

Using stacking to average Bayesian predictive distributions

Yuling Yao

*Department of Statistics,
Columbia University*

yy2619@columbia.edu

Aki Vehtari

*Helsinki Institute of Information Technology
Department of Computer Science, Aalto University*

Aki.Vehtari@aalto.fi

Daniel Simpson

*Department of Mathematical Sciences,
University of Bath*

dp.simpson@gmail.com

Andrew Gelman

*Department of Statistics and Department of Political Science,
Columbia University*

gelman@stat.columbia.edu

Abstract

The widely recommended procedure of Bayesian model averaging is flawed in the \mathcal{M} -open setting in which the true data-generating process is not one of the candidate models being fit. We take the idea of *stacking* from the point estimation literature and generalize to the combination of predictive distributions, extending the utility function to any proper scoring rule, using Pareto smoothed importance sampling to efficiently compute the required leave-one-out posterior distributions and regularization to get more stability. We compare stacking of predictive distributions to several alternatives: stacking of means, Bayesian model averaging (BMA), pseudo-BMA using AIC-type weighting, and a variant of pseudo-BMA that is stabilized using the Bayesian bootstrap. Based on simulations and real-data applications, we recommend stacking of predictive distributions, with BB-pseudo-BMA as an approximate alternative when computation cost is an issue.

Keywords: Bayesian model averaging, model combination, proper scoring rule, predictive distribution, stacking, Stan

1. Introduction

A general challenge in statistics is prediction in the presence of multiple candidate models or learning algorithms (M_1, \dots, M_K) . Model selection—picking one model that can give optimal performance for future data—can be unstable and wasteful of information (see, e.g., Piiroinen and Vehtari, 2015). An alternative is model averaging, which tries to find an optimal model combination in the space spanned by all individual models. In Bayesian context, the natural target for prediction is to find a predictive distribution that is close to the true data generating distribution (Gneiting and Raftery, 2007; Vehtari and Ojanen, 2012).

Ideally, we prefer to attack the Bayesian model combination problem via continuous model expansion—forming a larger bridging model that includes the separate models M_k as special cases (Gelman, 2004)—but in practice constructing such an expansion can require conceptual and computational effort, and so it makes sense to consider simpler tools that work with existing inferences from separately-fit models.

1.1 Bayesian model averaging

If the set of candidate models between them represents a full generative model, then the Bayesian solution is to simply average the separate models, weighing each by its marginal posterior probability. This is called *Bayesian model averaging* (BMA) and is optimal if the prior is correct, that is, if the method is evaluated based on its frequency properties evaluated over the joint prior distribution of the models and their internal parameters (Madigan et al., 1996; Hoeting et al., 1999). If $y = (y_1, \dots, y_n)$ represents the observed data, then the posterior distribution for any quantity of interest Δ is,

$$p(\Delta|y) = \sum_{k=1}^K p(\Delta|M_k, y)p(M_k|y).$$

In this expression, each model is weighted by its posterior probability,

$$p(M_k|y) = \frac{p(y|M_k)p(M_k)}{\sum_{k=1}^K p(y|M_k)p(M_k)},$$

and this expression depends crucially on the marginal likelihood under each model,

$$p(y|M_k) = \int p(y|\theta_k, M_k)p(\theta_k|M_k)d\theta_k.$$

BMA is appropriate if the true data generating process (DGP) can be identified exactly within the set of model classes M_k , that is the \mathcal{M} -closed case (Bernardo and Smith, 1994). In the \mathcal{M} -open case when the DGP does not exist in the model list, BMA will asymptotically select the one single model in the list that is closest in Kullback-Leibler (KL) divergence.

Furthermore, in BMA, the marginal likelihood depends sensitively on the specified prior $p(\theta_k|M_k)$ for each model. For example, consider a problem where a parameter has been assigned a normal prior distribution with center 0 and scale 10, and where its estimate is likely to be in the range $(-1, 1)$. The chosen prior is then essentially flat, as would also be the case if the scale were increased to 100 or 1000. But such a change would divide the posterior probability of the model by roughly a factor of 10 or 100.

1.2 Predictive accuracy

From another direction, one can simply aim to minimize out-of-sample predictive error, equivalently to maximize expected predictive accuracy. In this paper we propose a novel log score stacking method for combining Bayesian predictive distributions. As a side result we also propose a simple model weighting using Bayesian leave-one-out cross-validation.

