value for $\boldsymbol{\lambda}$. Alternatively, it could be made from the log-penalty using (3), yielding larger penalty values when the selection is made from the marginal posterior instead of the (parametrization invariant) marginal likelihood.

2.3 Asymmetric posterior for non-penalized parameters

Assume that $\boldsymbol{\xi} = (\boldsymbol{\gamma}^\top, \boldsymbol{\theta}^\top)^\top \in \mathbb{R}^{k_1+k_2}$ with $\boldsymbol{\gamma} \in \mathbb{R}^{k_1}$ suspected to have a non-symmetric marginal posterior distribution. Let $\hat{\boldsymbol{\xi}}_\lambda = (\hat{\boldsymbol{\gamma}}_\lambda^\top, \hat{\boldsymbol{\theta}}_\lambda^\top)^\top$ denote the posterior mode of $p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D})$ and $\hat{\Sigma}_\lambda^{-1}$ the observed information matrix structured in blocks as follows,

$$\hat{\Sigma}_\lambda = \begin{bmatrix} \hat{\Sigma}_\lambda^{\gamma\gamma} & \hat{\Sigma}_\lambda^{\gamma\theta} \\ \hat{\Sigma}_\lambda^{\theta\gamma} & \hat{\Sigma}_\lambda^{\theta\theta} \end{bmatrix}.$$

The conditional posterior of $(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D})$ has an approximate Normal distribution resulting from the the GMRF prior in (1) for $(\boldsymbol{\theta}|\boldsymbol{\lambda})$. One has

$$(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D}) \sim \mathcal{N}_{k_2} \left(\mathbb{E}(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D}), \hat{\Sigma}_\lambda^{\theta|\gamma} \right),$$

where

$$\begin{aligned}\mathbb{E}(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D}) &= \hat{\boldsymbol{\theta}}_\lambda + \hat{\Sigma}_\lambda^{\theta\gamma} \left(\hat{\Sigma}_\lambda^{\gamma\gamma} \right)^{-1} (\boldsymbol{\gamma} - \hat{\boldsymbol{\gamma}}_\lambda); \\ \hat{\Sigma}_\lambda^{\theta|\gamma} &= \hat{\Sigma}_\lambda^{\theta\theta} - \hat{\Sigma}_\lambda^{\theta\gamma} \left(\hat{\Sigma}_\lambda^{\gamma\gamma} \right)^{-1} \hat{\Sigma}_\lambda^{\gamma\theta}.\end{aligned}\quad (4)$$

Hence, starting from the following identity, $p_\gamma(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D}) = p(\boldsymbol{\gamma}, \boldsymbol{\theta}|\boldsymbol{\lambda}, \mathcal{D})/p(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D})$, one gets the approximation

$$p_\gamma(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D}) \approx \frac{p(\boldsymbol{\gamma}, \boldsymbol{\theta}|\boldsymbol{\lambda}, \mathcal{D})}{\tilde{p}_G(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D})} \propto p(\boldsymbol{\gamma}, \mathbb{E}(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D})|\boldsymbol{\lambda}, \mathcal{D}) \left| \hat{\Sigma}_\lambda^{\theta|\gamma} \right|^{\frac{1}{2}}, \quad (5)$$

see Eq. (2) in Tierney *et al.* (1989) for a similar expression. We propose to reparametrize $\boldsymbol{\gamma}$ by projecting it on the eigenvectors of the singular value decomposition (SVD) of $\hat{\Sigma}_\lambda^{\gamma\gamma} = \mathbf{VZV}^\top$

where $\mathbf{V} = [\mathbf{v}_1 \dots \mathbf{v}_{k_1}]$ denotes the matrix of orthonormal eigenvectors, $\boldsymbol{\zeta}$ the eigenvalues and $\mathbf{Z} = \text{diag}(\boldsymbol{\zeta})$. It yields $\tilde{\boldsymbol{\gamma}} = \mathbf{Z}^{-\frac{1}{2}} \mathbf{V}^\top (\boldsymbol{\gamma} - \hat{\boldsymbol{\gamma}}_\lambda)$ and $\boldsymbol{\gamma} = \hat{\boldsymbol{\gamma}}_\lambda + \mathbf{V} \mathbf{Z}^{\frac{1}{2}} \tilde{\boldsymbol{\gamma}}$, with

$$p_{\tilde{\boldsymbol{\gamma}}}(\tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D}) \propto p_{\boldsymbol{\gamma}}(\hat{\boldsymbol{\gamma}}_\lambda + \mathbf{V} \mathbf{Z}^{\frac{1}{2}} \tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D}). \quad (6)$$

