Statistical Science
2017, Vol. 32, No. 1, 36–40
DOI: 10.1214/16-STS603
Main article DOI: 10.1214/16-STS576

How Principled and Practical Are Penalised Complexity Priors?

Christian P. Robert and Judith Rousseau

1. INTRODUCTION

This note discusses the paper "Penalising model component complexity" by Simpson et al. (2017). We acknowledge the highly novel approach to prior construction and commend the authors for setting new all-encompassing principles that will certainly impact Bayesian modelling. We also perceive the potential connection with other branches of the literature. Nonetheless, we remain uncertain as to what extent the principles exposed in the paper can be developed outside specific models, given the lack of precision in the said principles. The very notions of model component, base model and overfitting prior, are for instance much more conceptual than mathematical. We thus fear the advocated concept of penalised complexity may not reach further than extending first-guess priors into larger families, thus failing to establish reference priors on a novel and sound ground.

> "On the other end of the hunt for the holy grail, 'objective' priors are data-dependent and are not uniformly accepted among Bayesians on philosophical grounds."

The most sensitive aspect of Bayesian modelling is undoubtedly the call to a prior distribution. From Fisher onwards (Zabell, 1992), up to this very day (Martin and Liu, 2015, Seaman, Seaman and Stamey, 2012), the concept of prior distribution has been criticised as being alien to the sampling model and critics have pointed out the arbitrariness of some or all aspects of chosen priors. This is most prominent in weakly informative settings when the context is deemed too poor to return an expert opinion, and thus build an informed prior. The whole branch of so-called objective (aka, reference or noninformative) Bayesian statistics (Berger, Bernardo and Sun, 2009) has been constructed to answer and bypass such criticisms, clearly not achieving a complete silencing of such criticisms.

> "Prior selection is the fundamental issue in Bayesian statistics. Priors are the Bayesian's greatest tool, but they are also the greatest point for criticism: the arbitrariness of prior selection procedures and the lack of realistic sensitivity analysis (...) are a serious argument against current Bayesian practice."

In this paper, the authors aim at providing some form of prior robust modelling, rather than noninformative principles that are so delicate to specify, as shown by the literature (Liseo, 2005). It is a highly timely and pertinent paper on the selection and construction of priors. It also shows that the field of "objective" Bayes theory is still central to Bayesian statistics and this constitutes a great argument to encourage more Bayesian researchers to consider this branch of our field. This attempt is most commendable and we hope it will induce others to enlarge and deepen the work in this direction.

The paper starts with a review of prior selection in connection with levels of prior information. The authors then advance some desirable principles for the construction of priors on a collection of models that is restricted to hierarchical additive models with a latent structure. Connections with other approaches abound, from Jeffreys' priors and the asymptotic developments of Bochkina and Green (2014), to the nonlocal priors of Johnson and Rossell (2010), and sparsity priors. (However, this may constitute the more tentative part of the paper.) The applications are the disease mapping model of Besag et al. (1991) and the multivariate probit model.

2. PC PRIORS

> "Most model components can be naturally regarded as a flexible version of a base model."

Christian P. Robert is Professor, CEREMADE, Université Paris-Dauphine, PSL, 75775 Paris cedex 16, France (e-mail: xian@ceremade.dauphine.fr). Judith Rousseau is Professor, CEREMADE, Université Paris-Dauphine, PSL, 75775 Paris cedex 16, France (e-mail: rousseau@ceremade.dauphine.fr). Both authors are members of Laboratoire de Statistique, CREST, Paris. C. P. Robert is also affiliated as a part-time professor in the Department of Statistics of the University of Warwick, Coventry, UK, and is a senior member of IUF for 2016-2021.

The starting point for the authors' modelling is the so-called *base model*. We understand this approach operates via the (specialised?) notion of model components, as modularity is obviously essential (if challenging) in devising reference or default priors. However, the obvious question it induces is to figure how easy it is to define this base model. For instance, the authors later make a connection between base models and hypothesis testing. One may wonder whether or not such a notion always translates into a null hypothesis formulation and whether or not this reformulation is pertinent (if only because it relates to tests). From a more global perspective, we remain rather skeptical that there could be an automated version of the derivation of a base model, just as there is no single version of an "objective" Bayes prior (Kass and Wasserman, 1996, Robert, 2001). We assume this derivation somewhat follows from the "block" idea but we wonder at how generic model construction by blocks can be. The authors do acknowledge the difficulty in Section 7, as an unrealistic expectation on the practitioners.