1.3 Stacking

Stacking is a direct approach for averaging point estimates from multiple models. The idea originates with Wolpert (1992), and Breiman (1996) gives more details for stacking weights under different conditions. In supervised learning where the data are $((x_i, y_i), i = 1, \dots, n)$, and each model M_k has a parametric form $\hat{y}_k = f_k(x|\theta_k)$, stacking is done in two steps (Ting and Witten, 1999). In the first, baseline-level, step, each model is fitted separately and the leave-one-out (LOO) predictor $\hat{f}_k^{(-i)}(x_i) = E(y_i|\hat{\theta}_{k,y_{-i}}, M_k)$ is obtained for each model k and each data point i . Cross-validation or bootstrapping is used to avoid overfitting (LeBlanc and Tibshirani, 1996). In the second, meta-level, step, a weight for each model is obtained by minimizing the mean squared error, treating the leave-one-out predictors from the previous stage as covariates:

$$\hat{w} = \arg \min_w \sum_{i=1}^n \left(y_i - \sum_k w_k \hat{f}_k^{(-i)}(x_i) \right)^2. \quad (1)$$

Breiman (1996) notes that a positive constraint $w_k \geq 0, k = 1, \dots, K$, or a $K - 1$ simplex constraint: $w_k \geq 0, \sum_k^K w_k = 1$ enforces a solution. Better predictions may be attainable using regularization (Merz and Pazzani, 1999; Yang and Dunson, 2014). Finally, the point prediction for a new data point with feature vector \tilde{x} is

$$\hat{y} = \sum_{k=1}^K \hat{w}_k f_k(\tilde{x}|\hat{\theta}_{k,y_{1:n}}).$$

It is not surprising that stacking typically outperforms BMA when the criterion is mean squared predictive error (Clarke, 2003), because BMA is not optimized to this task. On the other hand, stacking is not widely used in Bayesian model combination because it only works with point estimates, not the entire posterior distribution (Hoeting et al., 1999).

Clyde and Iversen (2013) give a Bayesian interpretation for stacking by considering model combination as a decision problem when the true model M_t is not in the model list. If the decision is of the form $a(y, w) = \sum_{k=1}^K w_k \hat{y}_k$, then the expected utility under quadratic loss is,

$$E_{\tilde{y}}(u(\tilde{y}, a(y, w))|y) = - \int \|\tilde{y} - \sum_{k=1}^K w_k \hat{y}_k\|^2 p(\tilde{y}|y, M_t) d\tilde{y},$$

where \hat{y}_k is the predictor of new data \tilde{y} in model k . The stacking weights are the solution to the LOO estimator:

$$\hat{w} = \arg \min_w \frac{1}{n} \sum_{i=1}^n u(y_i, a(y_{-i}, w)),$$

where $a(y_{-i}, w) = \sum_{k=1}^K w_k E(y_i|y_{-i}, M_k)$.

The stacking predictor for new data \tilde{y} is $\sum_{k=1}^K \hat{w}_k \hat{y}_k$. The predictor in the k -th model \hat{y}_k can be either the plug-in estimator ,

$$E_k(\tilde{y}|\hat{\theta}_k, y) = \int \tilde{y} p(\tilde{y}|\hat{\theta}_{k,y}, M_k) d\tilde{y},$$

or the posterior mean,

$$E_k(\tilde{y}|y) = \int \tilde{y}p(\tilde{y}|\theta_k, M_k)p(\theta_k|y, M_k)d\tilde{y}d\theta_k.$$

Le and Clarke (2016) prove that the stacking solution is asymptotically the Bayes solution. With some mild conditions on the distribution, the following asymptotic relation holds:

$$\int l(\tilde{y}, a(y, w))p(\tilde{y}|y)d\tilde{y} - \frac{1}{n} \sum_{i=1}^n l(y_i, a(y_{-i}, w)) \xrightarrow{L_2} 0,$$

where l is the squared loss, $l(\tilde{y}, a) = (\tilde{y} - a)^2$. They also prove when the action is a predictive distribution $a(y_{-i}, w) = \sum_{k=1}^K w_k p(y_i|y_{-i}, M_k)$, then the asymptotic relation still holds for negative logarithm scoring rules.

