On the marginal likelihood and cross-validation

E. Fong & C. C. Holmes

Department of Statistics, University of Oxford, OX1 3LB
edwin.fong@stats.ox.ac.uk & cholmes@stats.ox.ac.uk

September 24, 2019

Abstract

In Bayesian statistics, the marginal likelihood, also known as the evidence, is used to evaluate model fit as it quantifies the joint probability of the data under the prior. In contrast, non-Bayesian models are typically compared using cross-validation on held-out data, either through k-fold partitioning or leave-p-out subsampling. We show that the marginal likelihood is formally equivalent to exhaustive leave-p-out cross-validation averaged over all values of p and all held-out test sets when using the log posterior predictive probability as the scoring rule. Moreover, the log posterior predictive is the only coherent scoring rule under data exchangeability. This offers new insight into the marginal likelihood and cross-validation and highlights the potential sensitivity of the marginal likelihood to the choice of the prior. We suggest an alternative approach using cumulative cross-validation following a preparatory training phase. Our work has connections to prequential analysis and intrinsic Bayes factors but is motivated through a different course.

1 Introduction

Probabilistic model evaluation and selection is an important task in statistics and machine learning, particularly when multiple models are under initial consideration. In the non-Bayesian literature, models are typically compared using out-of-sample performance criteria such as cross-validation (Geisser and Eddy, 1979; Shao, 1993; Vehtari and Lampinen, 2002), or predictive information (Watanabe, 2010). Computing the leave-p-out cross-validation score requires n-choose-p test set evaluations for n data points, which in most cases is computationally unviable and hence approximations such as k-fold cross-validation are often used instead (Geisser, 1975). A survey is provided by Arlot and Celisse (2010), and a Bayesian perspective on cross-validation by Vehtari and Ojanen (2012); Gelman et al. (2014).

In Bayesian statistics, the marginal likelihood or model evidence is the natural measure of model fit. For a model $\mathcal{M}$ with likelihood function or sampling distribution $\{f_\theta(y) : \theta \in \Theta\}$ parameterized by θ, a prior $\pi(\theta)$, and observations $y_{1:n} \in \mathcal{Y}^n$, the marginal likelihood or the prior predictive is defined as

$$p_{\mathcal{M}}(y_{1:n}) = \int f_\theta(y_{1:n}) \, d\pi(\theta). \tag{1}$$

The marginal likelihood can be used to calculate the posterior probability of the model given the data, $p(\mathcal{M} \mid y_{1:n}) \propto p_{\mathcal{M}}(y_{1:n}) \, p(\mathcal{M})$, as it is the probability of the data being generated under the prior when the model is correctly specified (Robert, 2007, Chapter 7). The ratio of marginal likelihoods between models is known as the Bayes factor that quantifies the prior to posterior odds on observing

the data. The marginal likelihood can be difficult to compute if the likelihood is peaked with respect to the prior, although Monte Carlo solutions exist; see Robert and Wraith (2009) for a survey. Under vague priors, the marginal likelihood may also be highly sensitive to the prior dispersion even if the posterior is not; a well known example is Lindley's paradox (Lindley, 1957; O'Hagan and Forster, 2004; Robert, 2014). As a result, its approximations such as the Bayesian information criterion (Schwarz, 1978) or the deviance information criterion (Spiegelhalter et al., 2002) are widely used, see also Gelman et al. (2014).

For our work, it is useful to note from the property of probability distributions that the log marginal likelihood can be written as the sum of log conditionals,

$$\log p_{\mathcal{M}}(y_{1:n}) = \sum_{i=1}^{n} \log p_{\mathcal{M}}(y_i \mid y_{1:i-1}) \tag{2}$$

where $p_{\mathcal{M}}(y_i \mid y_{1:i-1}) = \int f_\theta(y_i) \, d\pi(\theta \mid y_{1:i-1})$ is the posterior predictive for $i > 1$, $p_{\mathcal{M}}(y_1 \mid y_{1:0})$ $= \int f_\theta(y_1) \, d\pi(\theta)$, and this representation is true for any permutation of the data indices.

While Bayesian inference formally assumes that the model space captures the truth, in the model misspecified or so called M-open scenario (Bernardo and Smith, 2009, Chapter 6) the log marginal likelihood can be simply interpreted as a predictive sequential, or prequential (Dawid, 1984), scoring rule of the form $S(y_{1:n}) = \sum_i s(y_i \mid y_{1:i-1})$ with score function $s(y_i \mid y_{1:i-1}) = \log p_{\mathcal{M}}(y_i \mid y_{1:i-1})$. This interpretation of the log marginal likelihood as a predictive score (Kass and Raftery, 1995; Gneiting and Raftery, 2007; Bernardo and Smith, 2009, Chapter 6) has resulted in alternative scoring functions for Bayesian model selection (Dawid and Musio, 2014, 2015; Watson and Holmes, 2016; Shao et al., 2019), and provides insight into the relationship between the marginal likelihood and posterior predictive methods (Vehtari and Ojanen, 2012). Key et al. (1999) considered cross-validation from an M-open perspective and introduced a mixture utility for model selection that trades off fidelity to data with predictive power.

2   Uniqueness of the marginal likelihood under coherent scoring

To begin, we prove that under an assumption of data exchangeability, the log posterior predictive is the only prequential scoring rule that guarantees coherent model evaluation. The coherence property under exchangeability, where the indices of the data points carry no information, refers to the principle that identical models on seeing the same data should be scored equally irrespective of data ordering.

