

Comment on Article by Müller and Mitra

Peter D. Hoff *

Abstract. Due to their great flexibility, nonparametric Bayes methods have proven to be a valuable tool for discovering complicated patterns in data. The term “nonparametric Bayes” suggests that these methods inherit model-free operating characteristics of classical nonparametric methods, as well as coherent uncertainty assessments provided by Bayesian procedures. However, as the authors say in the conclusion to their article, nonparametric Bayesian methods may be more aptly described as “massively parametric.” Furthermore, I argue that many of the default nonparametric Bayes procedures are only Bayesian in the weakest sense of the term, and cannot be assumed to provide honest assessments of uncertainty merely because they carry the Bayesian label. However useful such procedures may be, we should be cautious about advertising default nonparametric Bayes procedures as either being “assumption free” or providing descriptions of our uncertainty. If we want our nonparametric Bayes procedures to have a Bayesian interpretation, we should modify default NP Bayes methods to accommodate real prior information, or at the very least, carefully evaluate the effects of hyperparameters on posterior quantities of interest.

Keywords: marginal likelihood, model misspecification, prior specification, sandwich estimation

1 Parameteric and nonparametric approaches

Historically, a standard justification of Bayesian methods has been that they provide an internally consistent approach to updating information: If $\mathcal{P}_\Theta = \{p(y|\theta) : \theta \in \Theta\}$ expresses our beliefs about Y given θ , and $\pi(\theta)$ expresses our beliefs about θ , then $\pi(\theta|y) \propto \pi(\theta)p(y|\theta)$ expresses what we *should* believe about θ , having observed $Y = y$. From this subjective Bayesian point of view, for $\pi(\theta|y)$ to be of most use, both $p(y|\theta)$ and $\pi(\theta)$ should actually represent our beliefs, at least approximately. A criticism of parametric Bayesian methods is that commonly used models \mathcal{P}_Θ are often suspected of being wrong. Nonparametric Bayes methods appear to solve this problem by making \mathcal{P}_Θ so large that it includes essentially all relevant sampling distributions $p(y|\theta)$. The authors seem to suggest that NP Bayes methods therefore provide an “honest representation of uncertainties”. I would agree with this, to the extent that $\pi(\theta)$ actually represents prior beliefs.

How honest are parametric and nonparametric priors? Parametric priors are arguably inaccurate as they assign probability one to a simple parametric model. However, the great advantage of parametric Bayesian approaches is that they allow the prior to be specified in terms of parameters of interest, which often happen to be the parameters

*Departments of Statistics and Biostatistics, University of Washington, Seattle, WA.
www.stat.washington.edu/hoff

about which we have real prior information. As a very simple example, suppose we have a sample y_1, \dots, y_n of independent observations from a population for which we have prior information about the mean and variance. In this case, a limited form of subjective, robust Bayesian inference for the population mean θ can proceed via the posterior density obtained from a normal sampling model for y_1, \dots, y_n . While the likelihood may not be exactly correct, the resulting inferences for θ are robust to nonnormality, asymptotically correct and provide confidence intervals for which the asymptotic frequentist coverage equals the asymptotic Bayesian coverage. Perhaps most importantly, the inference is transparent: the approximate normality of \bar{y} is well understood, and the effect of the prior on the parameter is simple, especially if a conjugate prior is used. Even if the prior does not represent our actual prior information, at least we can understand what information it represents.

In contrast, I think it is safe to say that for most NP Bayes methods used in practice, the prior does not represent actual prior beliefs, even approximately. One difficulty is that standard NP Bayes priors include hyperparameters that directly control things that we are unlikely to have prior information about (the number of modes of a density) and only indirectly control things we might have information about (means, variances and correlations). For example, the choice of the hyperparameters in the prior for a Dirichlet process mixture model (DPMM) induces a prior on the mean and variance of the population, but the mapping from the hyperparameters to these induced priors can be very opaque (Yamato 1984; Lijoi and Regazzini 2004). Similarly, the Pólya tree priors discussed in Section 2.2 require the specification of a partition over the sample space, the choice of which will generally affect the posterior. The “solution” to this is the addition of a prior over the set of possible partitions. It is hard to imagine that such a prior represents actual prior information about the underlying population.

