

Improved Laplace approximation for marginal likelihoods

Erlis Ruli, Nicola Sartori and Laura Ventura

*Dipartimento di Scienze Statistiche
Università degli Studi di Padova
Via Cesare Battisti 241, 35134 Padova, Italy*

e-mail: ruli@stat.unipd.it; sartori@stat.unipd.it; ventura@stat.unipd.it

Abstract: Statistical applications often involve the calculation of intractable multidimensional integrals. The Laplace formula is widely used to approximate such integrals. However, in high-dimensional or small sample size problems, the shape of the integrand function may be far from that of the Gaussian density, and thus the standard Laplace approximation can be inaccurate. We propose an improved Laplace approximation that reduces the asymptotic error of the standard Laplace formula by one order of magnitude, thus leading to third-order accuracy. We also show, by means of practical examples of various complexity, that the proposed method is extremely accurate, even in high dimensions, improving over the standard Laplace formula. Such examples also demonstrate that the accuracy of the proposed method is comparable with that of other existing methods, which are computationally more demanding. An R implementation of the improved Laplace approximation is also provided through the R package `iLaplace` available on CRAN.

Keywords and phrases: Asymptotic expansions for integrals, Bayes factor, conditional minimisation, integrated likelihood, normalising constant, numerical integration.

Received January 2016.

Contents

1	Background	3987
2	The improved Laplace approximation	3989
3	Examples	3993
3.1	Gompertz distribution: Fixed d and $n \rightarrow \infty$	3993
3.2	Multivariate t /skew- t	3994
3.3	A comparison with INLA	3995
3.4	Nonlinear regression	3998
3.5	GLMM with crossed random effects	4000
3.6	GLMM with spatial random effects	4003
4	Discussion	4005
	Acknowledgments	4006
	References	4007

1. Background

Statistical applications often involve the evaluation of finite integrals of the form

$$I_n = \int_{\mathbb{R}^d} e^{-h_n(x)} dx, \quad (1.1)$$

where $h_n(x)$ is a smooth and concave real function indexed by $n > 0$, with x a d -dimensional real vector. For instance, in Bayesian analyses, $-h_n(\cdot)$ may be the log-likelihood or the log-posterior kernel and (1.1) is the Bayesian marginal likelihood or the posterior normalising constant. Furthermore, in Generalized Linear Mixed Models (GLMM) $-h_n(\cdot)$ may represent the log-likelihood plus the log-density of the random effects. In this case, (1.1) gives the marginal likelihood for the parameters (θ, θ_u) , which can be generally written as

$$\begin{aligned} L(\theta, \theta_u; y) &= \int L(\theta; u, y) f(u; \theta_u) du \\ &= \int \exp\{\log L(\theta; u, y) + \log f(u; \theta_u)\} du \\ &= \int e^{-h_n(u; \theta, \theta_u, y)} du, \end{aligned} \quad (1.2)$$

where $f(u; \theta_u)$ is the density of the random effects indexed by the parameter θ_u , and $L(\theta; u, y)$ is the likelihood for θ based on the conditional density of y given u . The quantity n is related to the information in the sample, and is often the sample size.

Integral (1.1) is frequently intractable but it can be approximated by several methods; see, e.g., [16]. Here, we focus on the Laplace approximation; see, e.g., [5, Chap. 8] and [42, Chap. 6]. Let $\hat{x} = (\hat{x}_1, \dots, \hat{x}_d)$ be the unique minimum of $h(\cdot)$, where to ease notation hereafter we drop n from I_n , $h_n(\cdot)$ and related quantities. In addition, we assume that the Hessian matrix of $h(\cdot)$ at \hat{x} , i.e.

$$\hat{V} = V(\hat{x}) = \left. \frac{\partial^2 h(x)}{\partial x \partial x^\top} \right|_{x=\hat{x}},$$

is positive definite. The Laplace approximation of (1.1) is second-order accurate, i.e., $I = \hat{I}^L \{1 + O(n^{-1})\}$, with

$$\hat{I}^L = (2\pi)^{d/2} |\hat{V}|^{-1/2} H(\hat{x}), \quad (1.3)$$

where $H(\cdot) = \exp\{-h(\cdot)\}$; see, e.g., [5, p. 335].

The Laplace approximation is widely used both in the Bayesian framework for approximating posterior densities and posterior moments [45, 36] or Bayes Factors [24], and in the frequentist framework for integrating out random effects in GLMM [7] or to compute marginal likelihoods in group models [2, 32]. In addition, it has also been used to approximate hypergeometric functions of matrix arguments [8]. Moreover, [38] propose a simulation algorithm which draws posterior samples by inverting the approximate cumulative distribution function

based on the Laplace approximation for marginal posterior densities. Lastly, [28] and [35] apply the Laplace method in the context of survival analysis and joint modelling of survival and longitudinal data, respectively.

In the standard asymptotic setting with d fixed and $n \rightarrow \infty$, the Laplace approximation (1.3) is second-order accurate. On the other hand, if also d is large, the asymptotic expansion requires more terms in order to achieve the same accuracy as in lower-dimensions. For instance, when (1.1) factorises as a product of d scalar identical integrals, the relative error is of order $O(d/n)$ [42, Sect. 6.9]. However, in practice both d and n are fixed, and when d is large relatively to n it may be necessary to improve the accuracy of the standard Laplace method. Moreover, an unappealing feature of the Laplace approximation is that it does not account for skewness or kurtosis in the integrand function. Therefore, when the shape of the integrand is far from that of the Gaussian density, which can happen especially in high-dimensional or in small sample size problems, the standard Laplace approximation can be severely inaccurate. Example 3.2 in Sect. 3 shows an example in which the Laplace approximation fails dramatically, and with inaccuracy that deteriorates with increasing dimensionality.

A possible way to improve the Laplace approximation is through the inclusion of higher-order derivatives of $h(\cdot)$ in the Taylor expansions. [27] uses this idea in a Bayesian context. [34] propose a higher-order Laplace approximation for GLMM, by considering derivatives of $h(\cdot)$ up to the sixth order. [32] use a similar approach for approximating marginal likelihoods in group models. However, when $d > 1$, the computation of higher-order derivatives can be tedious. Similar strategies are pursued by the Bayesian Bartlett correction proposed by [14], and by the corrected Laplace approximation of [41]. However, the former involves posterior expectations, which in practice must be approximated through Monte Carlo methods and the latter solution is designed for situations, such as models with crossed random effects, in which the standard Laplace approximation may not be asymptotically valid. Another improvement of the standard Laplace approximation is proposed by [30], in which (1.1) is approximated by a product of scalar blocks, after a preliminary variable transformation to achieve approximate orthogonality.

In this paper we propose an improved Laplace approximation for integrals of the form (1.1) that, unlike the standard Laplace formula, can account for skewness and non-Gaussian tails in the integrand function. Moreover, we show that the proposed method has relative approximation error of order $O(n^{-3/2})$, in a standard asymptotic setting in which the sample size n diverges and d is fixed. The core idea of the proposed method is to build an approximation of the normalised integrand through sequential and re-normalised ratios of Laplace approximations. Finally, an approximation of the target integral (1.1) is obtained indirectly by the ratio of the un-normalised integrand over the approximation of the normalised integrand, both evaluated at a specific point x . Essentially, the proposed approximation of (1.1) can be written as $I^{\text{IL}} = \hat{c}I^{\text{L}}$, where $\hat{c} > 0$ is the improvement over the standard Laplace approximation, and $I = I^{\text{IL}}\{1 + O(n^{-3/2})\}$.

Compared to the standard Laplace approximation, the proposed method requires repeated conditional minimisations and repeated evaluations of the log-integrand function and its Hessian matrix. Conditional minimisations can be computationally demanding in high dimensions. Therefore, an alternative version is introduced, which uses approximate conditional minima obtained through a first order Taylor series expansion around the global minimum. This alternative version reduces the computational time while keeping comparable accuracy with respect to the original version. Nevertheless, the most demanding task is the computation of the global minimum, which is a requirement also for the standard Laplace method.

The rest of the article is structured as follows. Section 2 introduces the improved Laplace approximation. Section 3 illustrates the method in examples in which comparison with alternative approximations are also given. Section 4 concludes with some final remarks.

2. The improved Laplace approximation

Let $p(x) = H(x)/I$ be the density function which corresponds to the kernel $H(x) = \exp\{-h(x)\}$ with normalising constant I . By the identity

$$I = \frac{H(x)}{p(x)}, \quad (2.1)$$

if $p(x)$ is known then I is readily available, for an arbitrary x . Alternatively, if a suitable estimate $\hat{p}(x)$ of $p(x)$ is available, (2.1) provides an estimate \hat{I} of I , given by $H(x)/\hat{p}(x)$. For instance, (2.1) has been used to estimate Bayesian marginal likelihoods in MCMC settings [9, 10, 21], and to approximate hidden Gaussian Markov random fields [37, 36].

While (2.1) holds for any x , it is advisable to locate such a point at a high density region [9]. One possibility is to choose $x = \hat{x}$, which may be also convenient from a computational point of view. In MCMC settings, [21] show that coordinate points other than \hat{x} may improve the approximation error. However, locating such points can be computationally intensive.