However, most early literature limited stacking to averaging *point* predictions, rather than *predictive distributions*. In this paper, we extend stacking from minimizing the square error to maximizing scoring rules, hence make stacking applicable to combining a set of Bayesian posterior predictive distribution. We argue this is the appropriate version of Bayesian model averaging in the \mathcal{M} -open situation.

1.4 Fast leave-one-out cross-validation

Exact leave-one-out cross validation can be computationally costly. For example, in the econometric literature, Geweke and Amisano (2011, 2012) suggest averaging prediction models by maximizing predictive log score, only considering time series due to the computational cost of exact LOO for general data structures. In the present paper we demonstrate that Pareto-smoothed importance sampling leave-one-out cross-validation (PSIS-LOO) (Vehtari et al., 2016) is a practically efficient way to calculate the needed leave-one-out predictive densities $p(y_i|y_{-i}, M_k)$ to compute log score stacking weights.

1.5 Akaike weights and pseudo Bayesian model averaging

Leave-one-out cross-validation is related to various information criteria (see, e.g. Vehtari and Ojanen, 2012). In case of maximum likelihood estimates, leave-one-out cross-validation is asymptotically equal to Akaike's information criterion (AIC) Stone (1977). Given $AIC = -2 \log(\text{maximum likelihood}) + 2(\text{number of parameters})$, Akaike (1978) proposed to use $\exp(-\frac{1}{2}AIC)$ for model weighting (see also Burnham and Anderson, 2002; Wagenmakers and Farrell, 2004). More recently we have seen also Watanabe-Akaike information criterion (WAIC) (Watanabe, 2010) and leave-one-out cross-validation estimates used to compute model weights following the idea of AIC weights.

In Bayesian setting Geisser and Eddy (1979; see also, Gelfand 1996) proposed pseudo Bayes factors where marginal likelihoods $p(y|M_k)$ are replaced with a product of Bayesian leave-one-out cross-validation predictive densities $\prod_{i=1}^n p(y_i|y_{-i}, M_k)$. Following the naming by Geisser and Eddy, we call AIC type weighting which uses Bayesian cross-validation predictive densities as pseudo Bayesian model averaging (Pseudo-BMA).

In this paper we show that the uncertainty in the future data distribution should be taken into account when computing such weights. We will propose a AIC type weighting

using the Bayesian bootstrap and the expected log predictive density (elpd), which we call Pseudo-BMA+ weighting. We show that although Pseudo-BMA+ weighting gives better results than regular BMA or Pseudo-BMA weighting (in \mathcal{M} -open setting), it is still inferior to the log score stacking. Due to its simplicity we use Pseudo-BMA+ weights as initial guess for optimization procedure in the log score stacking.

1.6 Other model weighting approaches

Besides BMA, stacking and AIC type weighting, some other methods have been introduced to combine Bayesian models. Gutiérrez-Peña and Walker (2005) consider using a nonparametric prior in the decision problem stated above. Essentially they are fitting a mixture model with a Dirichlet process, yielding a posterior expected utility of,

$$U_n(w_k, \theta_k) = \sum_{i=1}^n \sum_{k=1}^K w_k f_k(y_i | \theta_k).$$

They then solve for the optimal weights $\hat{w}_k = \arg \max_{w_k, \theta_k} U_n(w_k, \theta_k)$.

Li and Dunson (2016) propose model averaging using weights based on divergences from a reference model in \mathcal{M} -completed setting. If the true model is known to be f^* , then an AIC type (or Boltzmann-Gibbs type) weight can be defined as,

$$w_{k|M} = \frac{\exp(-n\text{KL}(f^*, f_k))}{\sum_{j=1}^K \exp(-n\text{KL}(f^*, f_j))}. \quad (2)$$