The posterior dependence between the components of $\tilde{\boldsymbol{\gamma}}$ is expected to be milder than under the original $\boldsymbol{\gamma}$ parametrization. Therefore, conditionally on $\boldsymbol{\lambda}$, we propose to approximate the joint posterior density of $(\tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D})$ by the product of the marginal densities of its components:

$$p_{\tilde{\boldsymbol{\gamma}}}(\tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D}) \approx \prod_{s=1}^{k_1} p_{\tilde{\gamma}_s}(\tilde{\gamma}_s|\boldsymbol{\lambda}, \mathcal{D}). \quad (7)$$

Under that working independence hypothesis, each univariate marginal in the product in (7) is equal to its conditional with the other components set equal to an arbitrary value. Combined with (6), it implies that

$$\begin{aligned} p_{\tilde{\gamma}_s}(\tilde{\gamma}_s|\boldsymbol{\lambda}, \mathcal{D}) &= p_{\tilde{\gamma}_s|\tilde{\boldsymbol{\gamma}}_{-s}}(\tilde{\gamma}_s|\tilde{\boldsymbol{\gamma}}_{-s} = 0, \boldsymbol{\lambda}, \mathcal{D}) \\ &\propto p_{\tilde{\boldsymbol{\gamma}}}(\tilde{\boldsymbol{\gamma}}_s \mathbf{e}_s|\boldsymbol{\lambda}, \mathcal{D}) \\ &\propto p_{\boldsymbol{\gamma}}(\hat{\boldsymbol{\gamma}}_\lambda + \tilde{\gamma}_s \sqrt{\zeta_s} \mathbf{v}_s|\boldsymbol{\lambda}, \mathcal{D}), \end{aligned} \quad (8)$$

where $\mathbf{e}_s$ denotes the s th unit vector in $\mathbb{R}^{k_1}$ such that $[\mathbf{e}_s]_k = \delta_{ks}$. The univariate marginal posterior density in (8) can be evaluated for any value of $\tilde{\gamma}_s$ using (5). We suggest to approximate it using a skew- t (ST) or a skew-normal (SN) density. Here, we provide details for the SN density, the generalization to the ST being straightforward. By definition, $X \sim \text{SN}(\psi, \omega^2, \alpha)$ if $X \in \mathbb{R}$ and has density

$$\varphi(x|\psi, \omega^2, \alpha) = \frac{2}{\omega} \phi\left(\frac{x - \psi}{\omega}\right) \Phi\left(\alpha \frac{x - \psi}{\omega}\right),$$

with location parameter ψ , scale parameter $\omega > 0$ and slant (or skewness) parameter $\alpha \in \mathbb{R}$, see e.g. [Azzalini and Capitanio \(2014\)](#) for more details on its definition and properties. Such an approximation can be obtained by matching the first three central moments of the target variable with those of the approximating SN distribution, yielding, after matching the mean with the posterior mode of the target distribution, $(\tilde{\gamma}_s|\boldsymbol{\lambda}, \mathcal{D}) \sim \text{SN}(\tilde{\psi}_s, \tilde{\omega}_s^2, \tilde{\alpha}_s)$. Substituting these SN densities in (7) provides an approximation to the joint posterior of $(\tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D})$ and an efficient method to sample from it using the independence of its components. An analytic form for the joint posterior density of $(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D})$ can be obtained by combining (7) with $\tilde{\boldsymbol{\gamma}} = \mathbf{Z}^{-\frac{1}{2}} \mathbf{V}^\top (\boldsymbol{\gamma} - \hat{\boldsymbol{\gamma}}_\lambda)$, yielding

$$p_{\boldsymbol{\gamma}}(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D}) = \prod_{s=1}^{k_1} \frac{1}{\sqrt{\zeta_s}} \varphi(\tilde{\gamma}_s|\tilde{\psi}_s, \tilde{\omega}_s^2, \tilde{\alpha}_s).$$