Let $x_{1:q} = (x_1, \dots, x_q)$ be the first q and $x_{q+1:d}$ the last $d - q$ components of x ($q < d$). Moreover, let \hat{x}_{x_1} be the conditional minimum of $h(\cdot)$ with x_1 fixed and let $\hat{x}_{\hat{x}_{1:q}, x_{q+1}}$ be the conditional minimum with $x_{1:q}$ fixed at $\hat{x}_{1:q}$ and x_{q+1} fixed. We require that $h(\cdot)$ satisfies the usual regularity conditions for the validity of the Laplace approximation; see, e.g., [25].

Write $p(x)$ as

$$\begin{aligned} p(x) &= p_{X_1}(x_1)p_{X_2|X_1}(x_2|x_1) \cdots p_{X_d|X_{1:d-1}}(x_d|x_{1:d-1}) \\ &= \frac{\int_{\mathbb{R}^{d-1}} H(x) dx_{2:d}}{\int_{\mathbb{R}^d} H(x) dx} \frac{\int_{\mathbb{R}^{d-2}} H(x) dx_{3:d}}{\int_{\mathbb{R}^{d-1}} H(x) dx_{2:d}} \cdots \frac{H(x)}{\int_{\mathbb{R}} H(x) dx_d}. \end{aligned} \quad (2.2)$$

An improved approximation of $p(x)$ can be obtained by approximating the integrals of each ratio on the right hand side of (2.2) through the Laplace formula.

Specifically, the Laplace approximation of the marginal density $p_{X_1}(x_1)$ is

$$\hat{p}_{X_1}(x_1) = \frac{H(x_1, \hat{x}_{x_1})}{H(\hat{x})} \left\{ \frac{|\hat{V}|}{2\pi|V_{2:d}(x_1, \hat{x}_{x_1})|} \right\}^{1/2}, \quad (2.3)$$

where $V_{2:d}(\cdot)$ is the block $(2:d, 2:d)$ of $V(\cdot)$. This result is due to [45]. For the q th conditional density in (2.2) ($2 \leq q < d$), we apply the Laplace approximation to the numerator and the denominator and obtain

$$\hat{p}_{X_q|X_{1:q-1}}(x_q|x_{1:q-1}) = \frac{\hat{p}_{X_{1:q}}(x_{1:q})}{\hat{p}_{X_{1:q-1}}(x_{1:q-1})}. \quad (2.4)$$

Finally, the conditional density $p_{X_d|X_{1:d-1}}(\cdot)$, approximated by applying the univariate Laplace method to the integral in the denominator, is

$$\hat{p}_{X_d|X_{1:d-1}}(x_d|x_{1:d-1}) = \frac{H(x)}{H(x_{1:d-1}, \hat{x}_{x_{1:d-1}})} \left\{ \frac{V_{d:d}(x_{1:d-1}, \hat{x}_{x_{1:d-1}})}{2\pi} \right\}^{1/2}, \quad (2.5)$$

where $V_{d:d}(\cdot)$ is the d th element of the diagonal of $V(\cdot)$. Using results and under the assumptions of [25] and of [45], it is possible to show that (2.3), (2.4) and (2.5) have overall relative error of order $O(n^{-1})$. Recalling that, to use identity (2.1) we only need an approximation $\hat{p}(\hat{x})$ of $p(\hat{x})$, we might be tempted to take as $\hat{p}(\hat{x})$ the product of (2.3) times (2.4) (for $2 \leq q < d$) times (2.5), all evaluated at \hat{x} . However, such a product, when replaced in (2.1), reproduces exactly (1.3), the Laplace approximation of I .

To achieve third-order accuracy we propose to re-normalise numerically (2.3), (2.4) and (2.5). Re-normalisation of (2.4) and (2.5) entails the evaluation of multi-dimensional numerical integrations. While this is true in general, in our case we only need an approximation for $p(\hat{x})$ and therefore it is still possible to re-normalise (2.4) and (2.5) by using only scalar numerical integration. The key point is to fix all the conditioning variables at the corresponding modal values prior to the re-normalisations, as explained in Scheme 1.

The product of the re-normalised versions of (2.3), (2.4) and (2.5), evaluated at \hat{x} , gives a third-order approximation of $p(\hat{x})$, as shown by the following theorem.

Theorem 2.1. *Under the assumptions of [25] for the regularity of the Laplace approximation, the improved Laplace approximation of $p(\hat{x})$ has third-order accuracy, i.e.*

$$p(\hat{x}) = \hat{p}^{iL}(\hat{x})\{1 + O(n^{-3/2})\}.$$

Proof. The first step is to show that the approximation error of (2.3) and (2.4) holds uniformly. The uniformity of (2.3) has been shown by [25, Theorem 6]. Furthermore, on the basis of Theorem 6 of [25], we can show that also (2.4) holds uniformly. This is immediate as (2.4) is the ratio of the Laplace approximation of the marginal density of $X_{1:q}$ over that of $X_{1:q-1}$ ($2 \leq q < d$), both with relative error of order $O(n^{-1})$ holding uniformly. Hence the error in (2.4) is also uniform.

<p>Step 1 Compute \hat{c}_1, the normalising constant of (2.3);</p> <p>Step 2 For each q ($2 \leq q < d$), compute \hat{c}_q, the normalising constant of (2.4) with all the conditioning variables fixed at their modal values;</p> <p>Step 3 Compute \hat{c}_d, the normalising constant of (2.5) with all the conditioning variables fixed at the corresponding modal values;</p> <p>Step 4 Set $\hat{p}^{\text{IL}}(\hat{x}) = \frac{\hat{p}(\hat{x}_1)}{\hat{c}_1} \left\{ \prod_{q=2}^{d-1} \frac{\hat{p}(\hat{x}_q \hat{x}_{1:q-1})}{\hat{c}_q} \right\} \frac{\hat{p}(\hat{x}_d \hat{x}_{1:d-1})}{\hat{c}_d}$ as the improved Laplace approximation of $p(\hat{x})$;</p> <p>Step 5 Finally, get the improved Laplace approximation $I^{\text{IL}} = H(\hat{x}) / \hat{p}^{\text{IL}}(\hat{x})$ of I.</p>

Scheme 1: Pseudo-code description of the improved Laplace method.

As shown by [45], after numerical re-normalisation, (2.3) has error of order $O(n^{-3/2})$. This is because the $O(n^{-1})$ term, when uniform, gets absorbed into the normalising constant. To complete the proof we need to show that also (2.4), after numerical re-normalisation with respect to x_q and with the conditioning variables fixed at the corresponding modal values, has error of order $O(n^{-3/2})$. To prove this, let $\hat{p}_{X_q|X_{1:q-1}}^*(x_q|\hat{x}_{1:q-1})$ be the re-normalised approximate conditional density of $X_q|X_{1:q-1}$ with the conditioning variables fixed at the modal values, i.e.

$$\hat{p}_{X_q|X_{1:q-1}}^*(x_q|\hat{x}_{1:q-1}) = \frac{c_q^{-1} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q)}{c_{q-1}^{-1} \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})},$$

where $c_q = \int_{\mathbb{R}^q} \hat{p}_{X_{1:q}}(x_{1:q}) dx_{1:q}$ and $c_{q-1} = \int_{\mathbb{R}^{q-1}} \hat{p}_{X_{1:q-1}}(x_{1:q-1}) dx_{1:q-1}$. This re-normalised approximate conditional density is third-order accurate, i.e.

$$\begin{aligned} \hat{p}_{X_q|X_{1:q-1}}(x_q|\hat{x}_{1:q-1}) &= \frac{c_q^{-1} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q) \{1 + O(n^{-3/2})\}}{c_{q-1}^{-1} \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1}) \{1 + O(n^{-3/2})\}} \\ &= \frac{c_q^{-1} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q)}{c_{q-1}^{-1} \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})} \{1 + O(n^{-3/2})\} \\ &= \hat{p}_{X_q|X_{1:q-1}}^*(x_q|\hat{x}_{1:q-1}) \{1 + O(n^{-3/2})\}. \end{aligned}$$

However, note that there is no need to compute c_{q-1} and c_q , because

$$\begin{aligned} \hat{p}_{X_q|X_{1:q}}^*(x_q|\hat{x}_{1:q-1}) &= \frac{\frac{c_q^{-1} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q)}{c_{q-1}^{-1} \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})}}{\int_{\mathbb{R}} \frac{c_q^{-1} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q)}{c_{q-1}^{-1} \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})} dx_q} \\ &= \frac{\hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q) / \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})}{\int_{\mathbb{R}} \hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q) / \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1}) dx_q} \\ &= \frac{\hat{p}_{X_{1:q}}(\hat{x}_{1:q-1}, x_q) / \hat{p}_{X_{1:q-1}}(\hat{x}_{1:q-1})}{\hat{c}_q}. \end{aligned}$$

That is, we need only to re-normalise (2.4) with the conditioning variables fixed prior to integration at their modal values, i.e. to compute \hat{c}_q . Note that, after

numerical re-normalisation with the conditioning variables fixed, approximation (2.5) becomes exact. \square

Finally, the replacement of $p(\hat{x})$ with $\hat{p}^{\text{iL}}(\hat{x})$ in (2.1) delivers the improved Laplace approximation of (1.1). Or equivalently, the improved Laplace approximation can be written as $I^{\text{iL}} = \hat{c}I^{\text{L}}$, with $\hat{c} = \prod_{i=1}^d \hat{c}_i$ and \hat{c}_i ($1 \leq i \leq d$) defined in Scheme 1.