The true model can be approximated with a reference model $f_0(\cdot | \theta_0)$ using nonparametric methods like Gaussian process or Dirichlet process, and $\text{KL}(f^*, f_j)$ can be estimated by its posterior mean,

$$\widetilde{\text{KL}}_1(f_0, f_k) = \int \int \text{KL}(f_0(\cdot | \theta_0), f_k(\cdot | \theta_k)) p(\theta_k | y) p(\theta_0 | y) d\theta_k d\theta_0,$$

or by the Kullback-Leibler divergence for posterior predictive distributions,

$$\widetilde{\text{KL}}_2(f_0, f_k) = \text{KL}\left(\int f_0(\cdot | \theta_0) p(\theta_0 | y) d\theta_0, \int f_k(\cdot | \theta_k) p(\theta_k | y) d\theta_k\right).$$

Here, $\widetilde{\text{KL}}_1$ corresponds to Gibbs utility, which can be criticized for not using the posterior predictive distributions (Vehtari and Ojanen, 2012), although asymptotically the two utilities are identical, and $\widetilde{\text{KL}}_1$ is often computationally simpler than $\widetilde{\text{KL}}_2$. As the entropy of the reference model is constant, the corresponding terms cancel out in KL terms in (2), leaving terms which are equal to the expected log predictive density (the expectation taken with respect to the reference model) and thus the method could be called Reference-Pseudo-BMA.

2. Theory and methods

We label classical stacking (1) as *stacking of means* because it combines the models by minimizing the mean squared error of the point estimate, which is the L_2 distance between the posterior mean and observed data. In general, we can use a proper scoring rule (or equivalently the underlying divergence) to compare distributions. After choosing that, stacking can be extended to combining the whole distribution.

2.1 Stacking using proper scoring rules

Adapting the notation of Gneiting and Raftery (2007), we label Y as the random variable on the sample space (Ω, \mathcal{A}) that can take values on $(-\infty, \infty)$. \mathcal{P} is a convex class of probability measure on Ω . Any member of \mathcal{P} is called a probabilistic forecast. A function $S : \mathcal{P} \times \Omega \rightarrow \bar{\mathbb{R}}$ defines a scoring rule if $S(P, \cdot)$ is \mathcal{P} quasi-integrable for all $P \in \mathcal{P}$. In the continuous case, distribution P is identified with density function p .

For two probabilistic forecasts P and Q , we write $S(P, Q) = \int S(P, \omega) dQ(\omega)$. A scoring rule S is called *proper* if $S(Q, Q) \geq S(P, Q)$ and *strictly proper* if the equation holds only when $P = Q$ almost surely. A proper scoring rule defines the divergence $d : \mathcal{P} \times \mathcal{P} \rightarrow (0, \infty)$ as $d(P, Q) = S(Q, Q) - S(P, Q)$. For continuous variables, some popularly used scoring rules include:

1. *Quadratic score*: $QS(p, y) = 2p(y) - \|p\|_2^2$ with the divergence $d(p, q) = \|p - q\|_2^2$.
2. *Logarithmic score*: $LogS(p, y) = \log(p(y))$ with $d(p, q) = KL(p, q)$ The logarithmic score is the only proper local score assuming regularity conditions.
3. *Continuous ranked score*: $CRPS(F, y) = -\int_{\mathbb{R}} (F(y') - 1(y' \geq y))^2 dy'$ with $d(F, G) = \int_{\mathbb{R}} (F(y) - G(y))^2 dy$, where F and G are the corresponding distribution function.
4. *Energy score*: $ES(P, y) = \frac{1}{2} E_P \|Y - Y'\|_2^\beta - E_P \|Y - y\|_2^\beta$, where Y and Y' are two independent random variables from distribution P . When $\beta = 2$, this becomes $ES(P, y) = -\|E_P(Y) - y\|^2$. The energy score is strictly proper when $\beta \in (0, 2)$ but not when $\beta = 2$.
5. *Scoring rules depending on first and second moments*: Examples include $S(P, y) = -\log \det(\Sigma_P) - (y - \mu_P)^T \Sigma_P^{-1} (y - \mu_P)$, where μ_P and Σ_P are the mean vector and covariance matrix of distribution P .