In demonstrating the uniqueness of the log posterior predictive, it is useful to introduce the notion of a general Bayesian model (Bissiri et al., 2016), which is a framework for Bayesian updating without the requirement of a true model. Define a parameter of interest by

$$\theta_0 = \arg\min_{\theta} \int l(\theta, y) dF_0(y) \tag{3}$$

where $F_0(y)$ is the unknown true sampling distribution giving rise to the data, and $l : \Theta \times \mathcal{Y} \to [0, \infty)$ is a loss function linking an observation y to the parameter θ. Bissiri et al. (2016) argue that after observing $y_{1:n}$, a coherent update of beliefs about θ_0 from a prior $\pi_G(\theta)$ to the posterior $\pi_G(\theta \mid y_{1:n})$ exists and must take on the form

$$\pi_G(\theta \mid y_{1:n}) \propto \exp\left\{-wl(\theta, y_{1:n})\right\} \pi_G(\theta) \tag{4}$$

where $l(\theta, y_{1:n}) = \sum_i l(\theta, y_i)$ is an additive loss function and $w > 0$ is a loss scale parameter; see Holmes and Walker (2017); Lyddon et al. (2019) on the selection of w. For $w = 1$ and $l(\theta, y) =$

$-\log f_\theta(y)$, we obtain traditional Bayesian updating without assuming the model $f_\theta(y)$ is true for some value of θ. From (3), M-open Bayesian inference is simply targeting the value of θ that minimizes the Kullback-Leibler divergence between $dF_0(y)$ and $f_\theta(y)$. The form (4) is uniquely implied by the assumptions in Theorem 1 of Bissiri et al. (2016), and we now focus on the coherence property of the update rule. An update function $\psi\{l(\theta, y), \pi_G(\theta)\} = \pi_G(\theta \mid y)$ is coherent if, for some inputs $y_{1:2}$, it satisfies

$$\psi[l(\theta, y_2), \psi\{l(\theta, y_1), \pi_G(\theta)\}] = \psi\{l(\theta, y_1) + l(\theta, y_2), \pi_G(\theta)\}.$$

This coherence condition is natural under an assumption of exchangeability as we expect posterior inferences about θ_0 to be unchanged whether we observe $y_{1:2}$ in any order or all at once, as it is in traditional Bayesian updating.

We now extend this coherence condition to general Bayesian model choice, where the goal is to evaluate the fit of the observed data under the general Bayesian model class $\mathcal{M}_G = \{l(\theta, y) : \theta \in \Theta\}$ with a prior $\pi_G(\theta)$. We treat w as a parameter outside of the model specification, as there are principled methods to select it from the model, prior and data. We define the log posterior predictive score as

$$s_G(\tilde{y} \mid y_{1:n}) = \log \int g\{l(\theta, \tilde{y})\} d\pi_G(\theta \mid y_{1:n})$$

where $g : [0, \infty) \to [0, \infty)$ is a continuous monotonically decreasing scoring function that transforms $l(\theta, y)$ into a predictive score for a test point $\tilde{y}$. We define the cumulative prequential log score as

$$S_G(y_{1:n}) = \sum_{i=1}^{n} s_G(y_i \mid y_{1:i-1})$$

where $s_G(y_1 \mid y_{1:0}) = \log \int g\{l(\theta, y_1)\} d\pi_G(\theta)$. The cumulative prequential log score sums the log posterior predictive score of each consecutive data point in a prequential manner, where a large score indicates that the model is predicting well. An intuitive choice for the scoring function might be the negative loss $g(l) = -l$, but we will see that this violates coherency, as defined below.

Definition 1. *The model scoring function $g(l)$ is coherent if it satisfies*

$$\sum_{i=1}^{n} s_G(y_i \mid y_{1:i-1}) = \log \int g\{l(\theta, y_{1:n})\} d\pi_G(\theta) \tag{5}$$

for all Θ, $\pi(\theta)$ and $n > 0$, such that $S_G(y_{1:n})$ is invariant to the ordering or partitioning of the observations.

We now present our main result on the uniqueness of the choice of g.

Proposition 1. *If the model scoring function $g : [0, \infty) \to [0, \infty)$ is continuous, monotonically decreasing and coherent, then the unique choice of scoring rule $g(l)$ is*

$$g(l) = \exp(-wl)$$

where w is the loss-scale in the general Bayesian posterior.

Proof. The proof is given in the Supplementary Material. □

This holds irrespective of whether the model is true or not. More importantly for us is the corollary below.

Corollary 1. *The marginal likelihood is the unique coherent marginal score for Bayesian inference.*

Proof. Let $w = 1$ and $l(\theta, y) = -\log f_\theta(y)$, and hence $g\{l(\theta, y)\} = f_\theta(y)$. □

The marginal likelihood arises naturally as the unique prequential scoring rule under coherent belief updating in the Bayesian framework. The coherence of the marginal likelihood implies an invariance to the permutation of the observations $y_{1:n}$ under exchangeability, including independent and identically distributed data, a property that is not shared by other prequential scoring rules, such as Dawid and Musio (2014); Grünwald and van Ommen (2017); Shao et al. (2019).

3 The marginal likelihood and cross-validation

3.1 Equivalence of the marginal likelihood and cumulative cross-validation

The leave-p-out cross-validation score is defined as

$$S_{CV}(y_{1:n}; p) = \frac{1}{\binom{n}{p}} \sum_{t=1}^{\binom{n}{p}} \frac{1}{p} \sum_{j=1}^{p} s\left(\tilde{y}_j^{(t)} \mid y_{1:n-p}^{(t)}\right) \tag{6}$$

where $\tilde{y}_{1:p}^{(t)}$ denotes the tth of n-choose-p possible held-out test sets, with $y_{1:n-p}^{(t)}$ the corresponding training set, such that $y_{1:n} = \{\tilde{y}^{(t)}, y^{(t)}\}$, and S_{CV} records the average predictive score per datum. Although leave-one-out cross-validation is a popular choice, it was shown in Shao (1993) that it is asymptotically inconsistent for a linear model selection problem, and requires $(p/n) \to 1$ as $n \to \infty$ for consistency. We will not go into further detail here but instead refer the reader to Arlot and Celisse (2010). Selecting a larger p has the interpretation of penalizing complexity (Vehtari and Ojanen, 2012), as complex models will tend to over-fit to a small training set. However, the number of test set evaluations grows rapidly with p and hence k-fold cross-validation is often adopted for computational convenience.