Remark 1. The factor \hat{c} is an index of the magnitude of the improvement of the proposed method over the standard Laplace approximation. Indeed, values of \hat{c} close to 1 indicate that the improved Laplace approximation is not improving over the standard Laplace. In this case, it is likely that the integrand is Gaussian-like. On the other hand, $\hat{c} \neq 1$ indicates that the integrand may not be Gaussian-like, e.g. it may be skewed and/or heavy-tailed. Note also that if I factors as the product of d scalar integrals, then the improved Laplace approximation of I corresponds to its computation via numerical integration.

Remark 2. An important difference of the proposed method from the integrated nested Laplace approximation (INLA) of [36] is that our method is specifically designed to approximate normalising constants or marginal likelihoods for general models, and for arbitrary components of the parameter for both Bayesian and frequentist inference. For instance, the method can be used to approximate (1.2) even when random effects are not necessarily Gaussian. On the other hand, INLA is designed for approximating marginal posterior distributions and can approximate only Bayesian marginal likelihoods of Gaussian latent fields; see Eq. (30) of [36]. Some numerical comparison with INLA are provided in Section 3.3.

Remark 3. The order (x_1, \dots, x_d) is arbitrary, that is, the asymptotic error of the improved Laplace is not affected by their permutation. In practice, however, it may be useful to order x according to the cardinality of the arguments of each element of $\nabla h(x)$, the gradient of $h(\cdot)$. In particular, the element of x for which the corresponding element of $\nabla h(x)$ depends on all elements of x may be placed as the first factor in (2.2). For instance, if $d = 3$ and $\partial h(x)/\partial x_1$ depends on x_1 , $\partial h(x)/\partial x_2$ depends on $x_{1:2}$ and $\partial h(x)/\partial x_3$ depends on x , then the order (x_3, x_2, x_1) , with x_3 being the first factor in (2.2), can simplify the conditional minimisations required by the proposed method. Obviously, when each element of $\nabla h(x)$ depends on $x_{1:3}$, there is no preferred ordering. In the examples considered in Section 3 we did not experience any practical difference in the results across different permutations of x .

Remark 4. Conditional minimisations can become computationally demanding when d is large. Nevertheless, it is possible to avoid them by considering a first-order Taylor series expansion of the conditional minima as in [11]. In particular, let $x = (y, z)$, where y is the fixed block and z the remaining part of x . Then \hat{z}_y , the conditional minimum of $h(\cdot)$ for fixed y , can be approximated by the linear regression

$$\tilde{z}_y = \hat{z} + V_{zz}^{-1}V_{zy}(\hat{y} - y), \quad (2.6)$$

which is such that $\hat{z}_y - \tilde{z}_y = O(n^{-1})$. Recently, [26] applied a similar idea in a different context and noted excellent performance (see also [38]). We explore the numerical performance of the improved Laplace approximation with approximate conditional minima \tilde{z}_y in place of their exact version \hat{z}_y in Examples 3.2, 3.3, 3.5 and 3.6.

Remark 5. Linear constraints $Ax = b$, with A and b being appropriate matrices and vectors respectively, can be handled through a change of variables problem and by applying the proposed method to the remaining free components of x . Finally, when the distribution of $x_1|x_2$ is more Gaussian-like and x_2 is low-dimensional, as it happens in the INLA framework, then it may be more sensible to approximate $x_1|x_2$ by the proposed method and integrate out x_2 by numerical integration.

3. Examples

The improved Laplace approximation is implemented in the R [44] package `iLaplace` [39], available on the CRAN repository. Except for the example of Section 3.1, computations with the improved Laplace approximation, with either exact or approximate conditional minima, are performed in parallel over 11 threads through the parallel implementation provided in the package `iLaplace`. Essentially, it is an embarrassingly parallel implementation in which each integral in Steps 1–3 of Scheme 1 is computed through a separate thread.

3.1. Gompertz distribution: Fixed d and $n \rightarrow \infty$

Consider the sequence of sample sizes $\{n_i = \lceil n_{i-1} + 1.2\sqrt{n_{i-1}} \rceil, n_1 = 20, \text{ and } i = 2, \dots, 30\}$, where the symbol $\lceil \cdot \rceil$ denotes the ceiling function. Let $y^i = (y_1, \dots, y_{n_i})$ be a random sample of size n_i ($i = 2, \dots, 30$) from the Gompertz distribution, with density

$$p(y; \theta) = \alpha\beta e^{\beta y} e^{-\alpha} \exp\{-\alpha e^{\beta y}\},$$

with $\theta = (\theta_1, \theta_2) = (\log \alpha, \log \beta) \in \mathbb{R}^2$ and $y > 0$. Moreover, for each i , consider 100 random datasets of size n_i from the Gompertz distribution with $\alpha = 2$ and $\beta = 3$. For each of these datasets, we compute the normalising constant of the posterior distribution $\pi(\theta|y^i) \propto L(\theta; y^i)\pi(\theta)$, where $L(\theta; y^i)$ is the likelihood function for θ based on data y^i and $\pi(\theta) = \pi(\theta_1)\pi(\theta_2)$ is the prior distribution with $\pi(\theta_1)$ and $\pi(\theta_2)$ both being $N(0, 100)$.

The aim is to compare the behaviour of the standard (\hat{I}^L) and the improved (\hat{I}^{iL}) Laplace approximations with the target value (I) computed by adaptive numerical integration, as the sample size n diverges. Similarly to [13] and [12], let $a_1 > 0$, $a_2 > 0$, $(b_1, b_2) \in \mathbb{R}^2$ and suppose that

$$I = \hat{I}^{iL}(1 + b_1 n^{-a_1}) + o(n^{-a_1}),$$

and that

$$I = \hat{I}^L(1 + b_2 n^{-a_2}) + o(n^{-a_2}),$$

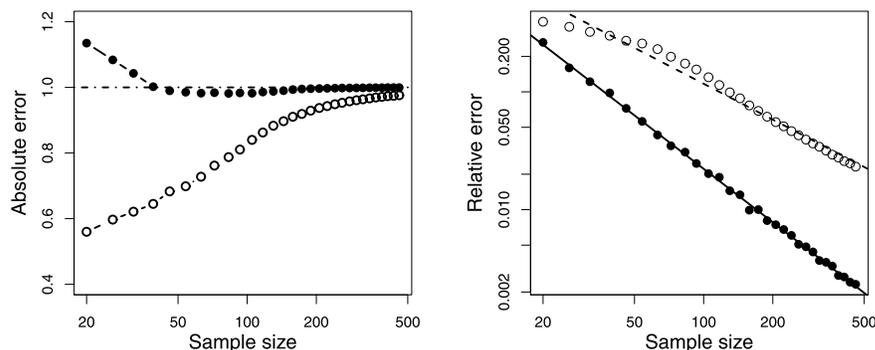


FIG 1. Numerical evaluation of the asymptotic error of the improved and standard Laplace methods. Left panel: log-log plot of (\hat{I}^{iL}/I) (\bullet) and (\hat{I}^L/I) (\circ) versus n . Right panel: log-log plot of the relative error $|\hat{I}^{iL}/I - 1|$ (\bullet) and $|\hat{I}^L/I - 1|$ (\circ) versus n ; the solid and dashed lines are the corresponding least-squares regression lines. The empirical slope is -1.51 , with 0.99 confidence interval $(-1.53, -1.48)$, for the solid line, and -1.01 , with 0.99 confidence interval $(-1.09, -0.93)$, for the dashed line.

for $n \rightarrow \infty$. Then, $\lim_{n \rightarrow \infty} \{\hat{I}^{iL}/I\} = 1$ and $\lim_{n \rightarrow \infty} \{\hat{I}^L/I\} = 1$, and if the improved Laplace approximation is more accurate than the standard Laplace, then the first limit should converge faster. Furthermore, a log-log graph of $|I/\hat{I}^{iL} - 1|$ ($|I/\hat{I}^L - 1|$) against n should be linear with slope $-a_1$ ($-a_2$) and intercept $\log |b_1|$ ($\log |b_2|$).

The left panel of Figure 1 reports the log-log plot of \hat{I}^{iL}/I and \hat{I}^L/I , both averaged across the 100 repetitions at each value of n and against the sample size n . This plot highlights that the improved Laplace approximation is more accurate than the standard Laplace method, since the visual convergence at 1 of \hat{I}^{iL}/I happens at a faster rate.

The log-log plot of the relative error averaged across the 100 repetitions at each value of the sample size n is shown on the right panel of Figure 1. This shows that the improved Laplace method achieves third-order accuracy whereas the standard Laplace formula is second-order accurate, e.g. $a_1 = 1.5$ and $a_2 = 1$.