Now return to the problem of model combination. The observed data are $y = (y_1, \dots, y_n)$. For simplicity, we remove all covariates x in the notation. Suppose we have a set of probabilistic models $\mathcal{M} = (M_1, \dots, M_K)$; then the goal in stacking is to find an optimal super-model in the convex linear combination with the form $\mathcal{C} = (\sum_{k=1}^K w_k p(\cdot | M_k) | \sum_k w_k = 1, w_k \geq 0)$, such that its divergence to the true data generating model, denoted by $p_t(\cdot | y)$, is minimized:

$$\min_w d\left(\sum_k w_k p(\cdot | y, M_k), p_t(\cdot | y)\right).$$

Or equivalently maximize the scoring rule of the predictive distribution,

$$\max_w S\left(\sum_{k=1}^K w_k p(\cdot | y), p_t(\cdot | y)\right), \tag{3}$$

where $p(\tilde{y} | y, M_k)$ is the predictive density of \tilde{y} in model k :

$$p(\tilde{y} | y, M_k) = \int p(\tilde{y} | \theta_k, y, M_k) p(\theta_k | y, M_k) d\theta_k.$$

We label the leave- i -out predictive density in model k as,

$$\hat{p}_{k,-i}(y_i) = \int p(y_i|\theta_k, M_k)p(\theta|y_{-i}, M_k)d\theta_k.$$

where $y_{-i} = (y_1, \dots, y_{i-1}, y_{i+1}, \dots, y_n)$.

Then we define the stacking weights as the solution to the following optimization problem:

$$\max_w \frac{1}{n} \sum_{i=1}^n S\left(\sum_{k=1}^K w_k \hat{p}_{k,-i}, y_i\right), \quad s.t. \quad w_k \geq 0, \quad \sum_{k=1}^K w_k = 1. \quad (4)$$

Eventually, the combined estimation of the predictive density is

$$\hat{p}(\tilde{y}|y) = \sum_{k=1}^K \hat{w}_k p(\tilde{y}|y, M_k). \quad (5)$$

When using logarithmic score, we call this *stacking of predictive distributions*:

$$\max_w \frac{1}{n} \sum_{i=1}^n \log \sum_{k=1}^K w_k p(y_i|y_{-i}, M_k), \quad s.t. \quad w_k \geq 0, \quad \sum_{k=1}^K w_k = 1.$$

The choice of scoring rule can depend on the underlying application. Stacking of means (1) corresponds to the energy score with $\beta = 2$. The reasons why we prefer stacking of predictive distributions (corresponding to the logarithmic score) to stacking of means are: (i) the energy score with $\beta = 2$ is not a strictly proper scoring rule and can give rise to identification problems, and (ii) every proper local scoring rule is equivalent to the logarithmic score (Gneiting and Raftery, 2007).

2.2 Asymptotic behavior of stacking

The stacking estimate (3) finds the optimal predictive distribution within the convex set $\mathcal{C} = \{\sum_{k=1}^K w_k p(\cdot|M_k) \mid \sum_k w_k = 1, w_k \geq 0\}$, that is the closest to the data generating process with respect to the chosen scoring rule. This is different from Bayesian model averaging, which asymptotically with probability 1 will select a single model: the one that is closest in KL divergence to the true data generating process.

Solving for the stacking weights in (4) is an M-estimation problem. Under some mild conditions (Le and Clarke, 2016; Clyde and Iversen, 2013; Key et al., 1999), for either the logarithmic scoring rule or energy score (negative squared error) and a given set of weights $w_1 \dots w_k$, as sample size $n \rightarrow \infty$, the following asymptotic limit holds:

$$\frac{1}{n} \sum_{i=1}^n S\left(\sum_{k=1}^K w_k \hat{p}_{k,-i}, y_i\right) - E_{\tilde{y}|y} S\left(\sum_{k=1}^K w_k p(\tilde{y}|y, M_k), \tilde{y}\right) \xrightarrow{L_2} 0.$$

Thus the leave-one-out-score is a consistent estimator of the posterior score. In this sense, the stacking weights is the optimal combination weights asymptotically.

In terms of Vehtari and Ojanen (2012, Section 3.3), the proposed stacking with log score is M_* -optimal projection of the information in the actual belief model M_* to \hat{w} , where explicit specification of M_* is avoided by re-using data as a proxy for the predictive distribution of the actual belief model and w_k are the free parameters.

