
Computing the quality of the Laplace approximation

Guillaume Dehaene

Ecole Polytechnique Federale de Lausanne
guillaume.dehaene@gmail.com

Abstract

Bayesian inference requires approximation methods to become computable, but for most of them it is impossible to quantify how close the approximation is to the true posterior. In this work, we present a theorem upper-bounding the KL divergence between a log-concave target density $f(\boldsymbol{\theta})$ and its Laplace approximation $g(\boldsymbol{\theta})$. The bound we present is computable: on the classical logistic regression model, we find our bound to be almost exact as long as the dimensionality of the parameter space is high.

The approach we followed in this work can be extended to other Gaussian approximations, as we will do in an extended version of this work, to be submitted to the Annals of Statistics. It will then become a critical tool for characterizing whether, for a given problem, a given Gaussian approximation is suitable, or whether a more precise alternative method should be used instead.

Bayesian inference requires the following challenging computations: given an unnormalized density $\tilde{f}(\boldsymbol{\theta}) = \exp(-\phi_f(\boldsymbol{\theta}))$, we must compute its integral Z and then compute various expected values under the normalized density $f = \tilde{f}/Z$. One possible approximation for these computations consists in computing the Laplace approximation of \tilde{f} (Bishop [2006]). This consists of finding the maximum $\boldsymbol{\theta}^*$ of $\tilde{f}(\boldsymbol{\theta})$ and using the following Gaussian approximation:

$$\tilde{f}(\boldsymbol{\theta}) \approx \tilde{g}(\boldsymbol{\theta}) = \tilde{f}(\boldsymbol{\theta}^*) \exp\left(-\frac{1}{2}(\boldsymbol{\theta} - \boldsymbol{\theta}^*) \underbrace{H\phi_f(\boldsymbol{\theta}^*)}_{=\phi_g(\boldsymbol{\theta})} (\boldsymbol{\theta} - \boldsymbol{\theta}^*)\right) \quad (1)$$

where $H\phi_f$ is the *Hessian* matrix of the second derivatives of ϕ_f .

There exists a theorem that justifies the use of this Laplace approximation: the Bernstein-von Mises theorem (BvM; Kleijn et al. [2012]). This theorem is derived from a frequentist analysis: we treat the data as random with some fixed probability distribution. The posterior $f(\boldsymbol{\theta})$ is then a function-valued random variable which, under fairly general assumptions, becomes approximately equal to the second function-valued random variable $g(\boldsymbol{\theta})$ in the large-data limit.

This theorem mostly offers only a qualitative reassurance to the computationally minded, as the assumptions are hard to check and involve inaccessible terms. This is because current statements of BvM theorems are heavily tailored towards showing that Bayesian inference is a valid frequentist method and, namely, that it coincides with Maximum Likelihood Estimation in the large-data limit. The object of this article is to instead give a Bernstein-von Mises theorem that is aimed at characterizing how good the Laplace approximation of one given $\tilde{f}(\boldsymbol{\theta})$, with no assumptions on how $\tilde{f}(\boldsymbol{\theta})$ was generated, and while involving only quantities that are computable from $\tilde{f}(\boldsymbol{\theta})$.

1 A deterministic Bernstein-von Mises theorem

Assumptions. In order to derive our theorem, we must make some assumptions on the target density $\tilde{f}(\boldsymbol{\theta})$. In practice, we use two which play very different roles:

1. A *local* assumption constraining the higher-derivatives of $\phi_f(\boldsymbol{\theta})$.
Indeed, the Laplace approximation is computed from a second-degree Taylor expansion of $\phi_f(\boldsymbol{\theta})$. In order for this to be valid, the derivatives of $\phi_f(\boldsymbol{\theta})$ need to be small.
2. A *global* assumption, constraining the overall shape of $\tilde{f}(\boldsymbol{\theta})$ to ensure that most of its mass resides in a close neighborhood of $\boldsymbol{\theta}^*$.

The object of this assumption is to keep us safe from trivial counter-examples such as the following:

$$\tilde{f}(\boldsymbol{\theta}) = \exp\left(-\frac{\theta^2}{2}\right) + \epsilon \exp\left(-\epsilon^2 \frac{\theta^2}{2}\right) \quad (2)$$

If ϵ is small, the second term is virtually invisible and the Laplace approximation appears to be good, but it is actually terrible since the second term contributes one-half of the overall mass of \tilde{f} .

Some shape assumptions actually make stating a BvM theorem straightforward. For example, if we are willing to assume that $\tilde{f}(\boldsymbol{\theta})$ is β -strongly log-concave, i.e: $H\phi_f(\boldsymbol{\theta}) \geq \beta$, we can then apply the log-Sobolev inequality (LSI; Otto and Villani [2000] Theorem 2) which upper-bounds the Kullback-Leibler divergence between the normalized densities g and f :

$$KL(g, f) \leq \frac{1}{2} E_g \left[(\nabla\phi_f(\boldsymbol{\theta}) - \nabla\phi_g(\boldsymbol{\theta}))^T \beta^{-1} (\nabla\phi_f(\boldsymbol{\theta}) - \nabla\phi_g(\boldsymbol{\theta})) \right] \quad (3)$$

By performing a Taylor expansion of $\nabla\phi_f(\boldsymbol{\theta})$ around $\boldsymbol{\theta}^*$ and computing various moments of g , we would obtain a complicated but computable upper-bound for the distance between g and f , dominated by the third derivative of ϕ_f . However, this inequality is almost as useless as it is easy to derive because most models never lead to strongly log-concave posterior distributions.

A more realistic assumption consists in assuming that $\tilde{f}(\boldsymbol{\theta})$ is log-concave, i.e: $H\phi_f(\boldsymbol{\theta}) > 0$ (see Saumard and Wellner [2014] for a review of the properties of log-concave densities). This assumption holds for any model such that the log-prior and the log-likelihood are both concave, e.g: logistic regression with a Gaussian prior. **This is our global assumption** which ensures that $\tilde{f}(\boldsymbol{\theta})$ has a single mode, and that its tails decay at least exponentially thus guaranteeing that most of its mass is in a close neighborhood of $\boldsymbol{\theta}^*$.

Strategy. We will only sketch our proof here and refer the interested reader to our appendix for the detailed proof and the full expression of our KL bound.

The key step is the following change of variable:

$$\begin{aligned} \boldsymbol{\theta} &\rightarrow (z, \mathbf{e}) \in \mathbb{R} * S^{d-1} \\ \boldsymbol{\theta} &= \boldsymbol{\theta}^* + z^2 (H\phi_f(\boldsymbol{\theta}^*))^{-1/2} \mathbf{e} \end{aligned}$$

where S^{d-1} is the d -dimensional unit sphere. This change of variable is such that the conditional density of the random variable $z|\mathbf{e}$ is strongly log-concave, thus enabling us to apply the LSI. We are then able to control the KL divergence between the random variable $z|\mathbf{e}$ under the approximate model g and the true model f . Denoting $\psi_{f,\mathbf{e}}(z)$ the negative log-density of $z|\mathbf{e}$ under f , we have:

$$KL(z_g|\mathbf{e}, z_f|\mathbf{e}) \leq \frac{E_g \left[\left(\psi'_{f,\mathbf{e}}(z) - 2z^3 \right)^2 \right]}{\min_z \left[\psi''_{f,\mathbf{e}}(z) \right]} \quad (4)$$

From this, we can approximate the density under f of the random variable \mathbf{e} , which requires an intractable marginalization of the variable z . Since we have an upper-bound of the KL divergence, we can use the Evidence Lower-Bound (Murphy [2012] chapter 21) to approximate the marginalization and have an upper-bound on the error:

$$\log[f(\mathbf{e})] \approx C + E_g \left[\frac{z^4}{2} - \psi_{f,\mathbf{e}}(z) \right] \quad (5)$$

We can finally turn this approximation of the density of e under model into an approximation of the KL divergence between g and f over this variable, as the variance of the ELBO as we sample e from the density g :

$$KL(e_g, e_f) \approx \text{var} \left[e_g \rightarrow \left(E_g \left[\frac{z^4}{2} - \psi_{f, e_g}(z) \right] \right) \right] \quad (6)$$

The total KL divergence between g and f is then found as the sum of the KL divergence caused by the e random-variable and the mean KL divergence caused by the conditional $z|e$:

$$KL(g, f) = KL(e_g, e_f) + E [KL(z_g|e_g, z_f|e_g)] \quad (7)$$

Measuring the derivatives. In our theorem, two key quantities control the size of the KL divergence between approximation and truth:

1. The KL divergence caused by the conditionals $KL(z_g|e_g, z_f|e_g)$ is controlled by a LSI.
2. The KL divergence caused by e is controlled by approximating the density $f(e)$ which is done through an Evidence Lower-Bound approximation.