Do such complications warrant abandoning NP Bayes methods and using simpler parametric approaches? It may depend on the data analysis objectives. NP Bayes methods provide a flexible means of representing high-dimensional data structure. Overfitting is avoided by a regularizer (the prior) that has a probabilistic interpretation. These features make NP Bayes methods an attractive set of tools for such tasks as prediction and clustering. However, if the data analysis objective is to describe our posterior information about a parameter of interest, then the appropriateness of NP Bayes is less clear. Example 1 from Müller and Mitra (2013) is a situation where the use of an NP Bayes method may be obfuscating the sources of information about the parameter of interest, $F(0)$. If we are to take the likelihood at face value, then it seems the data have little to say about the value of $F(0)$: Writing $f_k = F(k)$ and letting n_k be the number of T -cell sequences for which k replicates were observed, the likelihood can be expressed as $p(\mathbf{y}|f_0, \dots, f_4) = \prod_{k=1}^4 (f_k/[1-f_0])^{n_k}$. How much information do the data provide about f_0 ? One way to evaluate this is to consider the profile likelihood function of f_0 . For every fixed value of f_0 the likelihood $p(\mathbf{y}|f_0, \dots, f_4) = \prod_{k=1}^4 (f_k/[1-f_0])^{n_k}$ is maximized in f_1, \dots, f_4 at $\hat{f}_k = (1-f_0)n_k/n$. This gives a constant profile likelihood function for f_0 , equal to $\prod (n_k/n)^{n_k}$ for every value of f_0 . From a Bayesian perspective,

the posterior distribution of f_0 can be expressed as $\pi(f_0|\mathbf{y}) \propto \pi(f_0)p(\mathbf{y}|f_0)$, where

$$p(\mathbf{y}|f_0) = \int p(\mathbf{y}|f_0, \dots, f_4)\pi(f_1, \dots, f_4|f_0) df_1 \cdots df_4.$$

The profile likelihood argument suggests that $p(\mathbf{y}|f_0)$ will be fairly flat as a function of f_0 , especially if the prior over f_1, \dots, f_4 is “diffuse,” in which case $\pi(f_0|\mathbf{y}) \approx \pi(f_0)$. In this case, an “honest assessment” of posterior uncertainty about f_0 requires only an honest prior for f_0 . Alternatively, if we believe there to be a relationship among f_0, \dots, f_4 beyond the fact that $f_0 + \dots + f_4 < 1$, then again an honest assessment of posterior uncertainty about f_0 for these data requires only an honest specification of one’s joint beliefs about the five numbers f_0, \dots, f_4 . In either case, the DPMM over $\{f_0, f_1, \dots\}$ seems like a very indirect and opaque way to specify a prior over the relevant parameters.

One could argue that the DPMM in this example is helping us estimate other aspects of the unknown frequencies, such as the relative frequencies, perhaps. However, even though the DPMM in this example is billed as a “nonparametric Bayes” procedure, it is not without strong modeling assumptions. Specifically, this DPMM assumes the true distribution is a mixture of Poisson distributions - a class of distributions that does not contain all discrete distributions.

2 Partial remedies for particular situations

Alternative likelihoods: Consider a model $\{p(y|f) : f \in \mathcal{F}\}$ where f is a high dimensional parameter (such as a regression function or density). In situations where primary interest is in a low-dimensional parameter $\theta = \theta(f)$, the difficulties of infinite-dimensional prior specification can sometimes be avoided by using a likelihood that involves only θ . For example, in many problems there exists a statistic $t(y)$ whose distribution depends only on θ , and not the high-dimensional parameter f . In such cases, the likelihood can be expressed as

$$p(y|f) = p(t(y)|\theta) \times p(y|t(y), \theta, f).$$

The need to specify a prior over f can be avoided by constructing a posterior distribution for θ based only on the marginal likelihood $p(t(y)|\theta)$, i.e. $\pi(\theta|t(y)) \propto \pi(\theta)p(t(y)|\theta)$. Estimates based on such a posterior distribution could be inefficient, as they ignore any additional information about θ in $p(y|t(y), \theta, f)$, but they do not require specification of a prior for the high-dimensional nuisance parameter f . A concrete example of such a procedure is given in Hoff (2007) in the context of a semiparametric copula model, in which θ represents the parameters in a parametric dependence model and f parameterizes a set of unknown infinite-dimensional univariate marginal distributions.