An approximation to the marginal distribution of its s th component $(\gamma_s|\boldsymbol{\lambda}, \mathcal{D})$ can be obtained using fast Monte Carlo techniques. Indeed, an approximate large random sample from $(\tilde{\boldsymbol{\gamma}}|\boldsymbol{\lambda}, \mathcal{D})$ can first be generated by sampling its independent skew-normal components $\text{SN}(\tilde{\psi}_s, \tilde{\omega}_s^2, \tilde{\alpha}_s)$, yielding $\{\tilde{\boldsymbol{\gamma}}^{(m)} : m = 1, \dots, M\}$. The associated random sample for $(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D})$ is given by $\{\boldsymbol{\gamma}^{(m)} = \hat{\boldsymbol{\gamma}}_\lambda + \mathbf{V} \mathbf{Z}^{\frac{1}{2}} \tilde{\boldsymbol{\gamma}}^{(m)} : m = 1, \dots, M\}$. Then, a skew-normal approximation to the marginal posterior of γ_s can be fitted to $\{\gamma_s^{(m)} : m = 1, \dots, M\}$, yielding $(\gamma_s|\boldsymbol{\lambda}, \mathcal{D}) \sim \text{SN}(\psi_s, \omega_s^2, \alpha_s)$. Point estimates or credible regions for γ_s can be computed from it.

These different steps provide a convenient approximation to the joint posterior of the model parameters. Indeed, based on the factorization of the joint posterior density,

$$p(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D}) = p(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D}) p(\boldsymbol{\gamma}|\boldsymbol{\lambda}, \mathcal{D}) p(\boldsymbol{\lambda}|\mathcal{D}),$$

one has the following stochastic representation for $(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D})$,

$$(\boldsymbol{\xi}, \boldsymbol{\lambda}|\mathcal{D}) \sim \mathcal{N}_{k_2} \left(\mathbb{E}(\boldsymbol{\theta}|\boldsymbol{\gamma}, \boldsymbol{\lambda}, \mathcal{D}), \widehat{\Sigma}_{\boldsymbol{\lambda}}^{\boldsymbol{\theta}|\boldsymbol{\gamma}} \right) \times \prod_{s=1}^{k_1} \text{SN}(\tilde{\gamma}_s | \tilde{\psi}_s, \tilde{\omega}_s^2, \tilde{\alpha}_s) \times (\boldsymbol{\lambda}|\mathcal{D}) , \quad (9)$$

with moments in the first factor given in (4). It can be used to generate an arbitrarily large number of independent copies from the joint posterior incomparably faster than with MCMC. This is for example particularly useful to make inference on complicated functions of the model parameters or to visualize the posterior predictive distribution of the response for given covariates values.

3 Illustration

3.1 The additive proportional odds model for ordinal data

Denote by $\mathcal{E}_m = \{1, \dots, m\}$ the set containing the first m positive integers. Assume that n independent units are observed with data $\mathcal{D} = \{(y_i, \mathbf{x}_i) : i = 1, \dots, n\}$ where y is an ordinal random variable taking values in $\mathcal{E}_R$ and $\mathbf{x}$ a vector of covariates. The proportional odds (PO) model is a popular choice when the response is ordinal (Agresti, 2010). It assumes that

$$\text{logit}[P(Y \leq r|\mathbf{x})] = \eta_r = \gamma_r + \mathbf{x}^\top \boldsymbol{\theta} \quad (r \in \mathcal{E}_{R-1})$$

with a specific intercept γ_r for each cumulative logit, but a shared vector of regression parameters $\boldsymbol{\theta}$. Consequently,

$$\log \frac{\Pr(Y \leq r|\mathbf{x}_1)/\Pr(Y > r|\mathbf{x}_1)}{\Pr(Y \leq r|\mathbf{x}_2)/\Pr(Y > r|\mathbf{x}_2)} = \boldsymbol{\theta}(\mathbf{x}_1 - \mathbf{x}_2)$$

is independent of r and provides a clear interpretation to the regression parameters $\boldsymbol{\theta}$ with a change θ_k in the log-odds of Y taking values in the lower scale of the ordinal scale for every unit increase in the k th component of $\mathbf{x}$. Let $F_{ir} = P(Y_i \leq r|\mathbf{x}_i) = e^{\eta_{ir}}/(1 + e^{\eta_{ir}})$ for r in $\mathcal{E}_{R-1}$ and let $F_{i0} = 0$, $F_{iR} = 1$. Then, $\pi_{ir} = P(Y_i = r|\mathbf{x}_i) = F_{ir} - F_{i,r-1}$ for r in $\mathcal{E}_R$. Let $\boldsymbol{\xi} = (\boldsymbol{\gamma}^\top, \boldsymbol{\theta}^\top)^\top$ and assume a prior of the following form,