In both cases, we need to deal with expected values of differences of the log-densities: $\phi_f - \phi_g$, where ϕ_g is a Taylor expansion of ϕ_f to second-order. Thus, if we measure the strength of the higher-derivatives of ϕ_f , we can use it to deduce the size of the difference: $\phi_f - \phi_g$, and to control the KL divergence. Similarly, by controlling the higher-derivatives, we are able to bound the minimum curvature of $z|e$ under density f : $\min_z [\psi''_{f, e}(z)]$.

Furthermore, notice that the derivatives of ϕ_f only matter once we fix a direction e . Thus, we only need to measure the size of the derivatives of $\phi_f(\theta)$ along lines which go through θ^* . In this article, we will consider only derivatives up to fourth order, but our result could be extended to involve higher-derivatives to any arbitrary order. We define the following quantities:

$$\Delta_3(e) = \frac{\partial^3}{\partial r^3} \left[\phi_f \left(\theta^* + r(H\phi_f(\theta^*))^{-1/2} e \right) \right]_{r=0} \quad (8a)$$

$$\Delta_4(e) = \max_{r \geq 0} \left| \frac{\partial^4}{\partial r^4} \left[\phi_f \left(\theta^* + r(H\phi_f(\theta^*))^{-1/2} e \right) \right] \right| \quad (8b)$$

Second, notice that these derivatives impact the final result through an average over the random variable e_g , whether as the size of the typical oscillation of the ELBO approximation of $\log[f(e)]$, or as the size of the mean of $KL(z_g|e_g, z_f|e_g)$. Thus, instead of having to compute the maximum of $\Delta_3(e)$ (which we could only do in non-polynomial time; Hillar and Lim [2013]), we can simply sample from e_g and compute an empirical mean of $\Delta_3(e)$ instead, which is computationally much cheaper.

We are now finally ready to state our theorem.

Theorem 1. *A computable Bernstein-von Mises theorem.*

If the higher log-derivatives of f are small, as measured by eqs. 8a,8b, g is a good approximation of f :

- *It gives a good approximation of the conditional density of $z|e$:*

$$KL(z_g|e, z_f|e) \leq \frac{E_g \left[\left(\psi'_{f, e}(z) - 2z^3 \right)^2 \right]}{\min_z \left[\psi''_{f, e}(z) \right]} \approx \frac{2[\Delta_3(e)]^2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+3)/2]}{\Gamma[d/2]} \quad (9)$$

- *It gives a good approximation of the marginal density of e :*

$$\log[f(e)] \approx C + E_g \left[\frac{z^4}{2} - \psi_{f, e}(z) \right] \approx C + \frac{-(\sqrt{2})\Delta_3(e)}{3} \frac{\Gamma[(d+1)/2]}{\Gamma[d/2]} \quad (10)$$

$$KL(e_g, e_f) \approx \frac{1}{2} \text{var} [e_g \rightarrow \log[f(e_g)]] \quad (11)$$

$$\approx \frac{1}{9} \left[\frac{\Gamma[(d+1)/2]}{\Gamma[d/2]} \right]^2 E [\Delta_3(e_g)^2] \quad (12)$$

- And thus gives a good global approximation:

$$KL(g, f) \lesssim E \left[\Delta_3(e_g)^2 \right] \left(\frac{2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]} + \frac{1}{9} \left[\frac{\Gamma[(d+3)/2]}{\Gamma[d/2]} \right]^2 \right) \quad (13)$$

Please refer to our appendix for more detailed statements of this theorem.

Critically, note that the approximate expressions correspond to a rough first order approximation of the full bound which might not always be appropriate.

Computability. Critically, our upper-bound for the KL divergence is computable. In its approximate version, the only term that depends on the target density $\tilde{f}(\boldsymbol{\theta})$ is the term $E \left[\Delta_3(e_g)^2 \right]$. For each value of e , $\Delta_3(e)$ can be computed from the third-derivative tensor $\phi_f^{(3)}(\boldsymbol{\theta}^*)$. We are then left with approximating the expected value, which we can do by sampling from e_g . We could also find an explicit formula for this expected value, though this approach wouldn't work with the more exact forms of the bound (please refer to the appendix for the detailed expression of the bound).

Example: logistic regression. In order to demonstrate the applicability of our bound, we now show how it could be applied for linear logistic regression. If the data \mathcal{D} is composed of the n pairs: $(y_i, \mathbf{x}_i) \in \{-1, 1\} * \mathbb{R}^d$ and the prior is Gaussian (with mean 0 and covariance matrix $\sigma_0^2 I_d$), the posterior is:

$$p(\boldsymbol{\theta}|\mathcal{D}) = \frac{1}{Z} \exp \left(-\frac{1}{2} \sum_{j=1}^d \frac{(\theta_j)^2}{\sigma_0^2} \right) \prod_{i=1}^n \frac{1}{1 + \exp[-y_i(\boldsymbol{\theta} \cdot \mathbf{x}_i)]} \quad (14)$$

which is log-concave, with very regular higher-derivatives. We compared the approximate bound based on the third derivative (eq. 13) and the true KL divergence (computed by sampling methods; see appendix).

In the following table, we report the value of the KL divergence, of the approximate bound, and of their ratio in five typical examples. As can be observed, the bound is fairly tight: the ratio between the real KL divergence and the bound is consistently above 0.4.

d	n	σ_0	True KL	Approximate bound	Bound efficiency
5	20	10	0.31	0.82	0.38
5	100	10	0.057	0.11	0.5
5	1000	10	0.0065	0.011	0.55
50	100	10	528	1288	0.41
50	1000	10	0.38	0.46	0.83

2 Conclusion

We have presented an upper-bound for the KL divergence between the Laplace approximation and the true posterior distribution. We have applied successfully this bound to the classical logistic regression model and we have observed that it is tight even if we only use the terms which depend on the third derivative.

The theorem we have presented here represents only a small fraction of a longer article which we will submit to the Annals of Statistics. In this new article, we will present a more general version of the bound which can be applied to any Gaussian approximation. We will discuss why and when this bound is tight. We will also prove, through a frequentist analysis of the posterior, why and when the approximation with only the third derivative is valid. Finally, we will discuss how to extend this line of work to non log-concave approximations through other shape assumptions.

This work represents a major step forward for Gaussian approximation methods in Bayesian inference as it will enable the statistician to characterize, on a given problem, how good his approximation is through one cheap extra calculation. Should the approximation prove to be poor, he will then choose a more precise but more expensive approximation method.

References

- Christopher M Bishop. *Pattern recognition and machine learning*. springer, 2006.
- Christopher J Hillar and Lek-Heng Lim. Most tensor problems are np-hard. *Journal of the ACM (JACM)*, 60(6):45, 2013.
- BJK Kleijn, AW Van der Vaart, et al. The bernstein-von-mises theorem under misspecification. *Electronic Journal of Statistics*, 6:354–381, 2012.
- Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT press, 2012.
- Felix Otto and Cédric Villani. Generalization of an inequality by talagrand and links with the logarithmic sobolev inequality. *Journal of Functional Analysis*, 173(2):361–400, 2000.
- Adrien Saumard and Jon A Wellner. Log-concavity and strong log-concavity: a review. *Statistics surveys*, 8:45, 2014.

Appendix: proofs and detailed statement of the theorem

In this appendix, we give detailed versions of the theorem that we have stated in our main article, as well as a detailed proof of this theorem and of various important intermediate lemmas (section A). We also describe precisely how our example with logistic regression was handled (section B).

A detailed statement of our theorem can be found as Theorem 9 on page 16.

A Upper-bounding the KL divergence

This section focuses on proving our theorem that upper bounds the KL divergence between the Laplace approximation $g(\boldsymbol{\theta})$ and the target density $f(\boldsymbol{\theta})$.

Throughout this section, we will denote with $\boldsymbol{\theta}_g$ a Gaussian random variable with density $g(\boldsymbol{\theta})$ and with $\boldsymbol{\theta}_f$ a random variable with density $f(\boldsymbol{\theta})$. When we perform our change of variable, we will use the subscript to indicate which density is used.