From a Bayesian perspective it is natural to consider the log posterior predictive as the scoring function, $s(\tilde{y} \mid y) = \log \int f_\theta(\tilde{y}) d\pi(\theta \mid y)$, particularly as we have now shown that it is the only coherent scoring mechanism, which leads us to the following result.

Proposition 2. *The Bayesian marginal likelihood is equivalent to the cumulative leave-p-out cross-validation score using the log posterior predictive as the scoring rule, such that*

$$\log p_{\mathcal{M}}(y_{1:n}) = \sum_{p=1}^{n} S_{CV}(y_{1:n}; p) \tag{7}$$

with $s(\tilde{y}_j \mid y_{1:n-p}) = \log p_{\mathcal{M}}(\tilde{y}_j \mid y_{1:n-p}) = \log \int f_\theta(\tilde{y}_j) \, d\pi(\theta \mid y_{1:n-p})$.

Proof. This follows from the invariance of the marginal likelihood under arbitrary permutation of the sequence $y_{1:n}$ in (2). We provide a proof and an alternative proof by induction in the Supplementary Material. □

The Bayesian marginal likelihood is simply n times the average leave-p-out cross-validation score, $n \times (1/n) \sum_{p=1}^{n} S_{CV}(y_{1:n}; p)$, where the scaling by n is due to (6) being a per datum score. Bayesian models are evaluated through out-of-sample predictions on all $(2^n - 1)$ possible held-out test sets whereas cross-validation with fixed p only captures a snapshot of model performance. Evaluating the predictive performance on $(2^n - 1)$ test sets would appear intractable for most applications, but we see through (7) and (1) that it is computable as a single integral.

3.2 Sensitivity to the prior and preparatory training

The representation of the marginal likelihood as a cumulative cross-validation score (7) provides insight into the sensitivity to the prior. The last term in the right hand side of (7) involves no training data, $S_{CV}(y_{1:n}; n) = (1/n) \sum_{i=1}^{n} \log \int f_\theta(y_i) \, d\pi(\theta)$, which scores the model entirely on how well the analyst is able to specify the prior. In many situations, the analyst may not want this term to contribute to model evaluation. Moreover, there is tension between any desire to specify vague priors to safeguard their influence and the fact that diffuse priors can lead to an arbitrarily large and negative model score for real valued parameters from (7). It may seem inappropriate to penalize a model based on the subjective ability to specify the prior, or to compare models using a score that includes contributions from predictions made using only a handful of training points even with informative priors. For example, we see that 10% of terms contributing to the marginal likelihood come from out-of-sample predictions using, on average, less than 5% of available training data. This is related to the start-up problem in prequential analysis (Dawid, 1992).

A natural and obvious solution is to begin evaluating the model performance after a preparatory phase, for example using 10% or 50% of the data as preparatory training prior to testing. This leads to a Bayesian cumulative leave-P-out cross-validation score defined as

$$S_{CCV}(y_{1:n}; P) = \sum_{p=1}^{P} S_{CV}(y_{1:n}; p) \tag{8}$$

with a preparatory cross-validation score $S_{PCV}(y_{1:n}; P) = \sum_{p=P+1}^{n} S_{CV}(y_{1:n}; p)$, for $1 \leq P < n$. We suggest setting P to leave out $0.9n$, $0.5n$ or $\max(0.9n, n - 10d)$, where d is the total number of model parameters, as reasonable default choices, but clearly this is situation specific. One may be interested in reporting both S_{CCV} and S_{PCV}, as the latter can be regarded as an evaluation of the prior, but we suggest that only S_{CCV} is used for model evaluation from the arguments above. Although full coherency is now lost, we still have coherency conditioned on a preparatory training set, where permutation of the data within the training and test sets does not affect the score, and so we can write (8) as

$$S_{CCV}(y_{1:n}; P) = \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}} \left(\tilde{y}_{1:P}^{(t)} \mid y_{1:n-P}^{(t)} \right). \tag{9}$$

This equivalence is derived in the Supplementary Material in a similar fashion to Proposition 2. This has precisely the form of the the log geometric intrinsic Bayes factor of Berger and Pericchi (1996) but motivated by a different route. The intrinsic Bayes factor was developed in an objective Bayesian setting (Berger and Pericchi, 2001), where improper priors cause indeterminacies in the evaluation of the marginal likelihood. The intrinsic Bayes factor remedies this with a partition of the data into $y_{1:l}, y_{l+1:n}$, where $y_{1:l}$ is the minimum training sample used to convert an improper prior $\pi(\theta)$ into a proper prior $\pi(\theta \mid y_{1:l})$. In contrast, we set $n - P$ to provide preparatory training and $\pi(\theta)$ can be subjective. Moreover, in modern applications we often have $d \gg n$ where intrinsic Bayes factors cannot be applied in their original form.

We can approximate (9) through Monte Carlo where the training data sets $y_{1:n-P}^{(t)}$ are drawn uniformly at random, and for non-conjugate models the inner term must also be estimated, for example through

$$\hat{S}_{CCV}(y_{1:n}; P) = \frac{1}{T} \sum_{t=1}^{T} \log \left\{ \frac{1}{B} \sum_{b=1}^{B} f_{\theta_b^{(t)}} \left(\tilde{y}_{1:P}^{(t)} \right) \right\} \qquad (10)$$

where samples $\theta_b^{(t)} \sim \pi \left(\theta \mid y_{1:n-P}^{(t)} \right)$ are obtained via T Markov chain Monte Carlo samplers. If we assume that the number of samples B per chain is sufficiently large, then the variance of the estimate $\hat{S}_{CCV}$ is approximately of the form τ^2/T. However, fitting T models may be costly, but we can run the chains in parallel. To avoid the need for T Markov chain Monte Carlo chains in (10), we can instead take advantage of the fact that the partial posteriors for different training sets will be similar, and utilize importance sampling (Bhattacharya and Haslett, 2007; Vehtari et al., 2017) or sequential Monte Carlo (Bornn et al., 2010) to estimate the posterior predictives for computational savings. We provide further details on efficient computation of (10) in the Supplementary Material.