$$p(\boldsymbol{\xi}|\boldsymbol{\lambda}) \propto \exp \left(-\frac{1}{2} (\boldsymbol{\xi} - \mathbf{e})^\top \mathbf{K}_{\boldsymbol{\lambda}} (\boldsymbol{\xi} - \mathbf{e}) \right) , \quad (10)$$

conditionally on a vector of parameters $\boldsymbol{\lambda}$ and for a positive semi-definite matrix $\mathbf{K}_{\boldsymbol{\lambda}}$. The log-likelihood takes a simple form, $\ell(\boldsymbol{\xi}|\mathcal{D}) = \sum_{i=1}^n \log \pi_{iy_i}(\boldsymbol{\xi})$, with the resulting conditional posterior for $\boldsymbol{\xi}$,

$$\log p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D}) = \ell(\boldsymbol{\xi}|\mathcal{D}) - \frac{1}{2} (\boldsymbol{\xi} - \mathbf{e})^\top \mathbf{K}_{\boldsymbol{\lambda}} (\boldsymbol{\xi} - \mathbf{e}). \quad (11)$$

Explicit analytical forms can be obtained for the associated gradient and Hessian matrix, see Appendix A. Assume now for simplicity a model with J continuous covariates $x_1, \dots, x_J$ with smooth additive terms $f_j(x_j)$ ($j = 1, \dots, J$) describing their effects on the conditional log-odds,

$$\text{logit}[P(Y \leq r|\mathbf{x})] = \eta_r = \gamma_r + f_1(x_1) + \dots + f_J(x_J) .$$

Following Eilers and Marx (1996), consider now a basis of $(L + 1)$ cubic B-splines $\{s_{j\ell}^*(\cdot)\}_{\ell=1}^{L+1}$ associated to a generous number of equally spaced knots on the range $(x_j^{\min}, x_j^{\max})$ of values for x_j (Marx and Eilers, 1998). They are recentered for identification purposes in the additive model

using $s_{j\ell}(\cdot) = s_{j\ell}^*(\cdot) - \frac{1}{x_j^{\max} - x_j^{\min}} \int_{x_j^{\min}}^{x_j^{\max}} s_{j\ell}^*(u) du$ ($\ell = 1, \dots, L$). Then, the additive terms in the conditional model can be approximated using linear combinations of these recentered B-splines, $f_j(x_j) = \sum_{\ell=1}^L s_{j\ell}(x_{ij})\theta_{\ell j}$. In a Bayesian framework, as reminded in Section 2.1, smoothness can be forced on these additive terms by taking GMRF priors for the vectors of spline coefficients

$$p(\boldsymbol{\theta}_j | \lambda_j) \propto \exp \left(-\frac{1}{2} \boldsymbol{\theta}_j^\top (\lambda_j \mathbf{P}) \boldsymbol{\theta}_j \right),$$

where $\mathbf{P}$ stands for the penalty matrix. A multivariate Normal prior could be taken for $\boldsymbol{\gamma}$ to complete the model specification, $\boldsymbol{\gamma} \sim \mathcal{N}(\tilde{\mathbf{e}}, (\mathbf{Q})^{-1})$. Under the general formulation in (10), one has $\boldsymbol{\xi} = (\boldsymbol{\gamma}^\top, \boldsymbol{\theta}^\top)^\top$, $\boldsymbol{\theta} = (\boldsymbol{\theta}_1^\top, \dots, \boldsymbol{\theta}_J^\top)^\top$, $\mathbf{e} = (\tilde{\mathbf{e}}^\top, \mathbf{0}_{JL}^\top)^\top$, and a block-diagonal penalty matrix $\mathbf{K}_\lambda = \text{diag}(\mathbf{Q}, \mathcal{P}_\lambda)$ where $\mathcal{P}_\lambda = \boldsymbol{\Lambda} \otimes \mathbf{P}$ with $[\boldsymbol{\Lambda}]_{jj'} = \delta_{jj'} \lambda_j$. The conditional posterior for $\boldsymbol{\xi}$ is given by (11). With a Gamma prior for the penalty parameters, $\lambda_j \sim \mathcal{G}(a, b)$, one has $(\lambda_j | \boldsymbol{\xi}, \mathcal{D}) \sim \mathcal{G}(a + .5 \rho(\mathbf{P}), b + \boldsymbol{\theta}_j^\top \mathbf{P} \boldsymbol{\theta}_j)$. Starting from these conditional posterior distributions, a Metropolis-within-Gibbs algorithm can be set up to generate a random sample from the joint posterior for $(\boldsymbol{\xi}, \boldsymbol{\lambda})$, with Gibbs steps for the penalty parameters $\boldsymbol{\lambda}$ and Metropolis steps for the regression and splines parameters $\boldsymbol{\xi}$. Alternatively, proposals for $\boldsymbol{\xi}$ could be made using the modified Langevin (Roberts and Tweedie, 1996; Lambert and Eilers, 2009) or the Metropolis-Hastings algorithm with proposals based on the local topological information provided by the explicit analytic forms for the gradient and Hessian matrix (Gamerman, 1997). Such a sampling approach based on MCMC will be compared to the strategy proposed in Section 2.