3.2. Multivariate t /skew- t

To assess the accuracy of the proposed method even in extreme settings, we consider the multivariate t /skew- t distribution of [22], with density

$$\begin{aligned}
 p(y; \nu, a, c) &= \frac{\Gamma((\nu + d)/2)}{\Gamma((\nu + 1)/2)B(a, c)(a + c)^{1/2}2^{a+c-1}(\nu\pi)^{(d-1)/2}} \\
 &\times \frac{(1 + \nu^{-1}y_1^2)^{(\nu+1)/2} \left(1 + \frac{y_1}{(a+c+y_1^2)^{1/2}}\right)^{a+1/2}}{\{1 + \nu^{-1}(y_1^2 + \dots + y_d^2)\}^{(\nu+d)/2}} \\
 &\times \left(1 - \frac{y_1}{(a+c+y_1^2)^{1/2}}\right)^{c+1/2}. \tag{3.1}
 \end{aligned}$$

The positive parameters a and c determine the distribution of the skewed marginal, the parameter ν , i.e. the degrees of freedom (df), controls the tail behaviour of the distribution, and $B(\cdot, \cdot)$ and $\Gamma(\cdot)$ are the beta and gamma functions, respectively. This distribution is obtained from the multivariate Student's t -density centred at 0 and with identity scale matrix, where the marginal density of the first component is replaced with the univariate t /skew- t density. The case with $a = c = \nu/2$ leads to the ordinary multivariate Student's t -distribution with identity scale matrix and ν degrees of freedom.

We approximate the normalizing constant of the multivariate t /skew- t density, with the standard and the improved Laplace approximations, in two scenarios: the first with $a = c = 1.5$, and the second with $a = 12$ and $c = 0.5$. For each scenario, we consider multivariate t /skew- t densities with varying dimension and degrees of freedom. Results in the first row of Figure 2 show that the standard Laplace approximation rapidly deteriorates with increasing dimensionality, and increasing non-Gaussianity, i.e. low ν , higher skewness (large a and small c or vice versa). On the contrary, the improved Laplace approximation is reasonably accurate and stable across both increasing dimensionality and non-Gaussianity. A similar example has been considered also by [30], in order to test the accuracy of their modified Laplace approximation. However, their method is substantially less accurate than ours, only slightly improving the poor quality of the standard Laplace approximation. For instance, the normalising constant of the 10-variate t /skew- t density with $a = 4$, $c = 1$ and $\nu = 3$, is 0.013 with the standard Laplace, 0.02 with the modified Laplace method of [30] and 0.9981 with the improved Laplace approximation.

Now consider the same example but using the approximate conditional minima introduced in Remark 4 of Section 2.1. As shown in the second row of Figure 2, in this case, the results of the improved Laplace approximation with either actual or approximate conditional minima coincide.

3.3. A comparison with INLA

The following example has been considered by [17] and is known as challenging for the INLA methodology.

Let $y = (y_1, \dots, y_n)$ be conditionally independent binary values with

$$Y_i | u_i \sim \text{Bernoulli}(p_i),$$

$$\text{logit}(p_i) = \beta + u_i,$$

where $U_i \sim N(0, \sigma^2)$, for $i = 1, \dots, n$. We assume independent priors for β and $\nu = \sigma^{-2}$, with $\beta \sim N(0, 1)$ and $\nu \sim \text{Gamma}(1, 1)$ and we wish to obtain the marginal posterior distributions of β and ν .

We apply the improved and the standard Laplace methods to approximate $L(\beta, \nu; y)$, the marginal likelihood defined in (1.2) under the aforementioned modelling assumptions. Since $L(\beta, \nu; y)\pi(\beta, \nu)$ is bivariate, we use adaptive numerical integration for obtaining the marginal posteriors $\pi(\beta|y)$ and $\pi(\nu|y)$. For

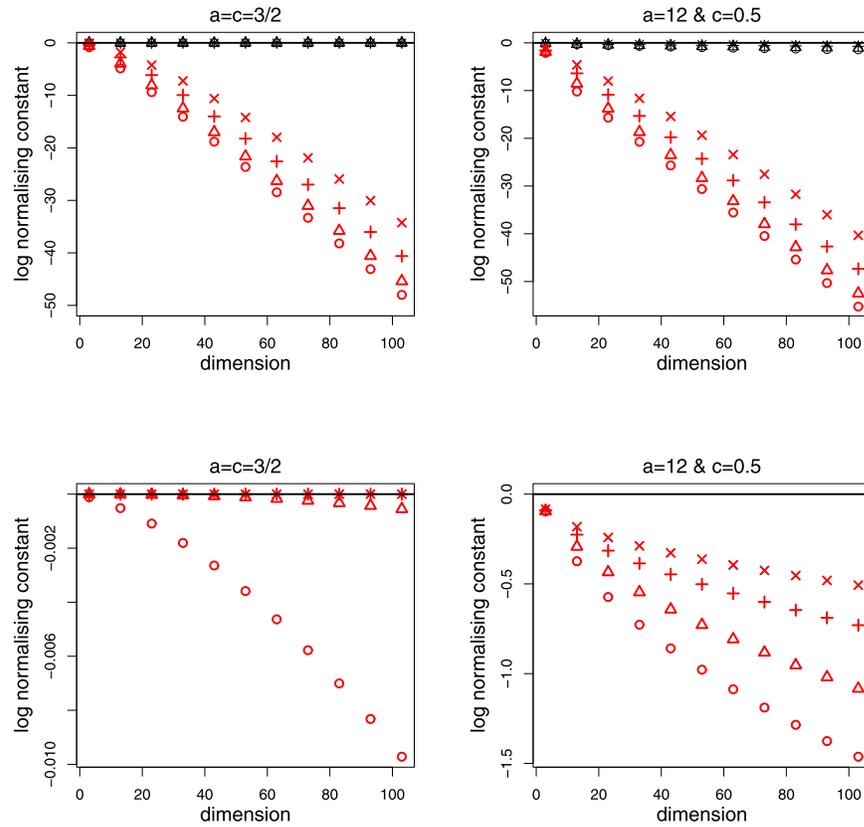


FIG 2. First row: standard (red coloured symbols) and improved Laplace (black) approximations of the normalizing constant (in logarithmic scale) of the multivariate t /skew- t density against dimension. Second row: improved Laplace approximation with either exact or approximate conditional minima overlap. Degrees of freedom are: 3 (\circ), 5 (\triangle), 10 ($+$) and 20 (\times).

comparison purposes, the marginal posteriors are also approximated by: MCMC, the standard version of INLA and the improved INLA proposed by [17]. For the MCMC approximation, we consider 10^6 final posterior samples with the JAGS software [33], after a burn-in of 10^6 samples. Computations with INLA are done using the associated R package R-INLA, using the default options.

Note that the integral over $u = (u_1, \dots, u_n)$, required for obtaining $L(\beta, \nu; y)$, can be factorised as product of n scalar integrals. Hence, in this case the improved Laplace method is as accurate in approximating $\pi(\beta, \nu|y)$ as is numerical integration.

As an illustration, consider a sample of size $n = 100$ generated from the model with $\sigma^2 = 1$ and $\beta = 2$. As a gold standard we use MCMC as implemented in the JAGS software. We compare it with the improved and the standard Laplace approximations and with INLA and the improved ILNA of [17] in the

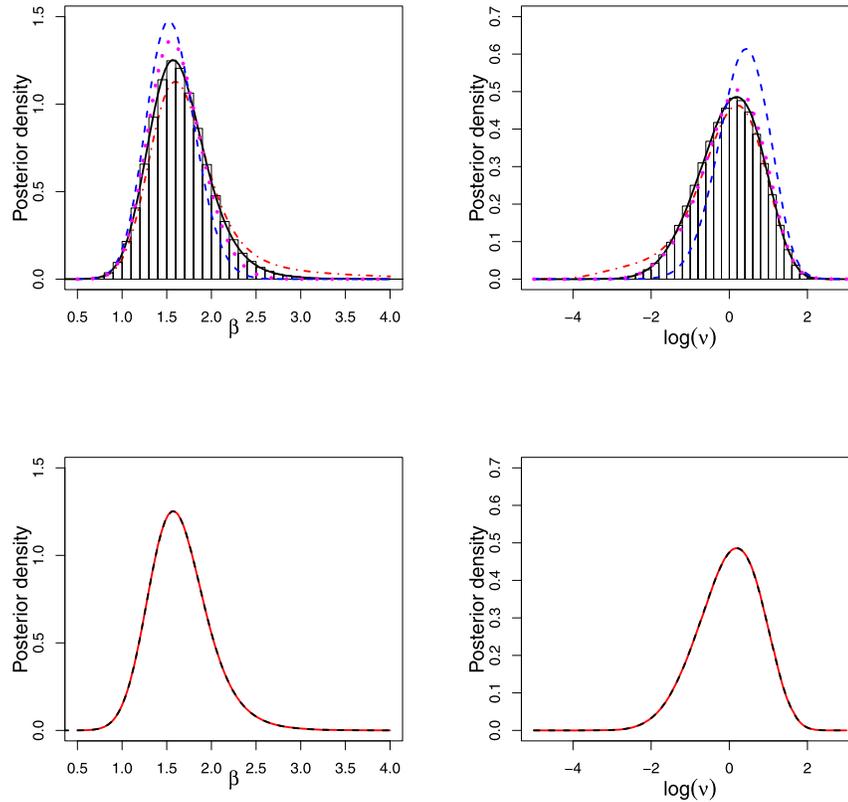


FIG 3. First row: marginal posterior distributions for β and $\log \nu$ approximated by: 10^6 MCMC samples with JAGS (histogram), Laplace (dash-dotted red curves), improved Laplace (solid black), INLA (dashed blue) and improved INLA (dotted magenta) methods. Second row: marginal posteriors of β and $\log \nu$ with the improved Laplace approximation with exact (black) and approximate conditional minima (red dashed).

first row of Figure 3. As expected, the improved Laplace approximation is virtually indistinguishable from the MCMC approximation. On the other hand, the Laplace approximation and both versions of INLA perform slightly worse than the improved Laplace approximation.