4 Illustration for the normal linear model

We illustrate the use of Bayesian cumulative cross-validation in a polynomial regression example, where the rth polynomial model is defined as

$$f_\theta(y \mid x, r) = \mathcal{N}\{y; \theta^{\mathrm{T}}\phi_r(x), \sigma^2\}, \quad \phi_r(x) = \begin{bmatrix} 1 & x & \ldots & x^{r-1} & x^r \end{bmatrix}^{\mathrm{T}}.$$

We observe the data $\{y_{1:n}, x_{1:n}\}$, and we place a fixed vague prior on the intercept term, $\theta_0 \sim \mathcal{N}(\theta_0; 0, 100^2)$, and $\theta_d \sim \mathcal{N}(\theta_d; 0, s^2)$ for $d \in \{1, \ldots, r\}$ on the remaining coefficients. In our example, we have $n = 100$ and the true model is $r = 1$, $\theta = \begin{bmatrix} 1 & 0.5 \end{bmatrix}^{\mathrm{T}}$ with known $\sigma^2 = 1$. For our prior, we vary the value of $s^2 \in \{10^{-1}, 10^0, 10^4\}$ to investigate the impact of the prior tails. For each prior setting, we calculate $\log p_{\mathcal{M}}(y_{1:n})$ and $S_{CCV}(y_{1:n}; P)$ for models $r \in \{0, 1, 2\}$. In this example, $\log p_{\mathcal{M}}(y_{1:n})$ is tractable, whereas S_{CCV} requires a Monte Carlo average over tractable log posterior predictives. We report the mean over 10 runs of estimating S_{CCV} with $T = 10^6$ random training/test splits. We calculate the Monte Carlo standard error over the 10 runs and report the maximum for each setting of P.

The results are shown in Table 1, where $\hat{S}_{CCV}$ is normalized to the same scale as $\log p_r(y_{1:n})$. Under the strong prior $s^2 = 10^{-1}$ and the moderate prior $s^2 = 10^0$, the marginal likelihood correctly identifies the true model, but when we increase s^2 to 10^4 it heavily over-penalizes the more complex models and prefers $r = 0$. In fact, the magnitude of the marginal likelihood and the discrepancy just described can be made arbitrarily large by simply increasing s^2, which should be guarded against when a modeller has weak prior beliefs. This issue is not observed with $\hat{S}_{CCV}$ for the values of P we consider. The vague prior does not impede the ability of $\hat{S}_{CCV}$ to correctly identify the true model $r = 1$ and the scores are stable within each column of P.

In the Supplementary Material, we present graphical tools for exploring the cumulative cross-validation and the effect of the choice of P on S_{CCV}. We provide an additional example using probit regression on the Pima Indian data set.

Table 1: Log marginal likelihoods and cumulative cross-validation scores for normal linear model

s^2	MODEL r	$\log p_r(y_{1:n})$	$\hat{S}_{CCV}(y_{1:n}; P) \times n/P$		
			$P = 0.9n$	$P = 0.5n$	$P = 0.1n$
10^{-1}	0	-158.82	-153.80	-153.21	-153.06
	1	-155.57	-150.39	-149.55	-149.27
	2	-156.12	-150.94	-149.81	-149.38
10^{0}	0	-158.82	-153.80	-153.21	-153.06
	1	-156.26	-150.77	-149.66	-149.34
	2	-157.80	-151.90	-150.04	-149.50
10^{4}	0	-158.82	-153.80	-153.21	-153.06
	1	-160.81	-150.91	-149.68	-149.35
	2	-166.93	-152.30	-150.08	-149.53
MAXIMUM STANDARD ERROR			0.002	0.008	0.023

5   Discussion

We have shown that for coherence, the unique scoring rule for Bayesian model evaluation in either M-open or M-closed is provided by the log posterior predictive probability, and that the marginal likelihood is equivalent to a cumulative cross-validation score over all training-test data partitions. The coherence flows from the fact that the scoring rule and the Bayesian update both use the same information, namely the likelihood function, which is appropriate as the alternative would be to learn and score under different criteria. If we are interested in an alternative loss function to the log likelihood, we advocate a general Bayesian update (Bissiri et al., 2016; Lyddon et al., 2019) that targets the parameters minimising the expected loss, with models evaluated using the corresponding coherent cumulative cross-validation score.

Acknowledgement

References

Arlot, S. and Celisse, A. (2010). A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4:40–79.

Berger, J. O. and Pericchi, L. R. (1996). The intrinsic Bayes factor for model selection and prediction. *Journal of the American Statistical Association*, 91(433):109–122.

Berger, J. O. and Pericchi, L. R. (2001). *Objective Bayesian Methods for Model Selection: Introduction and Comparison*, volume 38 of *Lecture Notes–Monograph Series*, pages 135–207. Institute of Mathematical Statistics, Beachwood, OH.

Bernardo, J. and Smith, A. (2009). *Bayesian Theory*. Wiley Series in Probability and Statistics. Wiley.

Bhattacharya, S. and Haslett, J. (2007). Importance re-sampling MCMC for cross-validation in inverse problems. *Bayesian Analysis*, 2(2):385–407.

Bissiri, P. G., Holmes, C. C., and Walker, S. G. (2016). A general framework for updating belief distributions. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 78(5):1103–1130.

Bornn, L., Doucet, A., and Gottardo, R. (2010). An efficient computational approach for prior sensitivity analysis and cross-validation. *Canadian Journal of Statistics*, 38(1):47–64.

Dawid, A. P. (1984). Present Position and Potential Developments: Some Personal Views: Statistical Theory: The Prequential Approach. *Journal of the Royal Statistical Society. Series A (General)*, 147(2):278.