Many researchers have considered other alternative likelihoods for robust or “non-parametric” Bayesian inference (Efron (1993); Lazar (2003); Greco et al. (2008), to name a few). Asymptotically correct likelihoods for parameters of interest can even be derived from misspecified models: Very generally, the limiting distribution of the MLE

$\hat{\theta}$ in a misspecified model is asymptotically normal, so that

$$\sqrt{n}(\theta^* - \hat{\theta}) \overset{\sim}{\sim} N(0, V(\theta^*))$$

where θ^* is the “pseudotrue” parameter and $V(\theta^*)$ is the “sandwich” variance (Huber 1967). In many cases (such as in exponential family models) θ^* is a population moment, and possibly the parameter of interest about which we may have prior information. In this case, nonparametric Bayesian inference can be obtained by combining a prior on θ^* with the asymptotic normal distribution of $\hat{\theta}$ as a likelihood. Such “Bayesian sandwich” procedures have been considered Szpiro et al. (2010), Müller (2012) and Hoff and Wakefield (2012).

Marginally specified priors: Such reductions of the parameter space are not feasible in applications such as prediction, where the high- or infinite-dimensional parameter f is of primary interest. In such cases it is important from a Bayesian perspective that the prior for f reflects known information as much as possible. Realistically, a statistician is unlikely to have informed opinions about all aspects of a high-dimensional parameter f , but may have real information about a finite-dimensional functional $\theta = \theta(f)$, such as the population mean or variance. Recently, Kessler et al. (2012) have proposed an approach for incorporating prior information about θ into a default NP Bayes prior for f . Specifically, let π_0 be a prior for f that is chosen arbitrarily or for computational convenience. This prior induces a marginal prior on θ , say P_0 , that may not reflect actual prior information, as quantified by a distribution P_1 . To remedy this problem, first express the default prior π_0 as

$$\pi_0(f \in A) = \int \pi_0(f \in A|\theta)P_0(d\theta).$$

To obtain a prior π_1 over f with the desired marginal distribution P_1 , simply replace P_0 with P_1 in the above expression. The resulting prior on f then becomes

$$\pi_1(f \in A) = \int \pi_0(f \in A|\theta)P_1(d\theta).$$

Such a prior generally retains the good large-support properties of the default NP Bayes prior π_0 , but has an induced prior over θ that matches the actual prior information P_1 . Computation of the posterior under such a prior is also generally available if an MCMC algorithm exists for the default prior π_0 . In this case, posterior approximation under π_1 can be made with the addition of a Metropolis-Hastings step.

Noninformative priors: In the absence of prior information, NP Bayes practitioners may attempt to produce “diffuse” priors by adjusting the hyperparameters in some way. However, intuition about what parameters correspond to “noninformativeness” can be misleading, partly due to the terminology used in the NP Bayes literature. For example, NP Bayes researchers should make clear that the total mass parameter α in a DPMM controls much more than “the uncertainty of” the mixing measure.

As DPMM researchers know, this hyperparameter controls such things as the entropy of the resulting probability density, the number of modes, etc. As another example, practitioners sometimes select overdispersed base measures in DPMMs in the hope that these reduce the effect of the prior on the analysis. [Bush et al. \(2010\)](#) have shown that such attempts generally lead to an unreasonably small number of mixture components, and propose a more nuanced version of the total mass hyperparameter to achieve a type of “noninformative” NP Bayes analysis.

3 Conclusion

Standard data analysis procedures can generally be described as techniques that convert a large set of numbers (the data) into a smaller set of numbers (parameter estimates, standard errors, etc.). Ideally, such a procedure is statistically meaningful and reasonably transparent: meaningful in that it has some desirable property in an idealized situation, and transparent so that its behavior can be understood outside of the idealized situation. Conjugate Bayesian estimation in an exponential family model is a good example of a meaningful and transparent procedure: The properties of such procedures are well-understood in the subjective Bayes framework, in an asymptotic framework and even in settings where the model is misspecified. In many cases, incorrect parametric models can provide meaningful, transparent and accurate inference for certain population parameters of interest, if not for all aspects of the population.

NP Bayes procedures typically convert a small set of numbers (the data) into a much larger set of numbers (the posterior distribution of the infinite dimensional parameter). What meaning can such extrapolative procedures have? Asymptotic results assure us that many NP Bayes procedures converge to the truth as fast as other nonparametric procedures, and perhaps faster if the the prior is close to the truth in a topological sense (see, for example, [Ghosal \(2001\)](#)). It might even be the case that 95% posterior confidence intervals have approximate 95% frequentist coverage, as they often do in parametric models ([Severini 1991](#)).