3.2 Application on survey data

Consider now an illustration of the proposed methodology on data coming from the European Social Survey (ESS Round 9, 2018) with a specific focus on the French speaking respondents from Wallonia, one of the three regions in Belgium. Each of the participants (aged at least 15) was asked to react to the following statement, *Gay men and lesbians should be free to live their own life as they wish*, with a positioning on a Likert scale going from 1 (=Agree strongly) to 5 (=Disagree strongly), with 3 labelled as *Neither agree nor disagree* (with relative frequencies 1: 54.9% ; 2: 30.4% ; 3: 8.2% ; 4: 5.4% ; 5: 1.1%). That ordinal response effectively recorded on $n = 552$ respondents was analyzed using the proportional odds model described above with the number of completed years of education (14.1 ± 4.4 years) and age (47.3 ± 18.5 years) entering as additive terms with $L = 10$ recentered B-splines spanning each covariate range. Starting from Gamma priors, $\lambda_j \sim \mathcal{G}(1, 10^{-4})$ ($j = 1, 2$), the penalty parameters λ_1 and λ_2 associated to $f_1(\text{eduyrs})$ and $f_2(\text{age})$, respectively, were selected by maximizing $p(\boldsymbol{\lambda} | \mathcal{D})$ in (2) using the Levenberg-Marquardt algorithm, yielding $\hat{\lambda}_1 = 191.8$ (e.d.f.=1.24), $\hat{\lambda}_2 = 18.4$ (e.d.f.=2.55), the value in brackets standing for the effective degrees of freedom. Alternatively the maximization of the marginal likelihood in (2) would yield a very large value for $\hat{\lambda}_1$ (e.d.f.=1.00), suggesting linearity for $f_1(\text{eduyrs})$, and $\hat{\lambda}_2 = 18.5$ (e.d.f.=2.55) practically unchanged. The fitted additive terms are visible on Fig. 3 with their pointwise 95% credible intervals, suggesting a statistically non-significant effect of **eduyrs**, but a tolerant perception of homosexuality tending to decrease with **age**, with a marked change in attitude revealed beyond age 60. To compare the merits of our methodological proposal, a MCMC algorithm was run to explore $p(\boldsymbol{\xi} | \hat{\boldsymbol{\lambda}}, \mathcal{D})$, the generated samples and their properties being compared to the analytical approximations suggested in Section 2.3. The estimated additive terms and their 95% credible intervals based on MCMC are practically identical to our estimates in Fig. 3, confirming the excellent quality of the Laplace approximation to the conditional posterior distribution of $(\boldsymbol{\theta} | \boldsymbol{\lambda}, \mathcal{D})$ underlying our calculations. Let us now focus on the non-penalized regression parameter in $\boldsymbol{\gamma}$ standing for the four intercepts in the proportional odds model. The scatterplot of the MCMC sample $\{\boldsymbol{\gamma}^{(m)} : m = 1, \dots, M\}$ generated from




Figure 3 – ESS dataset: fitted additive terms for `eduyrs` and `age` with pointwise 95% credible intervals.