The second row of Figure 3 shows the marginal posteriors of β and $\log \nu$ approximated by the improved Laplace method using the approximate conditional minima introduced in Remark 4. In this case, the improved Laplace approximation with either exact (black continued) or approximate (red dashed) conditional minima gives indistinguishable results.

A small simulation study is performed in order to assess the accuracy of the proposed method. In particular, we consider 100 datasets with sample size $n = 100$ drawn from the model and under the same parameter values as before. For each dataset, we compute the marginal posteriors of θ by MCMC, here

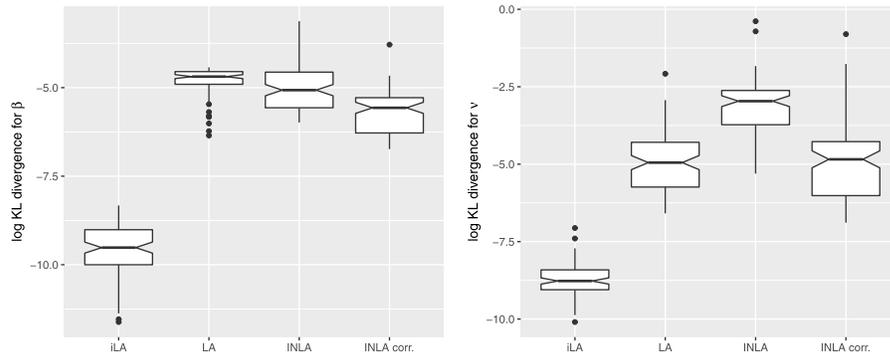


FIG 4. Log-KL divergences of marginal MCMC posteriors of β (left panel) and ν (right panel) from the corresponding ones approximated through the: Laplace (LA), improved Laplace (iLA), INLA and the corrected INLA methods in a repeated sampling context with 100 replications.

treated again as the gold standard, and by: the standard Laplace, the improved Laplace (with approximate conditional minima) and the original and corrected versions of INLA. The MCMC approximation is done by 10^6 samples, after a burn-in of 10^5 , and thinning equal to 10. As a measure of discrepancy, we compute the Kullback-Leibler (KL) divergence between the MCMC posterior and the other approximation methods. The KL divergence is defined as

$$KL(\pi; \tilde{\pi}) = \int_{\theta \in \Theta} \log \left\{ \frac{\pi(\theta|y)}{\tilde{\pi}(\theta|y)} \right\} \pi(\theta|y) d\theta,$$

where $\pi(\theta|y)$ is the MCMC posterior and $\tilde{\pi}(\theta|y)$ is the approximate posterior obtained with the other methods. For simplicity, we compute two marginal KL divergences, i.e. one for β and one for ν . The higher the KL, the worse is the approximation $\tilde{\pi}(\cdot|y)$. The MCMC marginal posteriors are computed with logspline density estimation using the `logspline` package of R. This tends to give smoother density estimates than usual kernel density estimators. The marginal posterior distributions obtained with either standard or improved Laplace approximation are available analytically, whereas those based on the two versions of INLA are build through smoothing splines.

The results in Figure 4 highlight that the marginal posteriors of β (left panel) and ν (right panel) approximated with the proposed method are the closest to the MCMC posteriors in terms of the KL divergence.

3.4. Nonlinear regression

Consider the nonlinear regression model

$$y_i = \beta_1 \exp\{1 - \exp(-x_i/\beta_2)\} + \sigma \varepsilon_i,$$

where y_i is the response variable, x_i is a covariate, (β_1, β_2) are unknown regression parameters, and the ϵ_i are independent error terms, $i = 1, \dots, n$. We focus on two possible distributions for the error term: the normal distribution and the Student's t -distribution, with unknown degrees of freedom τ . The aim is to choose among them through the Bayes Factor (BF), which in our case is given by the ratio of the posterior normalising constant of the normal model over that of the Student's t model.

As an example we consider the BOD2 dataset [3, p. 305], which concerns a study on biochemical oxygen concentration (y) as function of time (x). For both models, we assume the parameters are a priori independent. Moreover, a bivariate normal distribution with mean vector zero and scale matrix $10\mathbf{I}_2$, is assumed for β . Following the recommendations of [20], for the scale parameter σ we assume a half-Cauchy prior with scale equal to 10. Finally, the Jeffreys rule prior proposed by [18] is taken for τ . For numerical stability in the optimisations, σ and τ are considered in logarithmic scale.

Table 1 shows the log-marginal likelihoods and the Bayes factor approximated by the improved Laplace approximation and by: the standard Laplace, the Bartlett-corrected Laplace (DiCiccio et al. 14), importance sampling, the method of [10] and adaptive numerical integration as implemented in the R package `cubature`. The Chib and Jeliazkov's, IS and Bartlett-corrected Laplace approximations are replicated 500 times, where for each replication the MCMC algorithm is started at a different point. The final estimates of the log-marginal likelihood and of the BF (in decimal logarithmic scale) are obtained by averaging the 500 replications. At each replication, the Bartlett-corrected Laplace approximation is performed with 2×10^5 final MCMC posterior draws, after suitable burn-in and thinning to reduce autocorrelation. The same MCMC posterior sample is used also for computing the marginal likelihood with Chib and Jeliazkov's method. For the IS approximation we consider 10^6 draws from the multivariate Student's t -distribution with m degrees of freedom, centred at the posterior mode, with scale matrix equal to $c > 0$ times the inverse of the posterior Hessian at the modal value. Several values of c around 1 and $m \in [3, 50]$ were also considered. However, they gave very similar results, so we set $c = 1$ and $m = 3$. This choice permits to have an importance density with finite variance and with heavy tails (see, e.g., 16, Sect. 6.3). The standard deviation of the 500 log-marginal likelihoods divided by $\sqrt{500}$ is taken as a measure of Monte Carlo standard error.

Results in Table 1 indicate that the standard Laplace approximation and its Bartlett-corrected version are quite inaccurate, since both lead to substantial evidence in favor of the normal model (see [24] for the interpretation of the BF). Such an evidence is not confirmed by the BF approximated through adaptive numerical integration, here treated as the gold standard; neither IS and Chib and Jeliazkov's method confirm the aforementioned evidence. In addition, results of the improved Laplace method are in reasonable agreement with IS, Chib and Jeliazkov's approximation and adaptive numerical integration.

The inaccuracy of the standard Laplace approximation in the case of the Student's t model is most likely due to the non normality of the marginal pos-

TABLE 1

BOD2 data. Logarithm of Bayesian marginal likelihoods and BF computed with: adaptive numerical integration, standard Laplace, the improved Laplace, the Bartlett-corrected Laplace, importance sampling (IS) and Chib and Jeliazkov's method. For the last three methods, the point estimates are obtained by averaging over 500 replications; the quantity in parenthesis gives $3 \times SE_{MC}$, where SE_{MC} is the Monte Carlo standard error given by the standard deviation divided by $\sqrt{500}$.

Model	Adaptive integration	Laplace	Improved Laplace	Bartlett-corrected	Chib & Jeliazkov	IS
Normal	-2.539	-2.905	-2.540	-2.504 (0.0198)	-2.537 (0.003)	-2.539 (0.0001)
Student's t	-2.488	-5.179	-2.449	-4.188 (0.0199)	-2.478 (0.0031)	-2.457 (0.0001)
\log_{10} BF (Normal vs t)	-0.022	0.988	-0.039	0.731 (0.0103)	-0.027 (0.0057)	-0.036 (0.0183)

terior of $(\log \sigma, \log \nu)$. Indeed, a look at the bivariate kernel density estimate of this marginal bivariate posterior (not reported here) reveals that it is banana-shaped, and therefore it is quite far from being elliptical. Nevertheless, such a shape is well accommodated by the improved Laplace approximation.