How are NP Bayes procedures justified non-asymptotically? Small sample justifications of Bayesian procedures are often based on their optimality under a particular prior. In simple models, this justification is transparent, in that even if the prior doesn't represent one's actual beliefs, at least one understands what beliefs it represents. Prior specification for NP Bayes procedures are more opaque and harder to justify from a Bayesian perspective. Default prior distributions are not generally going to represent prior beliefs, making it difficult to interpret the corresponding posterior distributions as posterior beliefs (except perhaps asymptotically). In terms of transparency, it is certainly possible to gain a strong intuition for the effects of hyperparameters on posterior output. Yet expert nonparametric Bayesians frequently use terminology that may be misleading to less experienced NP Bayes practitioners: For example, referring to the mass parameter α in a DPMM as indexing uncertainty is an incomplete description at best. Referring to a DPMM as a method for “BNP inference” on a clustering overlooks, as the authors pointed out in Section 3.1, that the Pólya urn scheme is a very partic-

ular one-parameter partition model. Referring to a mixture of Poisson distributions as “nonparametric” may give the impression to the inexperienced reader that the resulting mixture model contains all discrete distributions.

Most importantly, a posterior distribution does not provide an honest assessment of uncertainty by virtue of being a posterior distribution. Such an assessment is obtained via either an honest prior or asymptotically. In the absence of an infinite sample size, considerable effort should be made to use a prior distribution that approximates as closely as possible any real prior information that is available. In the absence of prior information, a more complete (but tedious) description of uncertainty would include a sensitivity analysis over possible values of the hyperparameters.

I am certainly not arguing that such efforts regarding the prior be mandated for every application of an NP Bayes method. However, I feel that more effort in this direction is necessary if we want our posterior distributions to represent honest assessments of uncertainty.

References

- Bush, C., Lee, J., and MacEachern, S. (2010). “Minimally informative prior distributions for non-parametric Bayesian analysis.” *Journal of the Royal Statistical Society B*, 72(2): 253–268. 315
- Efron, B. (1993). “Bayes and likelihood calculations from confidence intervals.” *Biometrika*, 80(1): 3–26. 313
- Ghosal, S. (2001). “Convergence rates for density estimation with Bernstein polynomials.” *Annals of Statistics*, 29(5): 1264–1280. 315
- Greco, L., Racugno, W., and Ventura, L. (2008). “Robust likelihood functions in Bayesian inference.” *Journal of Statistical Planning and Inference*, 138(5): 1258–1270. 313
- Hoff, P. and Wakefield, J. (2012). “Bayesian sandwich posteriors for pseudo-true parameters.” Technical Report 605, Department of Statistics, University of Washington. 314
- Hoff, P. D. (2007). “Extending the rank likelihood for semiparametric copula estimation.” *Annals of Applied Statistics*, 1(1): 265–283. 313
- Huber, P. (1967). “The behavior of maximum likelihood estimators under non-standard conditions.” In LeCam, L. and Neyman, J. (eds.), *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, 221–233. University of California Press. 314
- Kessler, D., Hoff, P., and Dunson, D. (2012). “Marginally Specified Priors for Non-parametric Bayesian Estimation.” Technical Report 596, Department of Statistics, University of Washington. 314

- Lazar, N. A. (2003). “Bayesian empirical likelihood.” *Biometrika*, 90(2): 319–326. 313
- Lijoi, A. and Regazzini, E. (2004). “Means of a Dirichlet Process and Multiple Hypergeometric Functions.” *The Annals of Probability*, 32(2): 1469–1495. 312
- Müller, P. and Mitra, R. (2013). “Bayesian Nonparametric Inference - Why and How.” *Bayesian Analysis*. To appear. 312
- Müller, U. (2012). “Risk of Bayesian inference in misspecified models, and the sandwich covariance matrix.” *Unpublished Manuscript, Princeton University*. 314
- Severini, T. A. (1991). “On the relationship between Bayesian and non-Bayesian interval estimates.” *Journal of the Royal Statistical Society, B*, 53(3): 611–618. 315
- Szpiro, A. A., Rice, K. M., and Lumley, T. (2010). “Model-robust regression and a Bayesian “sandwich” estimator.” *Annals of Applied Statistics*, 4(4): 2099–2113. 314
- Yamato, H. (1984). “Characteristic functions of means of distributions chosen from a Dirichlet process.” *Annals of Probability*, 12(1): 262–267. 312