We first detail our change of variable: $\boldsymbol{\theta} \rightarrow (r, \mathbf{e}) \rightarrow (z, \mathbf{e})$, and next how it ensures that $\tilde{f}(z|\mathbf{e})$ is strongly log-concave. We then show how the LSI can be used to upper-bound the KL divergence between z_g and $z_f|\mathbf{e}$ for any value of \mathbf{e} . We next bound the divergence between \mathbf{e}_g and \mathbf{e}_f . Finally, we state our theorem that upper-bounds the KL divergence between g and f .

A.1 Two changes of variable

Our objective is to prove that the two random variables $\boldsymbol{\theta}_g$ and $\boldsymbol{\theta}_f$ have almost the same distribution, in that the KL divergence $KL(g, f)$ is close to 0.

In order to prove that, we will use a change of variable that considerably simplifies our work. Let $\boldsymbol{\theta}^*$ be the minimum of $\phi_f(\boldsymbol{\theta}) = -\log(\tilde{f}(\boldsymbol{\theta}))$, and $\Sigma = [H\phi_f(\boldsymbol{\theta}^*)]^{-1}$. Since $g(\boldsymbol{\theta})$ is the Laplace approximation of $f(\boldsymbol{\theta})$, $\boldsymbol{\theta}^*$ is the mean of $\boldsymbol{\theta}_g$ and Σ its covariance. We can thus perform a first change of variable to “standardize” $\boldsymbol{\theta}_g$, i.e: re-express it as a translated and scaled version of the standard Gaussian variable $\boldsymbol{\eta}$, with mean 0 and covariance the identity matrix:

$$\begin{aligned}\boldsymbol{\theta}_g &= \boldsymbol{\theta}^* + [H\phi_f(\boldsymbol{\theta}^*)]^{-1/2} \boldsymbol{\eta} \\ &= \boldsymbol{\theta}^* + \Sigma^{1/2} \boldsymbol{\eta}\end{aligned}$$

We then perform two further changes of variable on the standard Gaussian $\boldsymbol{\eta}$. We first re-express it as a product of a radius r_g and a direction \mathbf{e}_g . r_g takes values inside \mathbb{R}_+ and \mathbf{e} takes values inside the d -dimensional unit sphere:

$$\begin{aligned}\boldsymbol{\eta} &= r_g \mathbf{e}_g \\ r_g &= \|\boldsymbol{\eta}\| \\ \mathbf{e}_g &= \frac{\boldsymbol{\eta}}{\|\boldsymbol{\eta}\|}\end{aligned}$$

Since $\boldsymbol{\eta}$ is a standard Gaussian distribution, it is symmetric. Thus the random variable \mathbf{e}_g is uniformly distributed over the d -dimensional sphere: S^{d-1} (with density the inverse of the area of S^{d-1}). r_g also has a straightforward distribution: its distribution is χ_d : a *chi-distribution* (pronounced “ki”) with d -degrees of freedom. Furthermore, random variables \mathbf{e}_g and r_g are independent (see Lemma 2).

Our final change of variable consists in remapping the random variable r_g using a simple bijection.

$$r_g = z_g^2$$

The bijection makes it so that large values of r_g are compressed which ensures, as we prove in the next section, that we end up with a strongly log-concave density $\tilde{f}(z)$.

We perform the exact same change of variable on θ_f . We thus transform the comparison of θ_g and θ_f into that of the pairs (z_g, e_g) and (z_f, e_f) :

$$\theta_g = \theta^* + z_g^2 \left(\Sigma^{1/2} e_g \right) \quad (15a)$$

$$\theta_f = \theta^* + z_f^2 \left(\Sigma^{1/2} e_f \right) \quad (15b)$$

We will denote the log-density of z_g and $z_f|e$ using respectively ψ_g and $\psi_{f,e}$, i.e:

$$\begin{aligned} \psi_g(z_g) &= -\log [g(z_g)] \\ \psi_{f,e}(z_f) &= -\log [f(z_f|e)] \end{aligned}$$

Note that we compare the conditional random variable $z_f|e$ to the marginal z_g because z_g and e_g are independent. The conditional random variable $z_g|e$ and the marginal random variable z_g are thus the same.

A.2 Density of (z, e) under the two models

Let us now study the density of the pair of random variables (z, e) under our two models.

A.2.1 Density of (z_g, e_g)

First, under the Gaussian density $g(\theta)$, these two variables are independent, and have simple densities, as summed-up by the following lemma.

Lemma 2. *Density of (z_g, e_g) .*

The random variables (z_g, e_g) are independent.

The random variable z_g follows a $\chi_d^{1/2}$ distribution with density:

$$g(z_g) = \frac{1}{2^{d/2-2}\Gamma(d/2)} z^{2d-1} \exp\left(-\frac{z^4}{2}\right) \quad (16)$$

The random variable e_g is uniformly distributed over its support: the d -dimensional sphere S^{d-1} .

Furthermore, the density $g(z_g)$ is strongly log-concave.

Lemma 3. *Strong log-concavity of $g(z_g)$.*

The log-density $\psi_g(z_g)$ is strongly concave:

$$\min_z \left[\psi_g''(z) \right] = 2\sqrt{6}\sqrt{2d-1}$$

Proof. Let us now prove both of these lemma.

First, we compute the density of r, e from the density of the standard Gaussian η . This just follows from a straightforward change of variable formula:

$$\begin{aligned} g(r, e) &\propto r^{d-1} \exp\left(-\frac{1}{2}(re)^T I_d(re)\right) 1(r \in \mathbb{R}_+) 1(e \in S^{d-1}) \\ &\propto r^{d-1} \exp\left(-\frac{r^2}{2}\right) 1(r \in \mathbb{R}_+) 1(e \in S^{d-1}) \end{aligned}$$

We observe that this decomposes into a product of the two marginal densities. We further observe that r_g follows a χ_d distribution, for which we know the normalization constant to be $\frac{1}{2^{d/2-1}\Gamma(d/2)}$. A further change of variable yields the density of $z = \sqrt{r}$:

$$\begin{aligned} g(z) &= \frac{1}{2^{d/2-1}\Gamma(d/2)} (z^2)^{d-1} \exp\left(-\frac{z^4}{2}\right) 2z \\ &= \frac{1}{2^{d/2-2}\Gamma(d/2)} z^{2d-1} \exp\left(-\frac{z^4}{2}\right) \end{aligned}$$

Second, let us study the log-concavity of this density. The negative log-density (which we need to show is strongly convex) is:

$$\psi_g(z) = -(2d-1)\log z + \frac{z^4}{2} + \log \left[2^{d/2-2}\Gamma(d/2) \right]$$

The derivatives of $\psi_g(z)$ are straightforward to compute:

$$\begin{aligned}\psi'_g(z) &= -\frac{2d-1}{z} + 2z^3 \\ \psi''_g(z) &= \frac{2d-1}{z^2} + 6z^2 \\ \psi^{(3)}_g(z) &= -2\frac{2d-1}{z^3} + 12z\end{aligned}$$

Observe that $\psi''_g(z) > 0$ so that ψ_g is at least strictly convex. Furthermore, $\psi^{(4)}_g(z) > 0$ so that $\psi''_g(z)$ is also strictly convex. It thus reaches its unique minimum at the point for which $\psi^{(3)}_g(z) = 0$. This point is such that:

$$\begin{aligned}-2\frac{2d-1}{z^3} + 12z &= 0 \\ z^4 &= \frac{2d-1}{6} \\ z &= \left(\frac{2d-1}{6} \right)^{1/4}\end{aligned}$$

Thus, the minimum curvature of $\psi_g(z)$ is:

$$\begin{aligned}\min_z \left[\psi''_g(z) \right] &= \frac{2d-1}{\left(\frac{2d-1}{6} \right)^{1/2}} + 6 \left(\frac{2d-1}{6} \right)^{1/2} \\ &= 6^{1/2} (2d-1)^{1/2} + 6^{1/2} (2d-1)^{1/2} \\ &= 2\sqrt{6}\sqrt{2d-1}\end{aligned}\tag{17}$$

which concludes our proof. \square

A.2.2 Density of (z_f, \mathbf{e}_f)

We now investigate the density of the pair (z, \mathbf{e}) under the target density f . The best description for this pair of variable is a hierarchical description in which the direction \mathbf{e}_f is picked first according to its marginal density $f(\mathbf{e}_f)$. The ‘‘square-root-radius’’ z_f (or equivalently, the radius r_f) is then picked according to its conditional distribution $z_f|\mathbf{e}_f$.