3.5. GLMM with crossed random effects

We consider the problem of approximating the marginal likelihood for the fixed parameters in a model with crossed random effects [41, 40]. Such a model is useful, for instance, when analysing the Salamander mating data [29, p. 439]. These data have been analysed by [23, 40, 6, 4, 43], among others, and consist of three separate experiments, each performed according to the design given in McCullagh and Nelder [29, Table 14.3]. Each experiment involved matings among salamanders in two closed groups. Both groups contained five species R females, five species W females, five species R males and five species W males. Within each group, only 60 of the possible 100 heterosexual crosses were observed owing to time constraints. Thus, each experiment resulted in 120 binary observations indicating which matings were successful and which were not.

As in McCullagh and Nelder [29, p. 441], the data are modelled as if different sets of 20 male and 20 female salamanders were used in each experiment. Let y_{ij} be the indicator of a successful mating between female i and male j , for $i, j = 1, \dots, 60$, where only 360 of the (i, j) pairs are relevant. Let u_i^f denote the random effect that the i th female salamander has across matings in which she is involved, and define u_j^m similarly for the j th male. The data y_{ij} are assumed conditionally independent with

$$Y_{ij} | u_i^f, u_j^m \sim \text{Bernoulli}(p_{ij}),$$

$$\text{logit}(p_{ij}) = x_{ij}^\top \beta + u_i^f + u_j^m,$$

where x_{ij} is a 4-dimensional row vector of zeros and ones indicating the type of cross, β is the vector of fixed effects, $U_i^f \sim N(0, \sigma_f^2)$ and $U_j^m \sim N(0, \sigma_m^2)$.

TABLE 2

Salamander data. Comparison of the improved Laplace method (with exact and approximate conditional minima) with the standard Laplace and the modified Laplace approximation of [40]. "Sec." refers to the elapsed time in seconds and "Iter." refers to the number of iterations before convergence.

Methods	Approximate MLE						Sec.	Iter.
	β_0	β_{WS_f}	β_{WS_m}	$\beta_{WS_f \times WS_m}$	σ_f^2	σ_m^2		
Laplace:								
Exper. 1	1.34	-2.94	-0.42	3.18	1.58	0.073	0.92	22
Exper. 2	0.57	-2.46	-0.77	3.71	1.81	0.92	0.72	15
Exper. 3	1.02	-3.23	-0.82	3.82	0.35	1.85	1.05	22
Modified Laplace of [40] ^a :								
Exper. 1	1.37	-3.02	-0.44	3.27	1.72	0.185	–	–
Exper. 2	0.57	-2.53	-0.77	3.79	2.10	1.10	–	–
Exper. 3	1.04	-3.31	-0.83	3.90	0.46	2.07	–	–
Improved Laplace:								
Exper. 1	1.37	-3.02	-0.44	3.27	1.74	0.189	397	29
Exper. 2	0.56	-2.55	-0.79	3.77	2.12	1.14	209	17
Exper. 3	1.03	-3.30	-0.82	3.90	0.49	2.12	145	11
Improved Laplace with approximate conditional minima:								
Exper. 1	1.36	-2.99	-0.44	3.24	1.72	0.15	56	15
Exper. 2	0.56	-2.49	-0.75	3.72	2.07	1.05	64	16
Exper. 3	1.02	-3.27	-0.82	3.87	0.43	2.03	83	20

^aCorrected(1) values taken from [40]

As a first example we consider the estimation of β and (σ_f^2, σ_m^2) for each separate experiment – following the same model structure as in [40] – performed by maximising the approximate marginal likelihood. The aim is to compare the maximum likelihood estimate (MLE) based on the marginal likelihood approximated by the improved Laplace method with those based on the modified Laplace approximation proposed by [41] and [40]. It is well known that in models with crossed random effects the standard Laplace approximation is not asymptotically valid [41], and it may give poor results.

Let $\beta = (\beta_0, \beta_{WS_f}, \beta_{WS_m}, \beta_{WS_f \times WS_m})$ be the fixed effects, where β_0 is a constant, β_{WS_f} is the effect of the dummy variable WS_f which takes one if the observation is from a species W female and zero otherwise, and so on. The marginal likelihood has the form (1.2), with $\theta = \beta$ and $\theta_u = (\sigma_f^2, \sigma_m^2)$ and involves a 40 dimensional integral that cannot be reduced to a product of lower dimensional integrals, even though the random effects $u_f = (u_1^f, \dots, u_{20}^f)$ and $u_m = (u_1^m, \dots, u_{20}^m)$ have independent normal distributions. The approximate MLE for the three separate experiments (reported in Tab 2) are compared with those of Shun [40, Tab. 2 and Tab. 3].

From Table 2 we notice that the standard Laplace method is very fast but the resulting MLE are quite inaccurate as far as variance components parameters are concerned. On the other hand, approximate MLE obtained with the improved Laplace method, with either exact or approximate conditional min-

TABLE 3
Salamander mating data analysed jointly. Comparisons of the improved Laplace approximation (with either exact or approximate conditional minima) with the Monte Carlo Expectation-Maximisation (MC-EM) of [6], the Gibbs sampling of [23], the quasi-likelihood (PQL) approach of [7] and the standard Laplace approximation.

Methods	Approximate MLE					
	β_0	β_{WS_f}	β_{WS_m}	$\beta_{WS_f \times WS_m}$	σ_f^2	σ_m^2
Laplace	1.01	0.31	-1.90	0.99	1.17	1.04
Improved Laplace	1.02	0.32	-1.95	1.00	1.39	1.25
Improved Laplace (approximate cond. min.)	1.01	0.31	-1.92	0.98	1.34	1.19
MC-EM [6]	1.03	0.32	-1.95	0.99	1.40	1.25
Gibbs [23]	1.03	0.34	-1.98	1.07	1.50	1.36
PQL ^a	0.87	0.28	-1.69	0.95	1.35	0.93

^aValues from [6]

ima, are closer to those based on the modified Laplace approximation of [40]. However, the improved Laplace method is easier to compute since it does not require derivatives of the negative log-integrand beyond the second-order. In terms of computing time, the improved Laplace approximation with approximate conditional minima is much faster than the version with exact conditional minima, though slower than the standard Laplace approximation.

Consider now the joint analysis of the Salamander data, by independently combining the three experiments' data. In this case the marginal likelihood entails the computation of three 40-dimensional integrals. To compare our method with other results available in the literature, we consider a slightly modified version of the fixed effects. In particular, here $\beta = (\beta_{R/R}, \beta_{R/W}, \beta_{W/R}, \beta_{W/W})$, where $\beta_{R/R}$ denotes the effect of the cross between a species R female and a species R male, and so on.

The approximate MLE obtained from the improved Laplace approximation, the Monte Carlo Expectation-Maximisation (MC-EM) algorithm of [6], the quasi-likelihood approach of [7], the standard Laplace approximation and the posterior mean taken with the Gibbs sampling proposed by [23] are illustrated in Table 3. The standard Laplace approximation underestimates the variance parameters; see also [40]. The estimate of β and that of the variance parameters based on both versions of the improved Laplace approximation are quite similar to those of the MC-EM procedure of [6]; see also [43]. However, compared to MC-EM, the proposed method does not require tuning from the practitioner.

We notice that the approximate MLE based on the improved Laplace method with exact conditional minima is found within 9.5 minutes and after 14 iterations. Using approximate conditional minima, the approximate MLE is located within 3 minutes and after 15 iterations.

We can use (1.2) also for conducting full likelihood-based inference. For instance, in the case of the Salamander data analysed jointly, let us consider profile likelihood-based confidence intervals for σ_f^2 and σ_m^2 . Figure 5 depicts the aforementioned relative profile likelihoods, obtained with the standard Laplace and

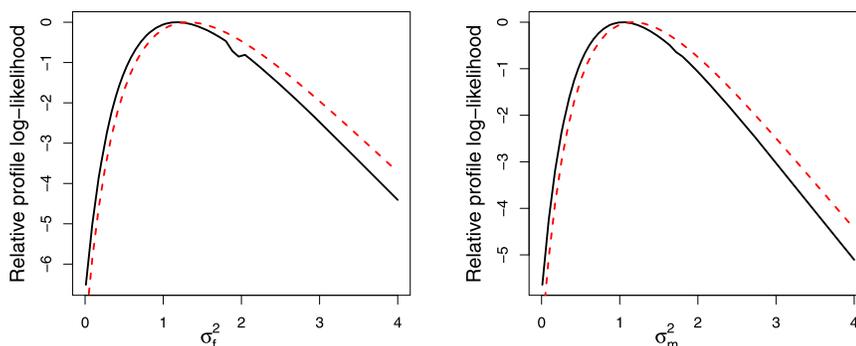


FIG 5. Salamander mating data analysed jointly. Approximate relative profile log-likelihoods for the variance components obtained with the standard Laplace approximation (black continued) and with the improved Laplace approximation with approximate conditional minima (red dashed).

with the improved Laplace with approximate conditional minima. From this plot we notice that standard Laplace-based profile likelihoods for the variance parameters are narrower than those based on the improved Laplace approximation. For instance, the 0.95 confidence interval found by inverting the profile likelihood of σ_f^2 and σ_m^2 with the standard Laplace approximation are (0.38, 2.7) and (0.31, 2.46) while those based on the improved Laplace approximation are (0.46, 2.98) and (0.38, 2.69).