In particular, defining a base model is typically done under a given parameterisation, which implies the whole approach fails to stay invariant under reparameterisation. To illustrate our point, consider the discussion in Section 4.5 on sparsity. The authors consider the model

$$\mathbf{y} = (y_1, \ldots, y_n) \sim \pi(\mathbf{y}|\beta), \quad \beta \sim \mathcal{N}_p(0, D),$$
$$D = \operatorname{diag}(D_{11}, \ldots, D_{pp}).$$

They define as their base model a sparse model, in essence corresponding to $\beta = 0$, that they translate into saying $D_{ii} = 0$ for all i and then into considering as a prior on D

$$D_{ii}^{-1} \sim \pi_\tau \tag{1}$$

where π_τ is the prior defined in (3.3). As the authors note, this does not lead to correct sparse behaviour (in other words to convincing shrinkage). They thus suggest to define as a sparse approach to the PC-prior a hierarchical prior similar in spirit to a spike and slab, where one first selects the number of nonzero components and then for these components only assumes (1). But it seems wrong to us to assume that the original problem translates into imposing that the D_{ii}'s have to be i.i.d. *sparse*. A more effective approach would be to consider

$$\mathbf{y} = (y_1, \ldots, y_n) \sim \pi(\mathbf{y}|\beta), \quad \beta \sim \mathcal{N}_p\big(0, \tau^{-1} D\big),$$
$$D = \operatorname{diag}(D_{11}, \ldots, D_{pp}),$$

where sparsity is expressed through $\tau \sim \pi_\tau$, as it then becomes a global notion. This modelling would have been much closer in spirit to the horseshoe prior approach, but with a different prior on τ implied by the PC approach and, being univariate, less problematic in the context of PC priors.

> "Occam's razor is the principle of parsimony, for which simpler model formulations should be preferred until there is enough support for a more complex model."

Assuming a base model has been constructed, we are definitely supportive of the idea of putting a prior on the distance from the base! This is even more the case because this concept is parameterisation invariant, at least at the hyperparameter level. Furthermore, it potentially associate a definitive meaning with the continuously invoked Occam's razor, even though we feel we could easily live without this constant reference to a vague notion proposed by a medieval monk in England. However, unless the hyperparameter ξ is one-dimensional, this approach fails to define a prior on ξ *per se*, which implies making further choices for the reference prior modelling.

We still wonder as to the particular Kullback–Leibler divergence chosen by the authors, $\mathrm{KLD}(\pi(\cdot|\xi) \parallel \pi(\cdot|\xi_0))$ as opposed to $\mathrm{KLD}(\pi(\cdot|\xi_0) \parallel \pi(\cdot|\xi))$. In terms of interpretation, it is not obvious to us that one conveys a better notion of complexity than the other. However, looking at the various examples in the paper, we realised that had the second choice been made the scaling argument in Section 3.3 would have failed and the PC prior would then be left undefined. Indeed, in this case, the Kullback–Leibler divergence between the distributions $\mathcal{N}(0, \varepsilon^2)$ and $\mathcal{N}(0, \sigma^2)$ is given by

$$\frac{1}{2}\left(\frac{\varepsilon^2}{\sigma^2} - 1 - \log\left(\frac{\varepsilon^2}{\sigma^2}\right)\right) \approx \log\left(\frac{\varepsilon}{\sigma}\right)$$

when $\varepsilon \approx 0$. Is there an explanation as to why one divergence is better than the other explanation that would be more illuminating than the mere observation than in the Gaussian case one works and the other fails?

We furthermore like equation (3.1) as this equation shows how the base constraint takes one away from Jeffreys' priors. However, equation (3.1) does not seem correct outside the unidimensional and bijective case, due to the differentiation term. Omitting this undefined Jacobian, one can opt for a uniform prior on the Kullback spheres of radius $d(\xi)$. The main part of the paper conveys a feeling of uni-dimensionality and we were

eager to see how it extends to models with many hyperparameters, which happens in Section 6. Similarly, if ξ was the mean of the Gaussian random variable and $\xi = 0$ the base model, then $\xi \to d(\xi)$ is not bijective although symmetry arguments suggest that it is enough to define a prior on $|\xi|$.

There is also a potential difficulty in the feature that $d(\xi)$ cannot be computed in a general setting. [Assuming that $d(\xi)$ has a nonvanishing Jacobian as on page 19 sounds rather unrealistic.] In Section 6 again, handling reference priors on correlation matrices $\mathbf{R}$ is a major endeavour. This should produce a steady flow of followers, even though it is certainly easier than contemplating the corresponding prior on a covariance matrix (Barnard, McCulloch and Meng, 2000).