Dawid, A. P. (1992). Prequential analysis, stochastic complexity and Bayesian inference. *Bayesian Statistics*, 4:109–125.

Dawid, A. P. and Musio, M. (2014). Theory and applications of proper scoring rules. *METRON*, 72(2):169–183.

Dawid, A. P. and Musio, M. (2015). Bayesian model selection based on proper scoring rules. *Bayesian Analysis*, 10(2):479–499.

Geisser, S. (1975). The predictive sample reuse method with applications. *Journal of the American Statistical Association*, 70(350):320–328.

Geisser, S. and Eddy, W. (1979). A predictive approach to model selection. *Journal of the American Statistical Association*, 74:153–160.

Gelman, A., Hwang, J., and Vehtari, A. (2014). Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24:997–1016.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.

Grünwald, P. and van Ommen, T. (2017). Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it. *Bayesian Analysis*, 12(4):1069–1103.

Holmes, C. C. and Walker, S. G. (2017). Assigning a value to a power likelihood in a general Bayesian model. *Biometrika*, 104(2):497–503.

Kass, R. E. and Raftery, A. E. (1995). Bayes factors. *Journal of the American Statistical Association*, 90(430):773–795.

Key, J., Pericchi, L., and Smith, A. (1999). Bayesian model choice: What and why? (with discussion). *Bayesian Statistics*, 5.

Lindley, D. V. (1957). A statistical paradox. *Biometrika*, 44(1-2):187–192.

Lyddon, S. P., Holmes, C. C., and Walker, S. G. (2019). General Bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106(2):465–478.

Marin, J.-M. and Robert, C. P. (2010). Importance sampling methods for Bayesian discrimination between embedded models. *Frontiers of Statistical Decision Making and Bayesian Analysis*, pages 513–527.

O'Hagan, A. and Forster, J. J. (2004). *Kendall's Advanced Theory of Statistics, volume 2B: Bayesian Inference*, volume 2. Arnold.

Rischard, M., Jacob, P. E., and Pillai, N. (2018). Unbiased estimation of log normalizing constants with applications to Bayesian cross-validation. *arXiv preprint arXiv:1810.01382*.

Robert, C. P. (2007). *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*. Springer, 2nd edition.

Robert, C. P. (2014). On the Jeffreys-Lindley paradox. *Philosophy of Science*, 81(2):216–232.

Robert, C. P. and Wraith, D. (2009). Computational methods for Bayesian model choice. In *AIP conference proceedings*, volume 1193, pages 251–262. AIP.

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2):461–464.

Shao, J. (1993). Linear model selection by cross-validation. *Journal of the American Statistical Association*, 88(422):486–494.

Shao, S., Jacob, P. E., Ding, J., and Tarokh, V. (2019). Bayesian model comparison with the Hyvärinen score: Computation and consistency. *Journal of the American Statistical Association*, pages 1–24.

Spiegelhalter, D. J., Best, N. G., Carlin, B. P., and Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(4):583–639.

Vehtari, A., Gelman, A., and Gabry, J. (2017). Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, 27(5):1413–1432.

Vehtari, A. and Lampinen, J. (2002). Bayesian model assessment and comparison using cross-validation predictive densities. *Neural Computation*, 14(10):2339–2468.

Vehtari, A. and Ojanen, J. (2012). A survey of Bayesian predictive methods for model assessment, selection and comparison. *Statistics Surveys*, 6:142–228.

Watanabe, S. (2010). Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory. *Journal of Machine Learning Research*, 11:3571–3594.

Watson, J. and Holmes, C. C. (2016). Approximate models and robust decisions. *Statistical Science*, 31(4):465–489.

A   Supplementary Material

A.1   Proof of Proposition 1

Proof. We look at the case where $\Theta = \{0, 1\}$, so the prior $\pi_G(\theta)$ is parametrized by $p \in [0, 1]$ with $\pi_G(\theta = 0) = p$. We let $n = 2$, denoting the observables as y_1, y_2. We further denote $l(0, y_1) = l_0$ and $l(1, y_1) = l_1$, and likewise $l(0, y_2) = h_0$ and $l(1, y_2) = h_1$. We write p_1 as the updated $\pi_G(\theta = 0 \mid y_1)$ obtained from the general Bayesian update (4). The function $g(l)$ must then satisfy

$$\{g(l_0)p + g(l_1)(1 - p)\} \{g(h_0)p_1 + g(h_1)(1 - p_1)\}$$
$$= \{g(l_0 + h_0)p + g(l_1 + h_1)(1 - p)\} \tag{A.1.1}$$

for all $0 \leq p \leq 1$ and for all $l_0, l_1, h_0, h_1 \in [0, \infty)$. If we let $p = 1$, then $p_1 = 1$, so this simplifies to

$$g(l_0)g(h_0) = g(l_0 + h_0).$$

As g is continuous and monotonically decreasing, to satisfy (5) it must take on the form

$$g(l) = \exp(-\lambda l) \tag{A.1.2}$$

for $\lambda \geq 0$. We now explicitly write out the form of p_1

$$p_1 = \frac{\exp(-wl_0)p}{\exp(-wl_0)p + \exp(-wl_1)(1 - p)} = \frac{\exp(-wl_0)p}{Z_1}. \tag{A.1.3}$$

If we plug (A.1.2), (A.1.3) into (A.1.1), we obtain

$$\{\exp(-\lambda l_0)p + \exp(-\lambda l_1)(1 - p)\} \{\exp(-\lambda h_0)\exp(-wl_0)p + \exp(-\lambda h_1)\exp(-wl_1)(1 - p)\}$$
$$= Z_1 \left[\exp\{-\lambda(l_0 + h_0)\}p + \exp\{-\lambda(l_1 + h_1)\}(1 - p)\right].$$