$p(\boldsymbol{\gamma}|\hat{\boldsymbol{\lambda}}, \mathcal{D})$ using the modified Langevin algorithm can be found in the left panel of Fig. 4 where the posterior dependence between the vector components clearly stands out. The reparametrization suggested in Section 2.3 along the principal axes corresponding to the eigenvectors of the SVD decomposition of $\hat{\Sigma}_{\boldsymbol{\lambda}}^{\boldsymbol{\gamma}}$ yields $\tilde{\boldsymbol{\gamma}}$, with the scatterplot of the associated MCMC sample $\{\tilde{\boldsymbol{\gamma}}^{(m)} = \mathbf{Z}^{-\frac{1}{2}} \mathbf{V}^{\top}(\boldsymbol{\gamma}^{(m)} - \hat{\boldsymbol{\gamma}}_{\boldsymbol{\lambda}}) : m = 1, \dots, M\}$, visible in the right panel of Fig. 4 confirming that the posterior dependence between the vector components of $\tilde{\boldsymbol{\gamma}}$ is very mild and probably negligible for most practical purposes. The suggested analytical approximation to $p_{\tilde{\gamma}_s}(\tilde{\gamma}_s|\boldsymbol{\lambda}, \mathcal{D})$ in (8) was evaluated and superposed to the MCMC sample taken as a trustful proxy of the true marginal posterior distribution of $(\tilde{\gamma}_s|\boldsymbol{\lambda}, \mathcal{D})$, see Fig. 5. The quality of the analytical approximations is excellent with a noticeable left asymmetry for $(\tilde{\gamma}_1|\boldsymbol{\lambda}, \mathcal{D})$ in particular. Transformed back in the $\boldsymbol{\gamma}$ -parametrization, one obtains the analytical forms approximated by skew-normal densities in Fig. 6 with again a very well anticipated distribution of the MCMC samples for $(\gamma_s|\boldsymbol{\lambda}, \mathcal{D})$. The positive skewness is non-negligible for the marginal posterior distribution of γ_4 : that asymmetry would not be captured by a simple Laplace approximation. It is caused by the small proportion of respondents in the survey expressing an explicit disagreement with the submitted statement on the freedom of gays and lesbians.

4 Discussion

In this paper, the Laplace P-spline approach has been extended to improve the accuracy of inference in a Bayesian framework. Indeed, when information is sparse, the posterior distribution of non-penalized parameters may exhibit a non-negligible skewness that can have adverse effects on inference or predictions when ignored. The proposed approximation to the joint posterior density in (9) takes a simple form that can be used in a much faster way than MCMC to make predictions or inference on functions of the model parameters.

An approximation to the marginal posterior distribution of the penalty parameters $\boldsymbol{\lambda}$ keeps playing an important role in the procedure. Point estimates for $\boldsymbol{\lambda}$ can be derived from it with a subsequent empirical Bayes approach (Carlin and Louis, 2000) to handle these hyperparameters, see Section 3 for an illustration. Alternatively, the uncertainty in the selection of $\boldsymbol{\lambda}$ could be accounted for by marginalizing over it with a Monte Carlo or a grid-based integration in $p(\boldsymbol{\xi}|\mathcal{D}) = \int_{\boldsymbol{\lambda}} p(\boldsymbol{\xi}|\boldsymbol{\lambda}, \mathcal{D}) p(\boldsymbol{\lambda}|\mathcal{D}) d\boldsymbol{\lambda}$. However, in the context of generalized additive models (Gresani and Lambert, 2021) and nonparametric double additive nonparametric location-scale models




Figure 4 – ESS dataset: scatterplots of the MCMC samples for $(\gamma|\lambda, \mathcal{D})$ and $(\tilde{\gamma}|\lambda, \mathcal{D})$ when $\lambda = \hat{\lambda}$.




Figure 5 – ESS dataset: approximated marginal posterior density for $(\tilde{\gamma}|\lambda, \mathcal{D})$ compared to MCMC samples when $\lambda = \hat{\lambda}$.




Figure 6 – ESS dataset: approximated marginal posterior density for $(\gamma | \lambda, \mathcal{D})$ compared to MCMC samples when $\lambda = \hat{\lambda}$.

(Lambert, 2021), simulation studies suggest that coverages of credible intervals resulting from an empirical Bayes approach for model parameter estimation are already close to their nominal values, even with moderate sample sizes.

The proposed methodology diverges from the proposal made by Rue *et al.* (2009) and underlying INLA where the size of the latent vector ξ increases with sample size and where the marginal distribution of the scalar components of ξ are the research focus. The joint distribution of $(\xi | \lambda, \mathcal{D})$ with asymmetric forms for the non-penalized components is here available for the whole vector ξ and completed by an approximation to the marginal posterior distribution of the hyperparameters λ .

Acknowledgements

Philippe Lambert acknowledges the support of the ARC project IMAL (grant 20/25-107) financed by the Wallonia-Brussels Federation and granted by the Académie Universitaire Louvain.