> "The current practice of prior specification is, to be honest, not in a good shape. While there has been a strong growth of Bayesian analysis in science, the research field of "practical prior specification" has been left behind." (p. 23)

There still are (numerous) quantities to specify and calibrate in PC priors, a feature that may actually be deemed a good thing by Bayesians (and even by some modellers). But overall we think this paper and its central message constitute a terrific step for Bayesian analysis and not solely for its foundations, provided a more directive approach is adopted.

3. PC DIFFICULTIES

A first point in the delicate implementation of the PC principle is that those PC priors rely on several choices made in the ordering of the relevance, complexity, nuisance level and so on, of the parameters, in that way quite similar to reference priors (Berger, Bernardo and Sun, 2009). While the first author also wrote a paper on Russian roulette (Lyne et al., 2015), there is a further "Russian doll" principle at work behind (or within) PC priors. Each shell of the Russian doll corresponds to a further level of complexity whose order need be decided by the modeller. This does not sound to be a very realistic assumption in a hierarchical model with several types of parameters having only local meaning.

A second point is that the construction of those PC priors reflects another Russian doll structure, namely one of embedded models, hence would and should lead to a natural multiple testing methodology. However, the first author of the paper clearly rejected this notion during his talk at the ISBA 2016 conference, by being opposed to testing *per se*.

4. FURTHER REMARKS

> "We do not know precisely the thinking that underscores the choice of prior, but we do know that they have been hugely influential. This is not a satisfactory state of affairs."

The paper repeats a well-known meme that computation killed the Bayesian spirit. Plus this less common notion that priors should not be influential. Why not? If (and since) there is no such thing as a noninformative prior, all priors are influential and produce different inference outputs. What matters in the (user's) end is to provide a way to calibrate this output.

> "While reference priors have been successfully used for classical models, they have a less triumphant history for hierarchical models."

This argument applies to any setting in the sense that reference priors do require some ordering in the importance or relevance of the parameters. If there is subjectivity at this level, this also reflects on the ill-defined nature of some components of a hierarchical model. Especially when considering models with latent variables.

> "To date, there has been no attempt to construct a general method for specifying weakly informative priors."

We have no qualm about the section on *ad hoc* priors since we agree with the assessment that mimicking an earlier study of a similar problem brings no justification to another use. We neither oppose the weakly informative section, where we agree that the lack of general principles makes the argument hard to sustain.

The base model concept is once more appealing as a concept, reminding us of exponentially tilted models, although it is hard to see how to face the huge arbitrariness in setting base and extension. In that respect, Definition 1 is not particularly helpful. And Informal Definition 1 makes things worse. Setting a base is clearly a subjective or personal prior choice that should be acknowledged as such. About the desirable conditions, to which the authors are welcome and free to set as they wish, *D1* is not pertinent until the authors define the very notion of *noninformative*. This, of course, sounds like a circular argument. Desideratum *D2* makes sense only provided the sampling model allows for a division of the parameters into blocks. And *D3* is also worth considering, if pretty vague. Desideratum *D4* should first specify what an *over-parameterised* model means,

while *D5* assumes identifiability is itself identifiable, which is not always the case (although it may be the case for additive models). Desiderata *D6* and *D7* seem to proceed from common sense, while *D8* is not especially constrictive. Overall, we find it fairly hard to build a theory around such vague principles, as they are too far from methodology, which brings back the earlier comments about the very notion of Occam's razor.

The principles set in Section 3 are making perfect sense, although we stress again that they do require a fair amount of calibration. And we feel that the debate about using Cauchy versus Student's t priors sounds like a bit of an old saw.

> "PC priors are not built to be hypothesis testing priors and we do not recommend their direct use as such."

There still is this feeling of a hypothesis test reformulation that occurs when considering the base model like the null. See also the link with Johnson and Rossell (2010) nonlocal priors. We forgot about Bernardo's (2011) reformulation, which is under-exploited.

> "PC priors are defined on individual components. This distinguishes PC priors, from reference priors, in which the priors depend on the global model structure. This global dependence is required to ensure a proper posterior. However, the modern applied Bayesian is far more likely to approach their modelling using a componentwise and often additive approach."

Without here attempting a defence of reference priors (Bernardo and Smith, 1994) or of other noninformative construction techniques, we have issues with this separation of the parameter in independent components. This feature remains an assumption and as such it may be a poor choice. It is also hard to think of components of the parameter as being meaningful by themselves. Thus, to contemplate building extended priors on some components irrespective of what happens to other parts of the parameter somewhat defies reason. The remark on the PC prior construction in the context of sparse models, as presented above, rules out to some extent the independence construction advocated by the authors.

> "Close to the base model, the PC prior is a tilted Jeffreys' prior for $\pi(x|\xi)$, where the amount of tilting is determined by the distance on the Riemannian manifold to the base model scaled by the parameter λ."