Expanding, cancelling terms, and simplifying we obtain

$$\exp(-\lambda l_1 - wl_0)\{\exp(-\lambda h_0) - \exp(-\lambda h_1)\}$$
$$= \exp(-\lambda l_0 - wl_1)\{\exp(-\lambda h_0) - \exp(-\lambda h_1)\}$$

and so we must have $\lambda = 0$ or $\lambda = w$, where only the latter solution is non-trivial. We have thus shown that for $n = 2, |\Theta| = 2$, the unique non-trivial solution to (5) is

$$g(l) = \exp(-wl). \tag{A.1.4}$$

The remainder of the proof involves showing that this choice of g satisfies (5) for all $n > 0$ and all Θ and $\pi(\theta)$. Subbing (A.1.4) into (5), we obtain

$$\prod_{i=1}^{n} \exp\{s_G(y_i \mid y_{1:i-1})\} = \prod_{i=1}^{n} \int \exp\{-wl(\theta, y_i)\} \frac{\exp\{-wl(\theta, y_{1:i-1})\} \, d\pi_G(\theta)}{\int \exp\{-wl(\theta', y_{1:i-1})\} \, d\pi_G(\theta')}$$
$$= \prod_{i=1}^{n} \frac{\int \exp\{-wl(\theta, y_{1:i})\} \, d\pi_G(\theta)}{\int \exp\{-wl(\theta', y_{1:i-1})\} \, d\pi_G(\theta')}$$
$$= \int \exp\{-wl(\theta, y_{1:n})\} \, d\pi_G(\theta)$$

where for convenience we denote $l(\theta, y_{1:0}) = 0$.   □

A.2 Proof of Proposition 2

Proof. Consider the $(n! \times n)$ matrix Z with elements $(Z)_{ti} = \log p_{\mathcal{M}}(y_i^{(t)} \mid y_{1:i-1}^{(t)})$, such that the tth row of Z records the prequential sequence of log posterior predictives under the tth of $n!$ permutations of $y_{1:n}$. By the property of conditional probabilities, we have that the row sums of Z are equal, $\sum_i(Z)_{ti} = \sum_i(Z)_{t'i}$ for all t, t', and hence

$$\log p_{\mathcal{M}}(y_{1:n}) = \frac{1}{n!} \sum_{t=1}^{n!} \sum_{i=1}^{n} (Z)_{ti} = \sum_{i=1}^{n} \frac{1}{n!} \sum_{t=1}^{n!} (Z)_{ti}.$$

Within each column of Z, the values $(Z)_{ti}$ are invariant to the permutation of $y_{1:i-1}$ in the preceding $i - 1$ columns under exchangeability. There are thus n-choose-$(i - 1)$ distinct training sets and $n - i + 1$ choices for y_i given the training set. For each column $i \in \{1, \ldots, n\}$, we can then write

$$\frac{1}{n!} \sum_{t=1}^{n!} (Z)_{ti} = \frac{1}{\binom{n}{i-1}} \sum_{t=1}^{\binom{n}{i-1}} \frac{1}{n-i+1} \sum_{j=1}^{n-i+1} s\left(\tilde{y}_j^{(t)} \mid y_{1:i-1}^{(t)}\right)$$

$$= S_{CV}(y_{1:n}; n - i + 1)$$

where $s\left(\tilde{y}_j^{(t)} \mid y_{1:i-1}^{(t)}\right) = \log p_{\mathcal{M}}\left(\tilde{y}_j^{(t)} \mid y_{1:i-1}^{(t)}\right)$. We have the result for $p = n - i + 1$. $\square$

A.3 Alternative proof of Proposition 2

To prove Proposition 2, we first begin by showing the following proposition.

Proposition 3. *For a preparatory cross-validation score, $S_{PCV}(y_{1:n}; P)$, defined as the sum of cross-validation terms from leave-$(P + 1)$-out to leave-n-out,*

$$S_{PCV}(y_{1:n}; P) = \sum_{p=P+1}^{n} S_{CV}(y_{1:n}; p),$$

we have the following equivalence relationship

$$S_{PCV}(y_{1:n}; P) = \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}}\left(y_{1:n-P}^{(t)}\right) \tag{A.3.1}$$

which states that S_{PCV} is the average log marginal likelihood over all choices of the training set.

Proof. To show this, we use a proof by induction. We see that (A.3.1) is trivially true for $P = n - 1$, as this is simply $S_{CV}(y_{1:n}; n)$. Assuming (A.3.1) holds for some $1 \le P \le n - 1$, we have

$$S_{PCV}(y_{1:n}; P - 1) = S_{PCV}(y_{1:n}; P) + S_{CV}(y_{1:n}; P)$$

$$= \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}}\left(y_{1:n-P}^{(t)}\right) + \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \frac{1}{P} \sum_{j=1}^{P} \log p_{\mathcal{M}}\left(\tilde{y}_j^{(t)} \mid y_{1:n-P}^{(t)}\right)$$

$$= \frac{1}{P\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \left\{ P \log p_{\mathcal{M}}\left(y_{1:n-P}^{(t)}\right) + \sum_{j=1}^{P} \log p_{\mathcal{M}}\left(\tilde{y}_j^{(t)} \mid y_{1:n-P}^{(t)}\right) \right\}.$$

11

From the properties of conditional probability, we can write

$$S_{PCV}(y_{1:n}; P-1) = \frac{1}{P\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \sum_{j=1}^{P} \log p_{\mathcal{M}}\left(\tilde{y}_j^{(t)}, y_{1:n-P}^{(t)}\right). \tag{A.3.2}$$

Again, the marginal likelihood is invariant to the permutation of the sequence under data exchangeability, so we have to consider the repetitions in the partitions $\tilde{y}_j^{(t)}, y_{1:n-P}^{(t)}$. For each of the n choose $(n-P+1)$ unordered sequences $y_{1:n-P+1}^{(t')}$, there are $(n-P+1)$ partitions into $\tilde{y}_j^{(t)}, y_{1:n-P}^{(t)}$, so there are $n-P+1$ repetitions of each unordered $y_{1:n-P+1}^{(t')}$ in (A.3.2). We can thus write

$$S_{PCV}(y_{1:n}; P-1) = \frac{(n-P+1)}{P\binom{n}{P}} \sum_{t'=1}^{\binom{n}{P-1}} \log p_{\mathcal{M}}\left(y_{1:n-P+1}^{(t')}\right)$$

$$= \frac{1}{\binom{n}{P-1}} \sum_{t'=1}^{\binom{n}{P-1}} \log p_{\mathcal{M}}\left(y_{1:n-P+1}^{(t')}\right)$$

and by induction we have (A.3.1). □

Proposition 2 then follows trivially by setting $P = 0$ in Proposition 3.