Appendices

A Gradient and Hessian in the PO model

Consider the proportional odds model defined in Section 3 and the notations therein. Let $v_{ir} = F_{ir}(1 - F_{ir})$, $w_{ir} = (1 + \pi_{ir} - 2F_{ir})$, $z_{ir} = (1 - 2F_{ir})v_{ir}$ for $r \in \mathcal{E}_R$ and take $s, t \in \mathcal{E}_{R-1}$. One has:

$$\begin{aligned} \frac{\partial \ell}{\partial \gamma_s} &= \sum_i \frac{\partial \log \pi_{iy_i}}{\partial \gamma_s} ; \quad \frac{\partial \log \pi_{ir}}{\partial \gamma_s} = \frac{1}{\pi_{ir}} (\delta_{rs} v_{ir} - \delta_{r-1,s} v_{i,r-1}) \\ \frac{\partial \ell}{\partial \theta_k} &= \sum_i \frac{\partial \log \pi_{iy_i}}{\partial \theta_k} ; \quad \frac{\partial \log \pi_{ir}}{\partial \theta_k} = x_{ik} w_{ir} \end{aligned}$$

and

$$\begin{aligned}\frac{\partial^2 \ell}{\partial \gamma_s \partial \gamma_t} &= \sum_i \frac{1}{\pi_{iy_i}} \{ \delta_{y_i, s, t} z_{i, y_i} - \delta_{y_i - 1, s, t} z_{i, y_i - 1} \} - \sum_i \frac{\partial \log \pi_{iy_i}}{\partial \gamma_s} \frac{\partial \log \pi_{iy_i}}{\partial \gamma_t} ; \\ \frac{\partial^2 \ell}{\partial^2 \theta_k \theta_\ell} &= \sum_i x_{ik} x_{i\ell} \left(\pi_{iy_i} w_{iy_i} - 2 \sum_{j=1}^{y_i} \pi_{ij} w_{ij} \right) ; \\ \frac{\partial^2 \ell}{\partial \theta_k \partial \gamma_s} &= - \sum_i x_{ik} (\delta_{y_i, s} v_{iy_i} + \delta_{y_i - 1, s} v_{i, y_i - 1}) .\end{aligned}$$

Therefore, given λ , one has

$$\mathbf{U}_\lambda = \frac{\partial \log p(\boldsymbol{\xi} | \lambda, \mathcal{D})}{\partial \boldsymbol{\xi}} = \frac{\partial \ell}{\partial \boldsymbol{\xi}} - \mathbf{K}_\lambda (\boldsymbol{\xi} - \mathbf{e}) \quad ; \quad \mathbf{H}_\lambda = \frac{\partial^2 \log p(\boldsymbol{\xi} | \lambda, \mathcal{D})}{\partial \boldsymbol{\xi} \partial \boldsymbol{\xi}^\top} = \frac{\partial^2 \ell}{\partial \boldsymbol{\xi} \partial \boldsymbol{\xi}^\top} - \mathbf{K}_\lambda .$$

References

- Agresti, A. (2010). *Analysis of Ordinal Categorical Data*. Hoboken, New Jersey: John Wiley & Sons, Inc.
- Azzalini, A. and Capitanio, A. (2014). *The Skew-Normal and Related Families*. New-York: Cambridge University Press.
- Brezger, A. and Lang, S. (2006). Generalized structured additive regression based on Bayesian P-splines. *Computational Statistics and Data Analysis*, **50**, 967–991.
- Carlin, B. and Louis, T. (2000). *Bayes and Empirical Bayes Methods for Data Analysis*. New-York: Chapman and Hall / CRC Press.
- Chiuchiolo, C., van Niekerk, J., and Rue, H. (2022). Joint posterior inference for latent Gaussian models with R-INLA. *Journal of Statistical Computation and Simulation*.
- Commenges, D., Jacqmin-Gadda, H., Proust, C., and Guedj, J. (2006). A Newton-like algorithm for likelihood maximization: the robust-variance scoring algorithm.
- Eilers, P. H. C. and Marx, B. D. (1996). Flexible smoothing with B-splines and penalties. *Statistical science*, **11**(2), 89–121.
- Eilers, P. H. C. and Marx, B. D. (2021). *Practical Smoothing: The Joys of P-splines*. Cambridge University Press.
- ESS Round 9 (2018). European Social Survey Round 9. Data file edition 3.1. Sikt - Norwegian Agency for shared services in education and research, Norway - Data Archive and distributor of ESS data for ESS ERIC. doi: 10.21338/NSD-ESS9-2018.
- Ferkingstad, E. and Rue, H. (2015). Improving the inla approach for approximate bayesian inference for latent gaussian models. *Electronic Journal of Statistics*, **9**(2), 2706–2731.
- Gamerman, D. (1997). Sampling from the posterior distribution in generalized linear mixed models. *Statistics and Computing*, **7**, 57–68.
- Gressani, O. and Lambert, P. (2018). Fast Bayesian inference using Laplace approximations in a flexible promotion time cure model based on P-splines. *Computational Statistics & Data Analysis*, **124**, 151–167.