3.6. GLMM with spatial random effects

We consider the application a Poisson geostatistical model to the Rongelap dataset [15]. This dataset reports counts y_i on radionuclide concentration over the length of time t_i , at the spatial location s_i , for $i = 1, \dots, 157$ different locations in Rongelap Island. On the basis of the theory of radioactive emissions, the count Y_i at the n locations can be treated approximately as realisations of independent random variables with mean $\lambda(s_i) = t_i \eta(s_i)$, where $\eta(s_i)$ measures the radioactivity at location s_i , $i = 1, \dots, n$. See [15] and reference therein for further details.

For these data, [15] propose the following geostatistical model

$$Y_i | \beta_0, \Sigma, u(s_i) \sim \text{Poisson}(t_i \eta(s_i)), \quad (3.2)$$

$$\log \eta(s_i) = \beta_0 + u(s_i),$$

where $(U(s_1), \dots, U(s_n)) \sim N_n(0, \Sigma)$ are spatial random effects, which are marginally normally distributed with mean zero and covariance matrix Σ . Typically, estimation of a full Σ is not possible, unless we place a proper prior on it, and some structure has to be imposed on it. Here we assume the exponential model, which implies that

$$\Sigma_{ij} = \text{cov}(U(s_i), U(s_j)) = \sigma^2 \exp\{-\|s_i - s_j\|_2 / \alpha\},$$

where σ^2 is a variance parameter, α controls the correlation function. In our computations the distance matrix of the locations is divided by 100 in order to avoid numerical overflow problems in the computation of Σ .

The aim is to fit model (3.2) to the Rongelap data by maximum likelihood estimation, where the parameter of interest is $(\beta_0, \sigma^2, \alpha)$. The marginal likelihood of $(\beta_0, \sigma^2, \alpha)$ can be recast in the form of (2), with $f_u(\cdot; \theta_u)$ given by the multivariate normal distribution with mean zero and the covariance matrix controlled by $\theta_u = (\sigma^2, \alpha)$ and with $L(\theta; u, y)$ the likelihood given by the conditional distribution of y given u and $\theta = \beta_0$.

To approximate the marginal likelihood we use the Laplace method and the improved Laplace approximation with approximate conditional minima. For comparison purposes, we approximate (2) also by importance sampling as proposed by [43]. Comparison with INLA in this case is not possible as the R-INLA package provides only numerical approximations to the marginal posterior distributions and not to the full marginal likelihood (2). We use the multivariate Student's t -distribution as importance density. The location of the importance density is fixed at the mode of the conditional density of u given y and $(\beta_0, \sigma^2, \alpha)$ and the scale matrix of the importance density is fixed at the Hessian matrix of the negative logarithm of the conditional density of u given y and $(\beta_0, \sigma^2, \alpha)$. The IS approximation of (2) for a fixed $(\beta_0, \sigma^2, \alpha)$ is

$$L_{IS}(\beta_0, \sigma^2, \alpha; y) = m^{-1} \sum_{j=1}^m \frac{L(\beta_0; u(s)_j, y) f_{u(s)}(u(s)_j; \Sigma)}{\tilde{f}(u(s)_j)},$$

where $u(s)_j$ is the j th random vector drawn from the importance distribution $\tilde{f}(\cdot)$, for $j = 1, \dots, m$ and m is the overall number of random draws. To have an importance density with heavy tails we fix the degrees of freedom to 5. Furthermore, we fix the random seed in order to obtain a smooth approximation for the likelihood function. An issue with the Monte Carlo approximation is that the resulting estimate is subject to stochastic variability. To take this into account we consider $m = 2 \times 10^4$ draws and compute the approximate MLE at 50 different seeds. Increasing m would give more stable results but at the cost of higher computing time. The final estimates are obtained by averaging the 50 approximate MLEs and the Monte Carlo standard error is also computed from these replications.

Results, shown in Table 4, highlight that the standard Laplace approximation tend to underestimate both σ^2 and α as compared to the IS approximation, here treated as more trustworthy. On the other hand, the improved Laplace approximation gives similar results to the IS approximation. Such a behaviour is perhaps more clear-cut if we look at the relative profile log-likelihood function of σ^2 , depicted in Figure 6 (the plots for α are similar and are omitted). This figure highlights that the standard Laplace approximation may produce misleading frequentist inference on the variance parameters, in terms of both, point estimation and profile likelihood-based interval estimation.

TABLE 4

Rongelap data. Improved Laplace approximation (with approximate conditional minima) the standard Laplace approximation and importance sampling estimates. For the IS method, the point estimates are obtained by averaging over 50 replications obtained with 50 different random seeds; the quantity in parenthesis gives $3 \times SE_{MC}$, where SE_{MC} is the Monte Carlo standard error given by the standard deviation divided by $\sqrt{50}$.

Approx. MLE	Methods		
	Laplace	Improved Laplace	Importance Sampling
β_0	1.830	1.830	1.830 (3.6×10^{-6})
σ^2	0.224	0.302	0.296 (2.2×10^{-4})
α	0.081	0.104	0.103 (4.5×10^{-5})

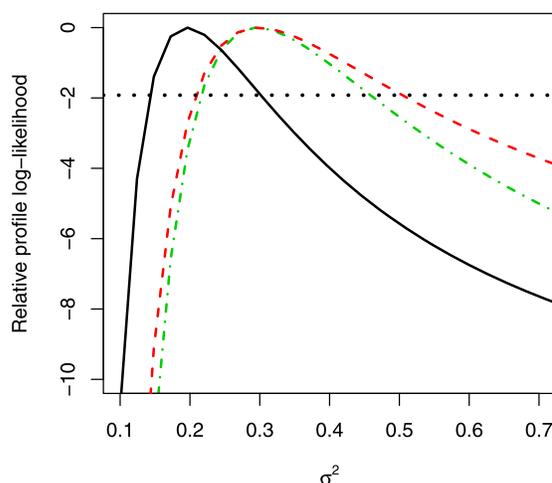


FIG 6. Rongelap data. Approximate relative profile log-likelihoods for σ^2 obtained with the standard Laplace approximation (black continued), with the improved Laplace approximation with approximate conditional minima (red dashed) and with IS (green dot-dashed). The horizontal dotted line is the level which gives the 0.95 confidence interval for σ^2 based on the log-profile likelihood ratio statistic.

4. Discussion

Although largely improving over the standard Laplace approximation, the proposed method is guaranteed to work only if the integrand is unimodal, with the mode being inside the domain of integration. This is because, if $h(\cdot)$ has either multiple minima or the minimum is not inside the domain of integration, then the determinants of blocks of its Hessian matrix may not be positive definite and the computation of (2.3) and (2.4) may break down. A possible way to deal with multimodal integrands is through a mixture of Laplace approximations, e.g. one Laplace approximation for each of mode, provided they can all be found. However, these issues are open problems and are left for future work.

A convenient feature of the proposed method is that the d integrals can be easily computed in parallel. The numerical re-normalisations are an additional

and difficult-to-quantify source of error. However, scalar numerical integration via carefully chosen adaptive quadratures is in general extremely accurate.

The main computational burden of the method is due to conditional minimisations and Hessian determinants. Both can be greatly simplified by considering analytical first and second-order derivatives of $h(\cdot)$. This is the strategy adopted in the `iLaplace` package [39] and throughout the examples. An alternative to analytical differentiation is the automatic differentiation, which provides on-line function differentiation during its evaluation; see, e.g., [19]. However, since automatic differentiation requires further programming efforts, we have not tried it in our package, though we plan to explore this possibility in future versions.

The version of the method with approximate conditional minima showed good performance and significant savings in terms of computing time in the examples considered. Another alternative to this could be to use approximate conditional minima as starting points for the computation of the exact ones. Although this would speed-up the computation of conditional minima, the method might not be as fast as when using approximate conditional minima in place of the actual one.

From a practical perspective, the improved Laplace approximation requires the integrand to be concave and unimodal but not necessarily symmetric or with Gaussian tails, though further assumptions are required to guarantee its asymptotic properties. In our experience, the standard Laplace approximation tends to work poorly when many variables of the integrand lay on the positive subset of the real numbers or when the dimensionality of the integrand increases with the sample size. Indeed, despite applying logarithmic transformations, such variables may still lead to asymmetric or heavy-tailed integrands. While in Bayesian applications $H(\cdot)$ may not always be unimodal, in GLMM it is often unimodal. In many instances, with independent random effects, the standard Laplace approximation or numerical integration with a few quadrature points are accurate enough for practical purposes. Indeed standard GLMM can now be fitted quite accurately by available R packages such as `lme4`. However, in models with complicated, dependent and/or crossed random effects, Laplace's method may perform poorly, and numerical integration may require a large number of quadrature points, hence leading to a higher computational overhead. Our method seems particularly suited for these contexts, as was also demonstrated by the examples of Sections 3.5 and 3.6.

Finally, the improved Laplace approximation can be used to compute slightly modified KL divergences that arise in the variational approximation framework [31]. In this context, the method can be useful for extending the usual Gaussian variational approach to the use of more flexible and non-conjugate densities, such as the skew- t [1]. This and the extension of the method to cases in which the mode lies outside the integration region are under investigation.