A.4 Derivation of S_{CCV} for Bayesian models

The following corollary follows easily from Propositions 2 and 3.

Corollary 2. *For the cumulative cross-validation score defined as*

$$S_{CCV}(y_{1:n}; P) = \sum_{p=1}^{P} S_{CV}(y_{1:n}; p), \tag{A.4.1}$$

we have the following equivalence relationship

$$S_{CCV}(y_{1:n}; P) = \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}}\left(\tilde{y}_{1:P}^{(t)} \mid y_{1:n-P}^{(t)}\right). \tag{A.4.2}$$

Proof. We note that $\log p_{\mathcal{M}}(y_{1:n}) = S_{CCV}(y_{1:n}; P) + S_{PCV}(y_{1:n}; P)$ from their definitions and Proposition 2. From the permutation invariance of the marginal likelihood, we can write

$$\log p_{\mathcal{M}}(y_{1:n}) = \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}}\left(\tilde{y}_{1:P}^{(t)}, y_{1:n-P}^{(t)}\right). \tag{A.4.3}$$

By subtracting (A.3.1) in Proposition 3 from (A.4.3) and regarding each term in the summation, we have

$$S_{CCV}(y_{1:n}; P) = \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \left\{\log p_{\mathcal{M}}\left(\tilde{y}_{1:P}^{(t)}, y_{1:n-P}^{(t)}\right) - \log p_{\mathcal{M}}\left(y_{1:n-P}^{(t)}\right)\right\}$$

$$= \frac{1}{\binom{n}{P}} \sum_{t=1}^{\binom{n}{P}} \log p_{\mathcal{M}}\left(\tilde{y}_{1:P}^{(t)} \mid y_{1:n-P}^{(t)}\right)$$

□

12

A.5   Computing S_{CCV}

We note that $\hat{S}_{CCV}$ in (10) is a biased estimate, and Rischard et al. (2018) provides unbiased estimators of $\log p_{\mathcal{M}}(\tilde{y}_{1:P} \mid y_{1:n-P})$ directly through unbiased Markov chain Monte Carlo and path sampling methods.

The arithmetic averaging over training/test splits $\hat{S}_{CCV}$ may also be inherently unstable, as demonstrated by the following example. Suppose that y is a binary random variable which takes on either 0 or 1 with equal probability, and we are attempting to estimate $S_{CCV}(y_{1:n}; n/2)$. For large n, it is likely that approximately half of the values in $y_{1:n}$ are equal to 0 and the other half to 1. There will thus exist a permutation of the sequence $y_{1:n}$ such that almost all the first $n/2$ values are equal to 0, with the remaining almost all equal to 1. The model will then be certain that $y = 0$ after observing the training set, and score the remaining $n/2$ points very poorly, giving a large negative log posterior predictive. This suggests that an arithmetic average may be unstable; the median or robust trimmed mean over permutations may be stabler alternatives.

The form in (A.4.2) relies on the conditional coherency of Bayesian updating and scoring. Without this, S_{CCV} still exists as defined in (A.4.1), and can be directly estimated for example through

$$\hat{S}_{CCV}(y_{1:n}; P) = \frac{P}{T} \sum_{t=1}^{T} \frac{1}{p^{(t)}} \sum_{j=1}^{p^{(t)}} s\left(\tilde{y}_j^{(t)} \mid y_{1:n-p^{(t)}}^{(t)}\right)$$

where $p^{(t)} \sim \mathcal{U}\{1, P\}$ and the training set $y_{1:n-p^{(t)}}^{(t)}$ is sampled uniformly at random conditioned on $p^{(t)}$. This facilities alternative choices for the belief updating model and $s(\tilde{y} \mid y)$.

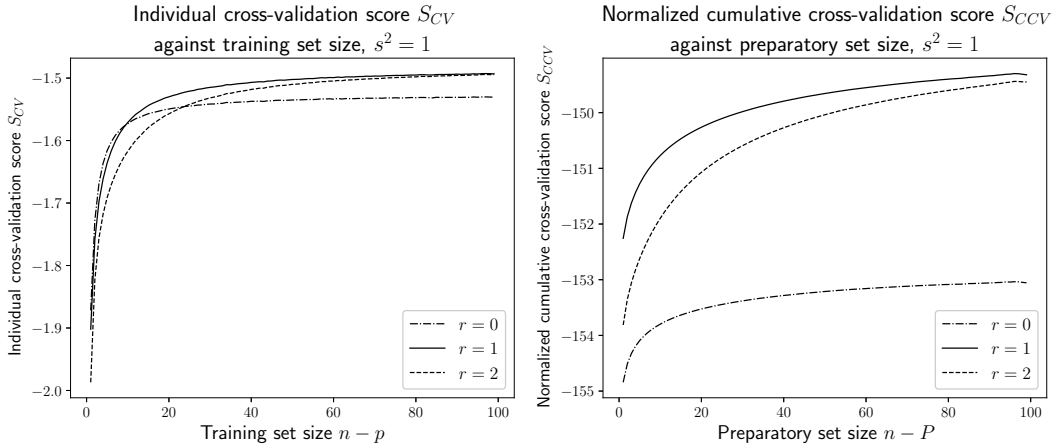## A.6   Visualization of cumulative cross-validation

Figure A.6.1: Leave-p-out cross-validation score $S_{CV}(y_{1:n}; p)$ against $n - p$ (left) and normalized cumulative cross-validation score $S_{CCV}(y_{1:n}; P) \times n/P$ against $n - P$ (right) for $s^2 = 1$ and $p, P \in \{1, \dots, 99\}$ in the polynomial regression example; the maximum standard error is 0.001 for S_{CV} and 0.005 for $\hat{S}_{CCV}$.