Our first lemma of this section describes the density marginal density of \mathbf{e}_f and the conditional density $z_f|\mathbf{e}_f$.

Lemma 4. *Density of (z_f, \mathbf{e}_f) .*

The conditional density of $z_f|\mathbf{e}_f$ is:

$$f(z_f|\mathbf{e}_f) \propto (z_f)^{2d-1} \exp\left(-\phi_f\left(\boldsymbol{\theta}^* + z^2\Sigma^{1/2}\mathbf{e}_f\right)\right)$$

The marginal distribution of \mathbf{e}_f is found by integrating out the conditional distribution of either $z_f|\mathbf{e}_f$ or $r|\mathbf{e}_f$:

$$\begin{aligned}f(\mathbf{e}_f) &\propto \int_{z \geq 0} (z_f)^{2d-1} \exp\left(-\phi_f\left(\boldsymbol{\theta}^* + z^2\Sigma^{1/2}\mathbf{e}_f\right)\right) dz_f \\ &\propto \int_{r \geq 0} (r_f)^{d-1} \exp\left(-\phi_f\left(\boldsymbol{\theta}^* + r\Sigma^{1/2}\mathbf{e}_f\right)\right) dr\end{aligned}$$

Furthermore, we can now finally highlight why the change of variable $r = z^2$ is important: as the next lemma asserts, this change of variable ensures that $f(z|e_f)$ is always strongly log-concave. In the limit where the higher derivatives of $\phi_f(\boldsymbol{\theta})$ along direction e_f become negligible, the minimum log-curvature of $f(z|e_f)$ even asymptotes to $\min_z \psi''(z)$.

In order to state this lemma, we will need additional notation. First, we need a shorter notation for the function $r \rightarrow \phi_f(\boldsymbol{\theta}^* + r\Sigma^{1/2}\mathbf{e})$ which we will denote with $\varphi_e(r)$. Second, we need to measure the derivatives of this function. Notice that we already know the first two derivatives of $\varphi_e(r)$:

$$\begin{aligned}\varphi'_e(r) &= \nabla\phi_f(\boldsymbol{\theta}^*)\Sigma^{1/2}\mathbf{e} = 0 \\ \varphi''_e(r) &= \mathbf{e}^T\Sigma^{1/2}H\phi_f(\boldsymbol{\theta}^*)\Sigma^{1/2}\mathbf{e} \\ &= \mathbf{e}^T\Sigma^{1/2}\Sigma^{-1}\Sigma^{1/2}\mathbf{e} \\ &= \mathbf{e}^T I_d \mathbf{e} \\ &= 1\end{aligned}$$

We will control the higher-derivatives using the following two quantities:

$$\Delta_3(\mathbf{e}) = \varphi_e^{(3)}(0) \tag{18a}$$

$$\Delta_4(\mathbf{e}) = \max_{r \geq 0} [\varphi_e^{(4)}(r)] \tag{18b}$$

Note that $\Delta_3(\mathbf{e})$ can be deduced from the third-derivative tensor: $\phi_f^{(3)}(\boldsymbol{\theta}^*)$ through:

$$\Delta_3(\mathbf{e}) = \phi_f^{(3)}(\boldsymbol{\theta}^*) \left[\Sigma^{1/2}\mathbf{e}, \Sigma^{1/2}\mathbf{e}, \Sigma^{1/2}\mathbf{e} \right] \tag{19}$$

Lemma 5. *Strong log-concavity of $f(z_f|\mathbf{e})$.*

For any \mathbf{e} , the conditional density $f(z_f|\mathbf{e})$ is strongly log-concave. The minimum curvature can be found numerically from the following formula:

$$r_0 = \frac{\Delta_3(\mathbf{e}) + \sqrt{[\Delta_3(\mathbf{e})]^2 + 2\Delta_4(\mathbf{e})}}{\Delta_4(\mathbf{e})}$$

$$\min \text{ curvature} = \min \left\{ r_0 + \Delta_3(\mathbf{e})(r_0)^2 - \Delta_4(\mathbf{e})\frac{(r_0)^3}{3}; \min_{0 \leq r \leq r_0} \frac{2d-1}{r} + 6r + 5\Delta_3(\mathbf{e})r^2 - \frac{7}{3}\Delta_4(\mathbf{e})r^4 \right\}$$

Proof. Now let us prove these two lemmas.

The first lemma is absolutely straightforward: it results from a straightforward change of variable formula.

The second lemma is more difficult. First, let us start by computing the derivatives of $\psi_{f,\mathbf{e}}(z)$:

$$\begin{aligned}\psi_{f,\mathbf{e}}(z) &= -(2d-1)\log(z) + \varphi_e(z^2) \\ \psi'_{f,\mathbf{e}}(z) &= -\frac{2d-1}{z} + 2z\varphi'_e(z^2) \\ \psi''_{f,\mathbf{e}}(z) &= \frac{2d-1}{z^2} + 2\varphi'_e(z^2) + 4z^2\varphi''_e(z^2)\end{aligned} \tag{20}$$

Let us consider the equation for $\psi''_{f,\mathbf{e}}(z)$. All of the terms are positive:

- The first term $(2d-1)/z^2 > 0$ is obvious
- The second term is also positive: since $\varphi''_e(r) > 0$, we have $\varphi'_e(r) > \varphi'_e(0) = 0$
- The final term is also positive: $\varphi''_e(r) > 0$ so that $4z^2\varphi''_e(z^2) > 0$

Furthermore, it is even true that $\psi''_{f,\mathbf{e}}(z)$ is lower-bounded, so that $\psi''_{f,\mathbf{e}}(z)$ is strongly log-concave. Indeed, we have:

$$\psi''_{f,\mathbf{e}}(z) > \frac{2d-1}{z^2} + 2\varphi'_e(z^2) \tag{21}$$

Critically, the first term is decreasing, and the second one increasing. For any z_0 , we have that, for all $z \geq z_0$:

$$\begin{aligned}\frac{2d-1}{z^2} + 2\varphi'_e(z^2) &\geq \frac{2d-1}{z_0^2} + 2\varphi'_e(z_0^2) \\ \psi''_{f,e}(z) &> 2\varphi'_e(z_0^2)\end{aligned}$$

Thus, for any z_0 , we have the following strictly positive lower bound for $\psi''_{f,e}(z)$:

$$\psi''_{f,e}(z) \geq \min \left\{ \min_{0 \leq z \leq z_0} \left[\frac{2d-1}{z^2} + 2\varphi'_e(z^2) + 4z^2\varphi''_e(z^2) \right]; 2\varphi'_e(z_0^2) \right\} > 0 \quad (22)$$

We have just proved that $f(z_f|e)$ is strongly log-concave. Let us now see how we should choose the value for z_0 .

We do so by first computing a Taylor expansion of $\varphi''_e(r)$:

$$\varphi''_e(r) \geq 1 + \Delta_3(e)r - \Delta_4(e)\frac{r^2}{2} \quad (23)$$

However, we have access to further information: we know that φ_e is strictly concave so that $\varphi''_e(r) > 0$.

Thus, there will be a critical value r_0 such that the lower-bound computed from the Taylor expansion is equal to 0. For $r \geq r_0$, the Taylor expansion bound gives no further information compared to simply knowing that $\varphi''_e(r) > 0$. This value r_0 is found by solving a second degree polynomial, yielding:

$$r_0 = \frac{\Delta_3(e) + \sqrt{[\Delta_3(e)]^2 + 2\Delta_4(e)}}{\Delta_4(e)} \quad (24)$$

This gives the limit of the zone for which our Taylor expansion is useful.