2.3 Pareto smoothed importance sampling

One challenge in calculating the stacking weights proposed in (4) is that we need the leave-one-out (LOO) predictive density,

$$p(y_i|y_{-i}, M_k) = \int p(y_i|\theta_k, M_k)p(\theta_k|y_{-i}, M_k)d\theta_k.$$

Exact LOO requires refitting each model n times. To avoid this onerous computation, we use the following approximate method. For the k -th model, we fit to all the data, obtaining simulation draws $\theta_k^s (s = 1, \dots, S)$ from the full posterior $p(\theta_k|y, M_k)$ and calculate

$$r_{i,k}^s = \frac{1}{p(y_i|\theta_k^s, M_k)} \propto \frac{p(\theta_k^s|y_{-i}, M_k)}{p(\theta_k^s|y, M_k)}.$$

The ratio $r_{i,k}^s$ has a density function in its denominator and can be unstable, a problem resolved using Pareto smoothed importance sampling (PSIS), which replaces the $r_{i,k}^s$ with smoothed importance weights $w_{i,k}^s$; for details, see Vehtari et al. (2016). Then the importance sampling is performed using

$$p(y_i|y_{-i}, M_k) \approx \frac{1}{\frac{1}{S} \sum_{s=1}^S w_{i,k}^s}.$$

When stacking on logarithmic score, we are combining each model's log predictive density. The PSIS estimate of the LOO expected log pointwise predictive density in the k -th model is,

$$\widehat{\text{elpd}}_{\text{loo}}^k = \sum_{i=1}^n \log \left(\frac{\sum_{s=1}^S w_{i,k}^s p(y_i|\theta_k^s, M_k)}{\sum_{s=1}^S w_{i,k}^s} \right).$$

The reliability of the PSIS approximation can be determined by the estimated shape parameter \hat{k} in the generalized Pareto distribution. For the left-out data points where $\hat{k} > 0.7$, Vehtari et al. (2016) suggest replacing the PSIS approximation of those problematic cases by the exact LOO or k -fold validation.

One potential drawback of LOO is the large variance when the sample size is small. We see in the simulation that when the ratio of relative sample size to effective number of parameters is small, the weighting can be unstable. How to adjust this small sample behavior is left for the future research.

2.4 Pseudo-BMA

In our paper, we also consider an AIC type weighting using leave-one-out cross-validation estimated expected log predictive densities. As mentioned in Section 1, these weights estimate the same quantities as (Li and Dunson, 2016) that used the divergence from the reference model based inference.

To maintain comparability with the given dataset and to get easier interpretation of the differences in scale of effective number of parameters, we define the expected log pointwise predictive density (elpd) for a new dataset \tilde{y} as a measure of predictive accuracy of a given model for the n data points taken one at a time (Gelman et al., 2014). For simplicity we suppress the dependence on k in the following notation:

$$\text{elpd} = \sum_{i=1}^n \int p_t(\tilde{y}_i) \log p(\tilde{y}_i|y) d\tilde{y}_i,$$

where $p_t(\tilde{y}_i)$ denotes the true distribution of \tilde{y}_i . We estimate this using LOO:

$$\text{elpd}_{\text{loo}} = \sum_{i=1}^n \log p(y_i|y_{-i}).$$

Then we define the Pseudo-BMA weighting for model k :

$$w_k = \frac{\exp(\widehat{\text{elpd}}_{\text{loo}}^k)}{\sum_{k=1}^K \exp(\widehat{\text{elpd}}_{\text{loo}}^k)}. \quad (6)$$

However, this estimation doesn't take into account the uncertainty resulted from having a finite number of proxy samples from the future data distribution. Taking into account the uncertainty would regularize the weights making them go further away from 0 and 1.

The computed estimate elpd is defined as the sum of n independent components so it is trivial to compute their standard errors by computing the standard deviation of the n pointwise values (Vehtari and Lampinen, 2002). Define

$$\widehat{\text{elpd}}_{\text{loo},i} = \log p(y_i|y_{-i}),$$

and then we can calculate

$$\text{se}(\widehat{\text{elpd}}_{\text{loo}}) = \sqrt{\sum_{i=1}^n (\widehat{\text{elpd}}_{\text{loo},i} - \widehat{\text{elpd}}_{\text{loo}}/n)^2}.$$