It is not clear to us that the approximation proposed

$$\pi(\xi) = I(\xi)^{1/2} \exp\big(-\lambda m(\xi)\big) + \cdots \tag{2}$$

is actually sharper than

$$\pi(\xi) = I(0)^{1/2} \lambda \exp\big(-\lambda \sqrt{I(0)}\xi\big) + \cdots \tag{3}$$

which is a direct consequence of the first approximation provided by the authors

$$\mathrm{KLD}\big(\pi(x|\xi) \parallel \pi(x|\xi = 0)\big) = \frac{I(0)\xi^2}{2} + \cdots .$$

Then locally the PC prior behaves like an exponential prior with parameter $\lambda\sqrt{I(0)}$. This representation is less sophisticated than the authors' presentation, but it abstains from conveying the wrong notion that it locally resembles Jeffreys' prior. Indeed, it seems to us that the PC prior shares no common feature with Jeffreys' priors, neither locally nor globally.

5. CONCLUSION

> "We still have to work them out on a case by case basis."

While the authors have uncovered several interesting and new avenues for exploring prior specification, we want to signal yet another avenue, associated with the notion of Bayesian robustness, as proposed by Watson and Holmes (2016) that we recently discussed in this journal (Robert and Rousseau, 2016).

In fine, we congratulate the authors for this radical proposal that has the merit of defining a natural collection of priors, while integrating the constraints of prior robustness. Formalising this most important aspect of Bayesian modelling is absolutely essential for methodological and practical purposes, even if the current proposal is unlikely to reach most practitioners. We do acknowledge that the proposals made in the paper are currently exploratory, rather than directive. Again, we stress that this proposal represents an important step in the rational (if not objective) construction of reference or weakly informative priors (Gelman et al., 2014).

ACKNOWLEDGEMENTS

A component of those comments appeared on the first author's blog in a preliminary format. Discussions with Dan Simpson and Havard Rue over several versions of the paper were most illuminating, starting at a 2014 workshop in Banff and continuing in Warwick and Valencià for O-Bayes 2015. We are grateful to James Watson for sharing with us his impressions on the paper.

REFERENCES

BARNARD, J., MCCULLOCH, R. and MENG, X.-L. (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statist. Sinica* **10** 1281–1311. MR1804544

BERGER, J. O., BERNARDO, J. M. and SUN, D. (2009). Natural induction: An objective Bayesian approach. *Rev. R. Acad. Cienc. Exactas Fís. Nat., Ser. A Mat.* **103** 125–135. MR2535278

BERNARDO, J.-M. and SMITH, A. F. M. (1994). *Bayesian Theory*. Wiley, Chichester. MR1274699

BOCHKINA, N. A. and GREEN, P. J. (2014). The Bernstein–von Mises theorem and nonregular models. *Ann. Statist.* **42** 1850–1878. MR3262470

GELMAN, A., CARLIN, J. B., STERN, H. S., DUNSON, D. B., VEHTARI, A. and RUBIN, D. B. (2014). *Bayesian Data Analysis*, 3rd ed. CRC Press, Boca Raton, FL. MR3235677

JOHNSON, V. and ROSSELL, D. (2010). On the use of non-local prior densities in Bayesian hypothesis tests. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **72** 143–170.

KASS, R. and WASSERMAN, L. (1996). Formal rules of selecting prior distributions: A review and annotated bibliography. *J. Amer. Statist. Assoc.* **91** 1343–1370.

LISEO, B. (2005). The elimination of nuisance parameters. In *Handbook of Statistics* 25 (D. Dey and C. Rao, eds.) 193–219. Elsevier/North-Holland, Amsterdam. MR2490527

LYNE, A.-M., GIROLAMI, M., ATCHADÉ, Y., STRATHMANN, H. and SIMPSON, D. (2015). On Russian roulette estimates for Bayesian inference with doubly-intractable likelihoods. *Statist. Sci.* **30** 443–467. MR3432836

MARTIN, R. and LIU, C. (2015). Marginal inferential models: Prior-free probabilistic inference on interest parameters. *J. Amer. Statist. Assoc.* **110** 1621–1631.

ROBERT, C. (2001). *The Bayesian Choice*, 2nd ed. Springer, New York.

SEAMAN, J. W. III, SEAMAN, J. W. JR. and STAMEY, J. D. (2012). Hidden dangers of specifying noninformative priors. *Amer. Statist.* **66** 77–84. MR2968002

WATSON, J. and HOLMES, C. (2016). Approximate models and robust decisions. *Statist. Sci.* **31** 465–489. MR3598725

ZABELL, S. L. (1992). R. A. Fisher and the fiducial argument. *Statist. Sci.* **7** 369–387. MR1181418