Now let us compute a Taylor expansion of $\varphi'_e(r)$. In the useful region $r \leq r_0$, we have:

$$\varphi'_e(r) \geq 0 + r + \Delta_3(e)\frac{r^2}{2} - \Delta_4(e)\frac{r^3}{3!} \quad (25)$$

For $r \geq r_0$, the only guarantee we have is that $\varphi'_e(r)$ is increasing so that:

$$\varphi'_e(r) \geq \varphi'_e(r_0) \geq r_0 + \Delta_3(e)\frac{(r_0)^2}{2} - \Delta_4(e)\frac{(r_0)^3}{3!} \quad (26)$$

We now combine these lower-bounds on $\varphi'_e(r)$ and $\varphi''_e(r)$ with the expression for $\psi''_{f,e}(z)$. For $z \leq z_0 = \sqrt{r_0}$, we have:

$$\begin{aligned}\psi''_{f,e}(z) &\geq \frac{2d-1}{z^2} + 2z^2 + \Delta_3(e)z^4 - \Delta_4(e)\frac{z^6}{3} + 4z^2 + \Delta_3(e)4z^4 - \Delta_4(e)2z^6 \\ &\geq \frac{2d-1}{z^2} + 6z^2 + 5\Delta_3(e)z^4 - \frac{7}{3}\Delta_4(e)z^6\end{aligned} \quad (27)$$

and for $z \geq z_0$, we only have the somewhat trivial bound:

$$\begin{aligned}\psi''_{f,e}(z) &> 2\varphi'_e(z_0^2) \\ \psi''_{f,e}(z) &\geq r_0 + \Delta_3(e)(r_0)^2 - \Delta_4(e)\frac{(r_0)^3}{3}\end{aligned} \quad (28)$$

At this point, we have an expression that is perfectly suitable for numerical optimization: we simply need to compute the extrema of a polynomial function over a finite range. The expression in the theorem is reached through the change of variable $r = z^2$. \square

A.3 Approximating z_f

We now turn to the task of computing whether the random variable z_g is a good approximation of $z_f|e$. More precisely, since we have proved that $f(z|e)$ is strongly log-concave, we can apply the Log-Sobolev Inequality (LSI; Otto and Villani [2000]) to upper-bound the KL divergence while avoiding the complicated task of upper-bounding the normalizing constant of $\tilde{f}(z|e)$.

The following lemma gives the result of applying the LSI. We express the results using properties of the distribution of the random variable r_g , which follows the more common χ_d distribution, instead of z_g . Moments of a χ_d random variable can be found in any thorough reference textbook on probability theory.

Lemma 6. $z_g \approx z_f|e$

The KL divergence between $g(z)$ and $f(z|e)$ is upper-bounded:

$$\begin{aligned} KL(z_g, z_f|e) &\leq \frac{E \left[4r_g \left(\varphi'_e(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} \left[\psi''_{f,e}(z) \right]} \\ &\leq \frac{[\Delta_3(e)]^2 E(r_g^5) + \frac{2}{3} |\Delta_3(e)| \Delta_4(e) E(r_g^6) + \frac{1}{9} [\Delta_4(e)]^2 E(r_g^7)}{\min_{z \geq 0} \left[\psi''_{f,e}(z) \right]} \\ &\gtrsim \frac{[\Delta_3(e)]^2 E(r_g^5)}{2\sqrt{6}\sqrt{2d-1}} \end{aligned}$$

Proof. This lemma is proved by a simple combination of the LSI with a Taylor expansion of $\varphi'_e(r)$ around 0.

First, we observe that $f(z|e)$ is a strongly log-concave density with minimal curvature $\min \left[\psi''_{f,e}(z) \right]$ (from lemma 5). Thus, we can apply the LSI:

$$\begin{aligned} KL(z_g, z_f|e) &\leq \frac{E \left[\left(\psi'_{f,e}(z_g) - \psi'_g(z_g) \right)^2 \right]}{\min_{z \geq 0} \left[\psi''_{f,e}(z) \right]} \\ &\leq \frac{E \left[4z_g^2 \left(\varphi'_e(z_g^2) - z_g^2 \right)^2 \right]}{\min_{z \geq 0} \left[\psi''_{f,e}(z) \right]} \\ &\leq \frac{E \left[4r_g \left(\varphi'_e(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} \left[\psi''_{f,e}(z) \right]} \end{aligned}$$

where the following lines correspond to simple substitutions: $\psi'_{f,e}(z) = 2z\varphi_e(z^2)$ and $r_g = z_g^2$.

We then turn to a Taylor expansion of $\varphi'_e(r_g) - r_g$ around 0. Critically, the first two terms are 0 because our Gaussian approximation is the Laplace approximation:

$$\begin{aligned} \left| \varphi'_e(r_g) - r_g - \Delta_3(\mathbf{e}) \frac{r_g^2}{2} \right| &\leq \Delta_4(\mathbf{e}) \frac{r_g^3}{3!} \\ \left| \varphi'_e(r_g) - r_g \right| &\leq |\Delta_3(\mathbf{e})| \frac{r_g^2}{2} + \Delta_4(\mathbf{e}) \frac{r_g^3}{3!} \\ \left(\varphi'_e(r_g) - r_g \right)^2 &\leq \left(|\Delta_3(\mathbf{e})| \frac{r_g^2}{2} + \Delta_4(\mathbf{e}) \frac{r_g^3}{3!} \right)^2 \\ 4r_g \left(\varphi'_e(r_g) - r_g \right)^2 &\leq r_g \left(|\Delta_3(\mathbf{e})| r_g^2 + \Delta_4(\mathbf{e}) \frac{r_g^3}{3} \right)^2 \end{aligned}$$

We can then easily compute the expected value of the last bound, yielding:

$$E \left[4r_g \left(\varphi'_e(r_g) - r_g \right)^2 \right] \leq [\Delta_3(\mathbf{e})]^2 E(r_g^5) + \frac{2}{3} |\Delta_3(\mathbf{e})| \Delta_4(\mathbf{e}) E(r_g^6) + \frac{1}{9} [\Delta_4(\mathbf{e})]^2 E(r_g^7)$$

which yields the claimed result. \square

A.4 Approximating e_f

We can now turn to the task of computing whether e_g gives a good approximation of e_f . This corresponds to checking whether e_f has an almost uniform distribution over the unit sphere S^{d-1} . Equivalently, we will check whether the log-density $\log [\tilde{f}(\mathbf{e})]$ has small oscillations.

The value of $\tilde{f}(\mathbf{e})$ is found by integrating out the unnormalized density $\tilde{f}(r|\mathbf{e})$ (or equivalently $\tilde{f}(z|\mathbf{e})$). The oscillations in $\xi(\mathbf{e}) = \log [\tilde{f}(\mathbf{e})]$ are caused by the fact that the higher-derivatives of $\varphi_e(r)$ differ depending on the direction \mathbf{e} : as the following lemma shows, $\Delta_3(\mathbf{e})$ is the main influence on $\xi(\mathbf{e}) = \log [\tilde{f}(\mathbf{e})]$.

Lemma 7. *Oscillations of $\log [\tilde{f}\mathbf{e}]$.*

$\log [\tilde{f}(\mathbf{e})]$ can be approximated using the ELBO:

$$\xi(\mathbf{e}) = \log [\tilde{f}(\mathbf{e})] = C + E_g \left(\frac{r_g^2}{2} - \varphi_e(r_g) \right) + \epsilon_1(\mathbf{e})$$

where $\epsilon_1(\mathbf{e})$ is a positive error equal precisely to $KL(z_g, z_f|\mathbf{e})$ and thus upper-bounded by lemma 6:

$$\epsilon_1(\mathbf{e}) = KL(z_g, z_f|\mathbf{e}) \leq \frac{E \left[4r_g \left(\varphi'_e(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f,\mathbf{e}}(z)]}$$

We can further perform a Taylor expansion of $\varphi_e(r_g)$ to get:

$$\begin{aligned} \xi(\mathbf{e}) &= C - \frac{\Delta_3(\mathbf{e}) E(r_g^3)}{6} + \epsilon_1(\mathbf{e}) + \epsilon_2(\mathbf{e}) \\ |\epsilon_2(\mathbf{e})| &\leq \frac{\Delta_4(\mathbf{e}) E(r_g^4)}{4!} \end{aligned}$$

If these derivatives are small on average, then the KL divergence $KL(e_g, e_f)$ is small, as the following lemma shows.

Lemma 8. $\mathbf{e}_g \approx \mathbf{e}_f$.