Simple modification of weights is to use the lognormal approximation:

$$w_k = \frac{\exp(\widehat{\text{elpd}}_{\text{loo}}^k - \frac{1}{2}\text{se}(\widehat{\text{elpd}}_{\text{loo}}^k))}{\sum_{k=1}^K \exp(\widehat{\text{elpd}}_{\text{loo}}^k - \frac{1}{2}\text{se}(\widehat{\text{elpd}}_{\text{loo}}^k))}.$$

Finally, the Bayesian bootstrap can be used to compute uncertainties related to LOO estimation (Vehtari and Lampinen, 2002). Bayesian bootstrap (BB) (Rubin, 1981) makes simple non-parametric approximation to the distribution of random variable. Having samples of z_1, \dots, z_n of a random variable Z , it is assumed that posterior probabilities for the z_i have the distribution Dirichlet(1, ..., 1) and values of Z that are not observed have zero posterior probability. Sampling from the uniform Dirichlet distribution gives BB samples from the distribution of the distribution of Z and thus samples of any parameter of this distribution can be obtained. As the distribution of $\widehat{\text{elpd}}_{\text{loo},i}^k$ is often highly skewed, BB is likely to work better than the Gaussian approximation.

In our model weighting, we can define,

$$z_i^k = \widehat{\text{elpd}}_{\text{loo},i}^k \quad i = 1, \dots, n.$$

We sample from the Dirichlet distribution vectors $(\alpha_{1,b}, \dots, \alpha_{n,b})$, where $b = 1, \dots, B$ and compute the weighted means,

$$\bar{z}_b^k = \sum_{i=1}^n w_{i,b} z_i^k, \quad k = 1, \dots, K.$$

Then the Bayesian bootstrap sample of the model weights is,

$$w_{k,b} = \frac{\exp(n\bar{z}_b^k)}{\sum_{k=1}^K \exp(n\bar{z}_b^k)},$$

and the final adjusted weights are,

$$w_k = \frac{1}{B} \sum_{b=1}^B w_{k,b}, \tag{7}$$

which we call Pseudo-BMA+ weights.

3. Simulation examples

In this section, we first illustrate the advantage of stacking of predictive distributions with a Gaussian mixture model. Then we compare stacking, BMA, Pseudo-BMA, Pseudo-BMA+ and other averaging methods through a series of linear regression simulation, where stacking gives the best performance in most cases. Finally we apply stacking to two real datasets, averaging multiple models so as to better explain US Senate voting or well-switching pattern in Bangladesh.

3.1 Gaussian mixture model

This simple example helps us understand how BMA and stacking behave differently. It also illustrates the importance of the choice of the scoring rules in combining distributions. Suppose the observed data $y = (y_i, i = 1, \dots, n)$ come iid from a normal distribution $N(3.4, 1)$, not known to the data analyst, and there are 8 candidate models, $N(1, 1)$, $N(2, 1), \dots, N(8, 1)$. This is an \mathcal{M} -open problem in that none of the candidates is the true model, and we have set the parameters so that the models are somewhat separate but not completely distinct in their predictive distributions.

For BMA with a uniform prior $\Pr(M_k) = \frac{1}{8}, k = 1, \dots, 8$, we can write the posterior distribution explicitly:

$$\hat{w}_k^{\text{BMA}} = P(M_k|y) = \frac{\exp(-\frac{1}{2} \sum_{i=1}^n (y_i - \mu_k)^2)}{\sum_{k'} \exp(-\frac{1}{2} \sum_{i=1}^n (y_i - \mu_{k'})^2)},$$