- Gressani, O. and Lambert, P. (2021). Laplace approximations for fast Bayesian inference in generalized additive models based on P-splines. *Computational Statistics & Data Analysis*, **154**, 107088.
- Gressani, O., Wallinga, J., Althaus, C., Hens, N., and Faes, C. (2021). EpiLPS: a fast and flexible Bayesian tool for near real-time estimation of the time-varying reproduction number. *medRxiv*.
- Gressani, O., Faes, C., and Hens, N. (2022). Laplacian-P-splines for Bayesian inference in the mixture cure model. *Statistics in Medicine*, **41**(14), 2602–2626.
- Jullion, A. and Lambert, P. (2007). Robust specification of the roughness penalty prior distribution in spatially adaptive Bayesian P-splines models. *Computational Statistics and Data Analysis*, **51**(5), 2542–2558.
- Lambert, P. (2021). Fast Bayesian inference using Laplace approximations in nonparametric double additive location-scale models with right-and interval-censored data. *Computational Statistics & Data Analysis*, **161**, 107250.
- Lambert, P. and Bremhorst, V. (2019). Estimation and identification issues in the promotion time cure model when the same covariates influence long- and short-term survival. *Biometrical Journal*, **61**(2), 275–289.
- Lambert, P. and Eilers, P. H. C. (2009). Bayesian density estimation from grouped continuous data. *Computational Statistics and Data Analysis*, **53**, 1388–1399.
- Lang, S. and Brezger, A. (2004). Bayesian P-splines. *Journal of Computational and Graphical Statistics*, **13**(1), 183–212.
- Laplace, P. S. (1774). Mémoire sur la probabilité des causes par les événements. *Mémoires de l’Académie Royale des Sciences présentés par divers savans*, **6**, 621–656.
- Leonard, T. (1982). A simple predictive density function: Comment. *Journal of the American Statistical Association*, **77**(379), 657–658.
- Marquardt, D. W. (1963). An algorithm for least-squares estimation of nonlinear parameters. *Applied Mathematics*, **11**, 431–441.
- Marx, B. D. and Eilers, P. H. (1998). Direct generalized additive modeling with penalized likelihood. *Computational Statistics and Data Analysis*, **28**, 193–209.
- Piegorsch, W. W. (1990). Maximum likelihood estimation for the negative binomial dispersion parameter. *Biometrics*, **46**(3), 863–867.
- Roberts, G. O. and Tweedie, R. L. (1996). Exponential Convergence of Langevin Distributions and Their Discrete Approximations. *Bernoulli*, **2**, 341–363.
- Rue, H. and Held, L. (2005). *Gaussian Markov Random Fields: Theory and Applications*. Boca Raton: CRC Press.
- 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**(2), 319–392.
- Ruli, E., Sartori, N., and Ventura, L. (2016). Improved Laplace approximation for marginal likelihoods. *Electronic Journal of Statistics*, **10**(2), 3986–4009.

- Tierney, L. and Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, **81**(393), 82–86.
- Tierney, L., Kass, R. E., and Kadane, J. B. (1989). Fully exponential Laplace approximations to expectations and variances of nonpositive functions. *Journal of the American Statistical Association*, **84**(407), 710–716.
- Van der Vaart, A. W. (2000). *Asymptotic statistics*, volume 3. Cambridge University Press.
- Wood, S. N. (2017). *Generalized Additive Models: An Introduction with R (2nd ed.)*. Boca Raton: CRC Press.
- Wood, S. N. and Fasiolo, M. (2017). A generalized Fellner-Schall method for smoothing parameter optimization with application to Tweedie location, scale and shape models. *Biometrics*, **73**, 1071–1081.