A visualization of the effects of the training/preparatory data size is shown in Figure A.6.1 for $s^2 = 1$ in the polynomial regression example. We omit $S_{CV}(y_{1:n}; n)$ and $S_{CCV}(y_{1:n}; n)$ for clarity of the plot, as both are significantly more negative than the other values. On the left we see that the individual cross-validation term $S_{CV}(y_{1:n}; p)$ prefers the simplest $r = 0$ model when the training set is very small as over-fitting is penalized, but as $n - p$ increases, the true $r = 1$ model overtakes it. The $r = 2$ model eventually overtakes the $r = 0$ model too, and we see the discrepancy between $r = 2$ and $r = 1$ decrease as over-fitting is penalized less and less. This latter effect is demonstrative of how leave-one-out cross-validation under-penalizes complex models as argued in Shao (1993), and why a value of $P > 1$ should be preferred. On the right, we observe a similar effect for the cumulative cross-validation score S_{CCV}, but the discrepancy between $r = 2$ and $r = 1$ remains more noticeable for moderate $n - P$ as a cumulative sum of S_{CV} terms is being taken.

A.7   Illustration for the probit model

To demonstrate the cumulative cross-validation score in an intractable example, we carry out model selection in the Pima Indian benchmark model with a probit model. We observe binary random variables $y_{1:n}$ with associated r-dimensional covariates $x_{1:n}$, and the probit model is defined as

$$f_\theta(y \mid x) = \{\Phi(\theta^{\mathrm{T}}\tilde{x})\}^y \{1 - \Phi(\theta^{\mathrm{T}}\tilde{x})\}^{1-y}$$

where Φ is the standard normal cumulative distribution function and $\tilde{x} = \begin{bmatrix} 1 & x^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}$. As suggested in Marin and Robert (2010), we elicit a g-prior $\pi(\theta) = \mathcal{N}\{\theta; 0_{r+1}, g(X^{\mathrm{T}}X)^{-1}\}$ where 0_{r+1} is a $r + 1$ vector of 0s and X is the n by $r + 1$ matrix with rows $\tilde{x}_i^{\mathrm{T}}$.

The dataset consists of $n = 332$ data points and we consider $r = 3$ covariates consisting of `glu`, `bp` and `ped`, which correspond to plasma glucose concentration from an oral glucose test, diastolic blood pressure and diabetes pedigree function respectively. We compare the full model $\mathcal{M}_0$: (`glu,bp,ped`) with $\mathcal{M}_1$: (`glu,bp`) through $\log p_{\mathcal{M}}(y_{1:n})$ and $S_{CCV}(y_{1:n}; P)$ to test for significance of `ped`. We standardize all covariates to have 0 mean and variance 1. We calculate $\log p_{\mathcal{M}}(y_{1:n})$ using importance sampling with a Gaussian proposal with 10^3 samples. The proposal mean is set to the maximum likelihood estimate of θ and proposal covariance to the estimated co-variance matrix of the maximum likelihood estimate as suggested in Marin and Robert (2010). For $S_{CCV}(y_{1:n}; P)$, we estimate each posterior predictive in (10) with the same importance sampling scheme where we temper the proposal such that its covariance matrix is divided by $(n - P)/n$. We also use 10^3 proposal samples and average over $T = 10^5$ random train/test splits. We carry out 10 runs of each and report the mean and maximum standard error as before.

We see in Table A.7.1 that for $g = n$, the simpler model with `ped` omitted performs worse for both scores, and there is thus strong evidence for `ped`. However, when we set $g = 10n$, we see that comparing models via the marginal likelihood suggests that `ped` is no longer significant, while the cumulative cross-validation score changes little with this increased variance of the prior. As a sanity check, we run a Gibbs sampler targeting $\pi(\theta \mid y_{1:n}, x_{1:n})$ for the two prior settings within the full model $\mathcal{M}_0$, and plot the marginal posterior of θ_{ped} in Figure A.7.1. For reference, the posterior means of $\theta_{\mathrm{glu}}, \theta_{\mathrm{bp}}$ are 0.70 and 0.12 respectively. The posteriors of θ_{ped} are indistinguishable for the two prior settings, with a significant mean for θ_{ped}. This agrees well with the cumulative cross-validation score $\hat{S}_{CCV}$ which is clearly robust to vague priors.

Table A.7.1: Log marginal likelihoods and cumulative cross-validation score for probit model

g	MODEL	$\log p_{\mathcal{M}}(y_{1:n})$	$\hat{S}_{CCV}(y_{1:n}; P) \times n/P$ $P = 0.9n$
n	(GLU, BP, PED)	-168.93	-165.87
	(GLU, BP)	-170.00	-167.37
$10n$	(GLU, BP, PED)	-173.10	-166.28
	(GLU, BP)	-173.05	-167.64
MAXIMUM STANDARD ERROR		0.004	0.02
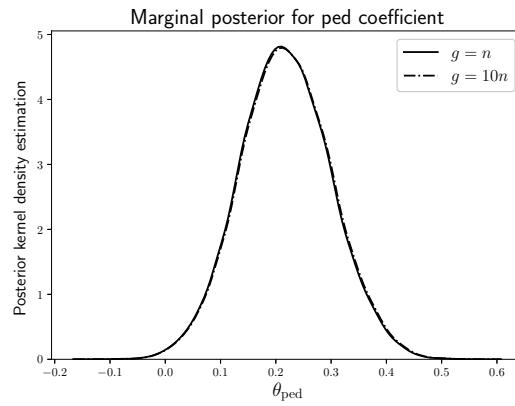
Figure A.7.1: Marginal posterior density plots for θ_{ped} for different prior scalings g.