The KL divergence can be re-expressed as:

$$KL(\mathbf{e}_g, \mathbf{e}_f) = \log [E (\exp [\xi (\mathbf{e}_g) - E (\xi (\mathbf{e}_g))])]$$

and then upper-bounded or approximated using $\text{var} [\xi (\mathbf{e}_g)]$:

$$\begin{aligned} KL(\mathbf{e}_g, \mathbf{e}_f) &\leq \log \left[1 + \frac{1}{2} \exp \left(\max_{\mathbf{e}} (\xi (\mathbf{e})) - E (\xi (\mathbf{e}_g)) \right) \text{var} [\xi (\mathbf{e}_g)] \right] \\ KL(\mathbf{e}_g, \mathbf{e}_f) &\approx \frac{1}{2} \text{var} [\xi (\mathbf{e}_g)] \\ &\approx \frac{1}{2} \left[\frac{E (r_g^3)}{6} \right]^2 \text{var} [\Delta_3 (\mathbf{e}_g)] \end{aligned}$$

Another useful upper-bound and approximation is found by separating the summands in $\xi (\mathbf{e})$. Noting

$\xi_{ELBO} (\mathbf{e}) = E_g \left(\frac{r_g^2}{2} - \varphi_e (r_g) \right)$ the ELBO approximation of $\xi (\mathbf{e})$, we have:

$$\begin{aligned} KL(\mathbf{e}_g, \mathbf{e}_f) &\leq \frac{1}{2} \log [E [\exp (2\xi_{ELBO} (\mathbf{e}_g) - 2E [\xi_{ELBO} (\mathbf{e}_g))])] + \frac{1}{2} \log [E [\exp (2\epsilon_1 (\mathbf{e}_g) - 2E [\epsilon_1 (\mathbf{e}_g))])] \\ &\lesssim \frac{1}{2} \log [E [\exp (2\xi_{ELBO} (\mathbf{e}_g) - 2E [\xi_{ELBO} (\mathbf{e}_g))])] + \text{var} (\epsilon_1 (\mathbf{e}_g)) \end{aligned}$$

We can also use $\xi (\mathbf{e}) = \frac{\Delta_3(\mathbf{e}_g)E(r_g^3)}{6} + \epsilon_1 (\mathbf{e}) + \epsilon_2 (\mathbf{e})$:

$$\begin{aligned} KL(\mathbf{e}_g, \mathbf{e}_f) &\leq \frac{1}{2} \log \left[E \left[\exp \left(2 \frac{\Delta_3 (\mathbf{e}_g) E (r_g^3)}{6} \right) \right] \right] + \frac{1}{2} \log [E [\exp (2\epsilon_1 (\mathbf{e}_g) + 2\epsilon_2 (\mathbf{e}_g) - 2E [\epsilon_1 (\mathbf{e}_g) + \epsilon_2 (\mathbf{e}_g))])] \\ &\lesssim \left[\frac{E (r_g^3)}{6} \right]^2 \text{var} [\Delta_3 (\mathbf{e}_g)] + \text{var} [\epsilon_1 (\mathbf{e}_g) + \epsilon_2 (\mathbf{e}_g)] \end{aligned}$$

Proof. Let us start by rewriting the KL divergence to make \tilde{f} and then $\xi = \log [\tilde{f}]$ appear:

$$\begin{aligned} KL(\mathbf{e}_g, \mathbf{e}_f) &= E \left(\log \frac{g(\mathbf{e}_g)}{f(\mathbf{e}_g)} \right) \\ &= E \left(\log \frac{g(\mathbf{e}_g) (f \tilde{f})}{\tilde{f}(\mathbf{e}_g)} \right) \\ &= E \left(\log \frac{g(\mathbf{e}_g)}{\tilde{f}(\mathbf{e}_g)} \right) + \log \left[\int \tilde{f} \right] \\ &= E \left(\log \frac{g(\mathbf{e}_g)}{\tilde{f}(\mathbf{e}_g)} \right) + \log \left[E \left(\frac{\tilde{f}(\mathbf{e}_g)}{g(\mathbf{e}_g)} \right) \right] \\ &= \log \left[E \left(\frac{\tilde{f}(\mathbf{e}_g)}{g(\mathbf{e}_g)} \right) \right] - E \left(\log \frac{\tilde{f}(\mathbf{e}_g)}{g(\mathbf{e}_g)} \right) \\ &= \log \left[E \left(\exp \left[\log \frac{\tilde{f}(\mathbf{e}_g)}{g(\mathbf{e}_g)} \right] \right) \right] - E \left(\log \frac{\tilde{f}(\mathbf{e}_g)}{g(\mathbf{e}_g)} \right) \end{aligned}$$

In this final expression, we can further remove $g(\mathbf{e}_g)$ which is constant, and bring $E(\log \tilde{f})$ into the first term:

$$KL(\mathbf{e}_g, \mathbf{e}_f) = \log [E (\exp [\xi (\mathbf{e}_g) - E (\xi (\mathbf{e}_g))])] \quad (29)$$

we are thus left with a final expression for the KL divergence which only involves expected values of $\xi (\mathbf{e}_g) = \log [\tilde{f} (\mathbf{e}_g)]$.

We now need to upper-bound this KL divergence to prove Lemma 8. We do so by upper-bounding the exponential function. We will use the following bound which holds for all $x \leq M$:

$$\exp(x) \leq 1 + x + \frac{x^2}{2} \exp(M)$$

This bound holds because these two functions have the same value and the same first derivative at 0, and their second derivatives verify:

$$\exp(x) \leq \exp(M)$$

Thus, their difference: $1 + x + \frac{x^2}{2} \exp(M) - \exp(x)$, is a convex function with minimum at 0, from which we deduce the inequality.

Armed with this upper-bound on the exponential function, let us return to the KL divergence, denoting $\xi(\mathbf{e}) = \log \left[\tilde{f}(\mathbf{e}) \right]$:

$$\begin{aligned} KL(\mathbf{e}_g, \mathbf{e}_f) &= \log \left[E \left(\exp \left[-\xi(\mathbf{e}_g) + E(\xi(\mathbf{e}_g)) \right] \right) \right] \\ &\leq \log \left[1 - E(\xi(\mathbf{e}_g) + E(\xi(\mathbf{e}_g))) + \exp \left(\max_{\mathbf{e}} (\xi(\mathbf{e})) - E(\xi(\mathbf{e}_g)) \right) \frac{E \left([\xi(\mathbf{e}_g) - E(\xi(\mathbf{e}_g))]^2 \right)}{2} \right] \\ &\leq \log \left[1 + 0 + \exp \left(\max_{\mathbf{e}} (\xi(\mathbf{e})) - E(\xi(\mathbf{e}_g)) \right) \frac{E \left([\xi(\mathbf{e}_g) - E(\xi(\mathbf{e}_g))]^2 \right)}{2} \right] \\ &\leq \log \left[1 + \frac{1}{2} \exp \left(\max_{\mathbf{e}} (\xi(\mathbf{e})) - E(\xi(\mathbf{e}_g)) \right) \text{var} [\xi(\mathbf{e}_g)] \right] \end{aligned}$$

We could also use a simple approximation of $\exp(x)$ and $\log(1+x)$ around 0. We would then get:

$$\begin{aligned} \exp(x) &\approx 1 + x + \frac{x^2}{2} \\ \log(1+y) &\approx y \\ KL(\mathbf{e}_g, \mathbf{e}_f) &\approx \log \left[1 + \frac{1}{2} \text{var} [\xi(\mathbf{e}_g)] \right] \\ &\approx \frac{1}{2} \text{var} [\xi(\mathbf{e}_g)] \end{aligned}$$

Now we will work on splitting $\xi(\mathbf{e})$. From Lemma 4, we know that:

$$\begin{aligned} \xi(\mathbf{e}) &= \log \left[\tilde{f}(\mathbf{e}) \right] = \log \left[\int \tilde{f}(r|\mathbf{e}) dr \right] \\ &= C + \log \left[\int r^{d-1} \exp(-\varphi_{\mathbf{e}}(r)) dr \right] \end{aligned}$$

where the constant C doesn't depend on \mathbf{e} .

A good approximation of integrals of these form is the Evidence Lower Bound (ELBO; Murphy [2012] Chapter 21):

$$\xi(\mathbf{e}) \approx E_g \left(\frac{r_g^2}{2} - \varphi_{\mathbf{e}}(r) \right)$$

Like its name indicates, the ELBO lower-bounds $\xi(\mathbf{e})$: the difference between the two is precisely equal to $KL(r_g, r_f|\mathbf{e})$.