Acknowledgments

We are grateful to the Editor, the Associate Editor and to three anonymous referees for their valuable comments which lead to a substantially improved

version of our initial draft. We thank Egil Ferkingstad for providing R code used for computations with INLA in the example described in Section 3.3. This work was partially supported by grants from the University of Padova (Progetti di Ricerca di Ateneo 2015), and from MIUR, 2015, Italy.

References

- [1] AZZALINI, A. and CAPITANIO, A. (2003). Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution. *Journal of the Royal Statistical Society: Series B* **65** 367–389. [MR1983753](#)
- [2] BARNDORFF-NIELSEN, O. E. and COX, D. R. (1994). *Inference and Asymptotics*. Chapman & Hall/CRC, London. [MR1317097](#)
- [3] BATES, D. M. and WATTS, D. G. (1988). *Nonlinear Regression Analysis and Its Applications*. Wiley Online Library, New York. [MR1060528](#)
- [4] BELLIO, R. and VARIN, C. (2005). A pairwise likelihood approach to generalized linear models with crossed random effects. *Statistical Modelling* **5** 217–227. [MR2210734](#)
- [5] BLEISTEIN, N. and HANDELSMAN, R. (1986). *Asymptotic Expansions of Integrals*. Dover, New York. [MR0863284](#)
- [6] BOOTH, J. G. and HOBERT, J. P. (1999). Maximizing generalized linear mixed model likelihoods with an automated Monte Carlo EM algorithm. *Journal of the Royal Statistical Society: Series B* **61** 265–285.
- [7] BRESLOW, N. E. and CLAYTON, D. G. (1993). Approximate inference in generalized linear mixed models. *Journal of the American Statistical Association* **88** 9–25.
- [8] BUTLER, R. W. and WOOD, A. T. A. (2002). Laplace approximations for hypergeometric functions with matrix argument. *The Annals of Statistics* **30** 1155–1177. [MR1926172](#)
- [9] CHIB, S. (1995). Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* **90** 1313–1321. [MR1379473](#)
- [10] CHIB, S. and JELIAZKOV, I. (2001). Marginal likelihood from the Metropolis–Hastings output. *Journal of the American Statistical Association* **96** 270–281. [MR1952737](#)
- [11] COX, D. R. and WERMUTH, N. (1990). An approximation to maximum likelihood estimates in reduced models. *Biometrika* **77** 747–761. [MR1086687](#)
- [12] DAVISON, A. C., FRASER, D. A. S. and REID, N. (2006). Improved likelihood inference for discrete data. *Journal of the Royal Statistical Society: Series B* **68** 495–508. [MR2278337](#)
- [13] DICICCIO, T. J. and YOUNG, G. A. (2008). Conditional properties of unconditional parametric bootstrap procedures for inference in exponential families. *Biometrika* **95** 747–758. [MR2443188](#)
- [14] DICICCIO, T. J., KASS, R. E., RAFTERY, A. and WASSERMAN, L. (1997). Computing Bayes factors by combining simulation and asymptotic approximations. *Journal of the American Statistical Association* **92** 903–915. [MR1482122](#)

- [15] DIGGLE, P. J., TAWN, J. A. and MOYEED, R. A. (1998). Model-based geostatistics. *Journal of the Royal Statistical Society: Series C* **47** 299–350. [MR1626544](#)
- [16] EVANS, M. and SWARTZ, T. (2000). *Approximating Integrals via Monte Carlo and Deterministic Methods*. Oxford University Press, Oxford. [MR1859163](#)
- [17] FERKINGSTAD, E. and RUE, H. (2015). Improving the INLA approach for approximate Bayesian inference for latent Gaussian models. *Electronic Journal of Statistics* **9** 2706–2731. [MR3433587](#)
- [18] FONSECA, T. C., FERREIRA, M. A. R. and MIGON, H. S. (2008). Objective Bayesian analysis for the Student-*t* regression model. *Biometrika* **95** 325–333. [MR2521587](#)
- [19] FOURNIER, D. A., SKAUG, H. J., ANCHETA, J., IANELLI, J., MAGNUSSON, A., MAUNDER, M. N., NIELSEN, A. and SIBERT, J. (2012). AD Model Builder: using automatic differentiation for statistical inference of highly parameterized complex nonlinear models. *Optimization Methods and Software* **27** 233–249. [MR2901959](#)
- [20] GELMAN, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis* **1** 515–534. [MR2221284](#)
- [21] HSIAO, C. K., HUANG, S.-Y. and CHANG, C.-W. (2004). Bayesian marginal inference via candidate’s formula. *Statistics and Computing* **14** 59–66. [MR2019906](#)
- [22] JONES, M. C. (2002). Marginal replacement in multivariate densities, with application to skewing spherically symmetric distributions. *Journal of Multivariate Analysis* **81** 85–99. [MR1901207](#)
- [23] KARIM, M. R. and ZEGER, S. L. (1992). Generalized linear models with random effects; salamander mating revisited. *Biometrics* **48** 631.
- [24] KASS, R. E. and RAFTERY, A. E. (1995). Bayes factors. *Journal of the American Statistical Association* **90** 773–795. [MR3363402](#)
- [25] KASS, R. E., TIERNEY, L. and KADANE, J. B. (1990). The validity of posterior expansions based on Laplace’s method. In *Bayesian and Likelihood Methods in Statistics and Econometrics: Essays in Honor of George A. Barnard* (S. GEISSER, J. HODGES, S. PRESS and A. ZELLNER, eds.) 473–487. North Holland.
- [26] KHARROUBI, S. A. and SWEETING, T. J. (2016). Exponential tilting in Bayesian asymptotics. *Biometrika* **103** 337–349. [MR3509890](#)
- [27] LINDLEY, D. V. (1980). Approximate Bayesian methods. *Trabajos de Estadística Y de Investigación Operativa* **31** 223–245. [MR0638879](#)
- [28] MARTINO, S., AKERKAR, R. and RUE, H. (2011). Approximate Bayesian inference for survival models. *Scandinavian Journal of Statistics* **38** 514–528. [MR2833844](#)
- [29] MCCULLAGH, P. and NELDER, J. A. (1989). *Generalized Linear Models*, 2 ed. Chapman and Hall, London. [MR3223057](#)
- [30] NOTT, D. J., FIELDING, M. and LEONTE, D. (2009). On a generalization of the Laplace approximation. *Statistics & Probability Letters* **79** 1397–1403. [MR2750887](#)

- [31] ORMEROD, J. T. and WAND, M. P. (2010). Explaining variational approximations. *The American Statistician* **64** 140–153. [MR2757005](#)
- [32] PACE, L., SALVAN, A. and VENTURA, L. (2006). Likelihood-based discrimination between separate scale and regression models. *Journal of Statistical Planning and Inference* **136** 3539–3553. [MR2256275](#)
- [33] PLUMMER, M. (2013). JAGS version 3.4.0. <http://mcmc-jags.sourceforge.net>.
- [34] RAUDENBUSH, S. W., YANG, M.-L. and YOSEF, M. (2000). Maximum likelihood for generalized linear models with nested random effects via high-order, multivariate Laplace approximation. *Journal of Computational and Graphical Statistics* **9** 141–157. [MR1826278](#)
- [35] RIZOPOULOS, D., VERBEKE, G. and LESAFFRE, E. (2009). Fully exponential Laplace approximations for the joint modelling of survival and longitudinal data. *Journal of the Royal Statistical Society: Series B* **71** 637–654. [MR2749911](#)
- [36] 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* **71** 319–392. [MR2649602](#)
- [37] RUE, H., STEINSLAND, I. and ERLAND, S. (2004). Approximating hidden Gaussian Markov random fields. *Journal of the Royal Statistical Society: Series B* **66** 877–892. [MR2102470](#)
- [38] RULI, E., SARTORI, N. and VENTURA, L. (2014). Marginal posterior simulation via higher-order tail area approximations. *Bayesian Analysis* **9** 129–146. [MR3188302](#)
- [39] RULI, E., SARTORI, N. and VENTURA, L. (2016). iLaplace: Improved Laplace Approximation for Integrals of Unimodal Functions R package version 1.1.0.
- [40] SHUN, Z. (1997). Another look at the salamander mating data: A modified Laplace approximation approach. *Journal of the American Statistical Association* **92** 341–349.
- [41] SHUN, Z. and McCULLAGH, P. (1995). Laplace approximation of high dimensional integrals. *Journal of the Royal Statistical Society: Series B* **57** 749–760. [MR1354079](#)
- [42] SMALL, C. G. (2010). *Expansions and Asymptotics for Statistics*. Chapman & Hall/CRC, Boca Raton, Florida. [MR2681183](#)
- [43] SUNG, Y. J. and GEYER, C. J. (2007). Monte Carlo likelihood inference for missing data models. *The Annals of Statistics* **35** 990–1011. [MR2341695](#)
- [44] R CORE TEAM (2016). R: A Language and Environment for Statistical Computing R Foundation for Statistical Computing, Vienna, Austria.
- [45] TIERNEY, L. and KADANE, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association* **81** 82–86. [MR0830567](#)