We have bounded the KL divergence between r_g and $r_f|e$. Thus, we can precisely control the error using the KL divergence upper-bound of lemma 6. Thus :

$$\begin{aligned}\xi(e) &= C_2 + E\left(\frac{r_g^2}{2} - \varphi_e(r_g)\right) + \epsilon_1(e) \\ \epsilon_1(e) &= KL(r_g, r_f|e) \\ 0 \leq \epsilon_1(e) &\leq \frac{E\left[4r_g\left(\varphi'_e(r_g) - r_g\right)^2\right]}{\min_{z \geq 0} [\psi''_{f,e}(z)]}\end{aligned}$$

Note that the constant C_2 doesn't depend on e and will vanish when in the KL divergence since we are interested in $\xi(e_g) - E(\xi(e_g))$.

The term that causes the majority of the oscillations of $\xi(e)$ is $E\left(\frac{r_g^2}{2} - \varphi_e(r_g)\right)$. Let us now compute the size of this term:

$$\begin{aligned}\left|\frac{r_g^2}{2} - \varphi_e(r_g) + \Delta_3(e) \frac{r_g^3}{3!}\right| &\leq \Delta_4(e) \frac{r_g^4}{4!} \\ \left|E\left(\frac{r_g^2}{2} - \varphi_e(r_g)\right) - \frac{\Delta_3(e)}{3!} E(r_g^3)\right| &\leq \frac{\Delta_4(e)}{4!} E(r_g^4)\end{aligned}$$

We can then rewrite $\xi(e)$ as:

$$\begin{aligned}\xi(e) &= C_2 - \frac{\Delta_3(e) E(r_g^3)}{6} + \epsilon_1(e) + \epsilon_2(e) \\ 0 \leq \epsilon_1(e) &\leq \frac{E\left[4r_g\left(\varphi'_e(r_g) - r_g\right)^2\right]}{\min_{z \geq 0} [\psi''_{f,e}(z)]} \\ |\epsilon_2(e)| &\leq \frac{\Delta_4(e)}{4!} E(r_g^4)\end{aligned}$$

Finally, we can use this decomposition of $\xi(e)$ into multiple terms when considering the KL divergence. To prove this, consider the following function:

$$K(\lambda) = \log[E(\exp[\lambda A(e_g) + (1-\lambda) B(e_g)])]$$

where $A(e_g)$ and $B(e_g)$ are arbitrary functions. The first derivative of K is:

$$\begin{aligned}K'(\lambda) &= \frac{\int g(e) de [A(e) - B(e)] \exp[\lambda A(e) + (1-\lambda) B(e)]}{\int g(e) de \exp[\lambda A(e) + (1-\lambda) B(e)]} \\ &= E[A(e) - B(e) | \lambda]\end{aligned}$$

where the expected value is computed against the normalized density

$$g(e|\lambda) = \frac{g(e) \exp[\lambda A(e) + (1-\lambda) B(e)]}{\int g(e) de \exp[\lambda A(e) + (1-\lambda) B(e)]}$$

The second derivative of K is:

$$\begin{aligned}K''(\lambda) &= E\left[(A(e) - B(e))^2 | \lambda\right] - E[A(e) - B(e) | \lambda]^2 \\ &= \text{var}[A(e) - B(e) | \lambda]\end{aligned}$$

Critically: $K''(\lambda) \geq 0$ and K is thus a convex function.

By the convexity of $K(\lambda)$, we have:

$$K(0.5) \leq K(0) + K(1)$$

Now, take $A(\mathbf{e}) = 2 \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} - E \left(2 \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} \right)$ and $B(\mathbf{e}) = 2\epsilon_1(\mathbf{e}) + 2\epsilon_2(\mathbf{e}) - 2E[\epsilon_1(\mathbf{e}) + \epsilon_2(\mathbf{e})]$. By combining the symmetry of \mathbf{e}_g with the asymmetry of $\Delta_3(\mathbf{e}_g)$, we have that $E \left(2 \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} \right) = 0$.

The values of K are then:

$$\begin{aligned} K(0.5) &= \log [E [\exp (\xi(\mathbf{e}) - E(\xi(\mathbf{e})))] \\ K(0) &= \log \left[E \left[\exp \left(2 \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} \right) \right] \right] \\ K(1) &= \log \left[E \left[\exp \left(2\epsilon(\mathbf{e}) + 2\epsilon'(\mathbf{e}) - 2E[\epsilon(\mathbf{e}) + \epsilon'(\mathbf{e})] \right) \right] \right] \end{aligned}$$

and we find the final inequality:

$$\frac{1}{2} \log \left[E \left[\exp \left(2 \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} \right) \right] \right] + \frac{1}{2} \log \left[E \left[\exp \left(2\epsilon(\mathbf{e}_g) + 2\epsilon'(\mathbf{e}_g) - 2E[\epsilon(\mathbf{e}_g) + \epsilon'(\mathbf{e}_g)] \right) \right] \right]$$

Note that for any $\alpha \in]0, 1[$, we could write:

$$\xi(\mathbf{e}) - E(\xi(\mathbf{e})) = \alpha \frac{1}{\alpha} \frac{\Delta_3(\mathbf{e}_g) E(r_g^3)}{6} + (1 - \alpha) \frac{1}{1 - \alpha} \left(\epsilon(\mathbf{e}_g) + \epsilon'(\mathbf{e}_g) - E[\epsilon(\mathbf{e}_g) + \epsilon'(\mathbf{e}_g)] \right)$$

which can give other useful upper-bounds on the KL divergence. \square

A.5 Approximating f

We are now finally ready to combine all of the preceding lemmas to state the full form of our theorem:

1. We have determined the density of the pairs (z_g, \mathbf{e}_g) and (z_f, \mathbf{e}_f) .
2. We have computed the KL divergence between z_g and the conditional distribution $z_f | \mathbf{e}$.
3. We have computed the KL divergence between \mathbf{e}_g and \mathbf{e}_f .

Now, the only step that remains consists in combining those results, which we are able to do because the KL divergence is invariant to changes of variables.

Theorem 9. *A detailed upper-bound on the KL divergence.*

The KL divergence between g and f can be decomposed as:

$$KL(g, f) = KL(\mathbf{e}_f, \mathbf{e}_g) + E(KL(z_g, z_f | \mathbf{e}_g))$$

where the second expected value computes the mean of the function $\mathbf{e} \rightarrow KL(z_g, z_f | \mathbf{e})$ over the random variable \mathbf{e}_g .

The KL divergence can be upper-bounded as:

$$\begin{aligned} KL(g, f) &\leq \frac{1}{2} \log [E [\exp (2\xi_{ELBO}(\mathbf{e}_g) - 2E[\xi_{ELBO}(\mathbf{e}_g)])]] + \frac{1}{2} \log [E [\exp (2\epsilon_1(\mathbf{e}_g) - 2E[\epsilon_1(\mathbf{e}_g)])]] \\ &\quad + E_{\mathbf{e}_g} \left[\frac{E_{r_g} \left[4r_g \left(\varphi'_{\mathbf{e}_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, \mathbf{e}_g}(z)]} \right] \end{aligned} \quad (30)$$

A computable approximation is found by approximating the term involving $\epsilon_1(\mathbf{e}_g)$:

$$\begin{aligned} KL(g, f) &\leq \frac{1}{2} \log [E [\exp (2\xi_{ELBO}(\mathbf{e}_g) - 2E[\xi_{ELBO}(\mathbf{e}_g)])]] + E \left[\left(\frac{E_{r_g} \left[4r_g \left(\varphi'_{\mathbf{e}_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, \mathbf{e}_g}(z)]} \right)^2 \right] \\ &\quad + E_{\mathbf{e}_g} \left[\frac{E_{r_g} \left[4r_g \left(\varphi'_{\mathbf{e}_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, \mathbf{e}_g}(z)]} \right] \end{aligned} \quad (31)$$

in which all expected values need to be approximated by sampling from e_g .

We can also approximate this bound by keeping only the terms that depend on $\Delta_3(e)$:

$$KL(g, f) \lesssim E \left[\Delta_3(e_g)^2 \right] \left(\frac{2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]} + \frac{1}{9} \left[\frac{\Gamma[(d+3)/2]}{\Gamma[d/2]} \right]^2 \right) \quad (32)$$

Proof. Let us prove the formula for the decomposition of the KL divergence. We start from the normal formula for the KL divergence:

$$KL(g, f) = \int g(\boldsymbol{\theta}) \log \left[\frac{g(\boldsymbol{\theta})}{f(\boldsymbol{\theta})} \right] d\boldsymbol{\theta}$$

We then perform the change of variable: $\boldsymbol{\theta} \rightarrow (z, e)$ and get:

$$\begin{aligned} KL(g, f) &= \int g(z) g(e) \log \left[\frac{g(z)g(e)}{f(z, e)} \right] dz de \\ &= \int g(z) g(e) \log \left[\frac{g(z)g(e)}{f(z|e)f(e)} \right] dz de \\ &= \int g(z) g(e) \left(\log \left[\frac{g(e)}{f(e)} \right] + \log \left[\frac{g(z)}{f(z|e)} \right] \right) dz de \\ &= \left(\int g(e) \log \left[\frac{g(e)}{f(e)} \right] \right) + \int g(e) \left(\int g(z) \log \left[\frac{g(z)}{f(z|e)} \right] dz \right) de \\ &= KL(e_g, e_f) + \int g(e) (KL(z_g, z_f|e)) de \\ &= KL(e_g, e_f) + E[KL(z_g, z_f|e_g)] \end{aligned}$$

We then combine this expression for the KL divergence with Lemmas 6 and 8. Using the most general approximation form from these theorems, we get:

$$\begin{aligned} KL(g, f) &\leq \frac{1}{2} \log [E[\exp(2\xi_{ELBO}(e_g) - 2E[\xi_{ELBO}(e_g)])]] + \frac{1}{2} \log [E[\exp(2\epsilon_1(e_g) - 2E[\epsilon_1(e_g)])]] \\ &\quad + E_{e_g} \left[\frac{E_{r_g} \left[4r_g \left(\varphi'_{e_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, e_g}(z)]} \right] \end{aligned}$$

in which $\epsilon_1(e) = KL(z_g, z_f|e)$ is bounded:

$$0 \leq \epsilon_1(e) \leq \frac{E_{r_g} \left[4r_g \left(\varphi'_{e_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, e_g}(z)]}$$

The term $\frac{1}{2} \log [E[\exp(2\epsilon_1(e_g) - 2E[\epsilon_1(e_g)])]]$ is problematic since we only have an upper-bound on each term $\epsilon_1(e)$. One solution consists in using the approximation (see the proof of Lemma 8 in which we have detailed further how to approximate this term):

$$\begin{aligned} \frac{1}{2} \log [E[\exp(2\epsilon(e_g) - 2E[\epsilon(e_g)])]] &\approx \text{var}[\epsilon_1(e_g)] \\ &\lesssim E[(\epsilon_1(e_g))^2] \\ &\lesssim E \left[\left(\frac{E_{r_g} \left[4r_g \left(\varphi'_{e_g}(r_g) - r_g \right)^2 \right]}{\min_{z \geq 0} [\psi''_{f, e_g}(z)]} \right)^2 \right] \end{aligned}$$

The final formula of the theorem is simply found by combining the approximate forms of Lemmas 6 and 8, yielding:

$$KL(g, f) \lesssim \frac{1}{2} \left[\frac{E(r_g^3)}{6} \right]^2 E[\Delta_3(e_g)^2] + \frac{E(r_g^5)}{2\sqrt{6}\sqrt{2d-1}} E[\Delta_3(e_g)^2]$$

We then use the expression for the moments of r_g :

$$E(r_g^3) = 2\sqrt{2} \frac{\Gamma[(d+3)/2]}{\Gamma[d/2]}$$

$$E(r_g^5) = 4\sqrt{2} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]}$$

We finally obtain:

$$\begin{aligned} KL(g, f) &\lesssim \frac{1}{2} \left[\frac{2\sqrt{2} \frac{\Gamma[(d+3)/2]}{\Gamma[d/2]}}{6} \right]^2 E[\Delta_3(e_g)^2] + \frac{4\sqrt{2} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]}}{2\sqrt{6}\sqrt{2d-1}} E[\Delta_3(e_g)^2] \\ &\lesssim \frac{1}{2} \left[\frac{\sqrt{2}\Gamma[(d+3)/2]}{3\Gamma[d/2]} \right]^2 E[\Delta_3(e_g)^2] + \frac{2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]} E[\Delta_3(e_g)^2] \\ &\lesssim E[\Delta_3(e_g)^2] \left(\frac{1}{2} \frac{2}{9} \left[\frac{\Gamma[(d+3)/2]}{\Gamma[d/2]} \right]^2 + \frac{2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]} \right) \\ &\lesssim E[\Delta_3(e_g)^2] \left(\frac{2}{\sqrt{3}\sqrt{2d-1}} \frac{\Gamma[(d+5)/2]}{\Gamma[d/2]} + \frac{1}{9} \left[\frac{\Gamma[(d+3)/2]}{\Gamma[d/2]} \right]^2 \right) \end{aligned}$$

□

B Details of the logistic regression example

In order to assess whether the rough approximation of the bound could prove useful, we have tested it in a simple example: logistic regression. In this model, the data \mathcal{D} is composed of the n pairs: $(y_i, \mathbf{x}_i) \in \{-1, 1\} * \mathbb{R}^d$ corresponding to class labels y_i and predictors (or covariates) \mathbf{x}_i . If we choose the prior to be Gaussian (with mean 0 and covariance matrix $\sigma_0^2 I_d$), the posterior is then:

$$f(\boldsymbol{\theta}) = \frac{1}{Z} \exp \left(-\frac{1}{2} \sum_{j=1}^d \frac{(\theta_j)^2}{\sigma_0^2} \right) \prod_{i=1}^n \frac{1}{1 + \exp[-y_i(\boldsymbol{\theta} \cdot \mathbf{x}_i)]} \quad (33)$$

which is log-concave.

In order to test our theorem, we needed to:

- Compute the Laplace approximation of the posterior, i.e: compute $\boldsymbol{\theta}^*$ and $H\phi_f(\boldsymbol{\theta}^*)$.
- Compute the third derivative tensor: $\phi_f^{(3)}(\boldsymbol{\theta}^*)$ and use it to deduce $\Delta_3(e)$:

$$\Delta_3(e) = \phi_f^{(3)}(\boldsymbol{\theta}^*) \left[\Sigma^{1/2} e, \Sigma^{1/2} e, \Sigma^{1/2} e \right]$$

- Approximate the real value of the KL divergence

$$KL(g, f) = \int g(\boldsymbol{\theta}) \log \left[\frac{g(\boldsymbol{\theta}) Z}{\tilde{f}(\boldsymbol{\theta})} \right] d\boldsymbol{\theta}$$

which requires approximating the normalizing constant Z .

We computed these values in the following way:

- The MAP value $\boldsymbol{\theta}^*$ was computed using a line-search gradient descent.

- The normalizing constant was approximated using a Markov Chain Monte Carlo sampling method:
 - Using standard Metropolis-Hastings, we generated k samples from $f(\boldsymbol{\theta})$: $\boldsymbol{\theta}_i$.
In the examples reported, we generated 10^7 samples and we then kept one in 1000 thus giving $k = 10^4$.
 - Using the samples, we approximated:

$$\frac{1}{Z} = E \left(\frac{g(\boldsymbol{\theta}_f)}{\tilde{f}(\boldsymbol{\theta}_f)} \right) \approx \frac{1}{k} \sum_i \frac{g(\boldsymbol{\theta}_i)}{\tilde{f}(\boldsymbol{\theta}_i)}$$

- The KL divergence was also computed by sampling:
 - We generated k_2 samples from the Gaussian density $g(\boldsymbol{\theta})$.
In the examples reported, $k_2 = 10^5$.
 - We used those samples to approximate the KL divergence:

$$KL(g, f) = E \left(\log \left[\frac{g(\boldsymbol{\theta}_g) Z}{\tilde{f}(\boldsymbol{\theta}_g)} \right] \right)$$

in which we reused the approximation of Z .

The data was generated from the logistic regression model itself:

- First, we picked the covariates \boldsymbol{x}_i according to a Gaussian distribution with mean 0 and variance the identity matrix.
- Then, we picked the true value of the parameter $\boldsymbol{\theta}_0$ from a Gaussian distribution with mean 0 and variance $d^{-1/2}I_d$. This scaling ensures that the values $\boldsymbol{\theta} \cdot \boldsymbol{x}_i$ are of order 1 no matter the dimension d .
- Finally, the label y_i was picked according to the logistic density:

$$\mathbb{P}(y_i = \pm 1) = \frac{1}{1 + \exp[-y_i(\boldsymbol{\theta} \cdot \boldsymbol{x}_i)]}$$