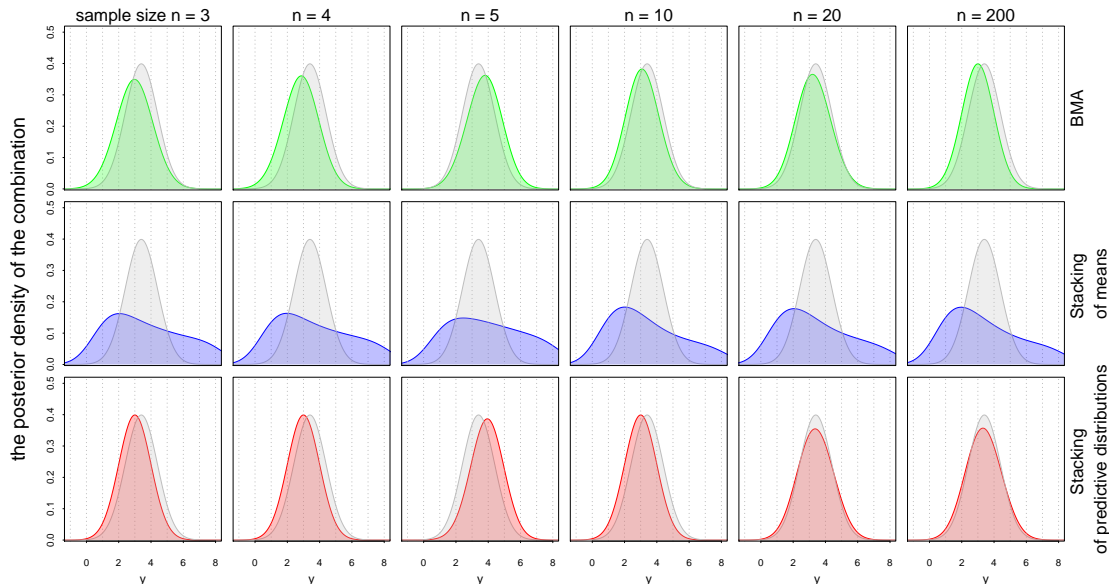


Figure 1: For the Gaussian mixture example, the predictive distribution $p(\tilde{y}|y)$ of BMA (green curve), stacking of means (blue) and stacking of predictive distributions (red). In each graph, the gray distribution represents the true model $N(3.4, 1)$. Stacking of means matches the first moment but can ignore the distribution. For this \mathcal{M} -open problem, stacking of predictive distributions outperforms BMA as sample size increases.

from which we see that $\hat{w}_3^{\text{BMA}} \xrightarrow{P} 1$ and $\hat{w}_k^{\text{BMA}} \xrightarrow{P} 0$ for $k \neq 3$ as sample size $n \rightarrow \infty$. Furthermore, for any given n ,

$$\begin{aligned} E_{y \sim N(\mu, 1)}[\hat{w}_k^{\text{BMA}}] &\propto E_y \left(\exp\left(-\frac{1}{2} \sum_{i=1}^n (y_i - \mu_k)^2\right) \right) \\ &\propto \left(\int_{-\infty}^{\infty} \exp\left(-\frac{1}{2} ((y - \mu_k)^2 + (y - \mu)^2)\right) dy \right)^n \\ &\propto \exp\left(-\frac{n(\mu_k - \mu)^2}{4}\right), \end{aligned}$$

where $\mu = 3.4$ and $\mu_k = k$ in this setting.

This example is simple in that there is no parameter to estimate within each of the models: $p(\tilde{y}|y, M_k) = p(\tilde{y}|M_k)$. Hence, in this case the weights from Pseudo-BMA and Pseudo-BMA+ are the same as the BMA weights, $\exp(\widehat{\text{elpd}}_{100}^k) / \sum_{k'} \exp(\widehat{\text{elpd}}_{100}^{k'})$.

For stacking, if we maximize the energy score (that is, minimizing the mean squared error of the point estimator), then the problem is nonidentifiable because the predictor of any data point in model k is always k . Meanwhile, the transformed parameter $\theta = \sum_{k=1}^8 w_k k$ is identifiable. With a uniform prior: $\sum_k w_k = 1, w_k \geq 0$, we can still make inference on the




Figure 2: (a) The left panel shows the expected log predictive density of the predictive distribution under BMA, stacking of means and stacking of predictive distributions. Stacking of predictive distributions performs best for moderate and large sample sizes. (b) The middle panel shows the main squared error treating the posterior mean of \hat{y} as a point estimation. Stacking of predictive distributions gives almost the same optimal mean squared error as stacking of means, both of which perform better than BMA. (c) The right panel shows the expected log predictive density of stacking and BMA when adding some more $N(4,1)$ models to the model list, where sample size is fixed to be 15. All average log scores and errors are calculated through 500 repeated simulation and 200 test data generating from the true distribution.

distribution of y by simulating the model,

$$y_i \sim N\left(\sum_{k=1}^8 w_k k, \sigma^2\right), \quad p(w_1, \dots, w_8, \sigma) = 1.$$

For stacking of predictive distributions, the weights are,

$$\hat{w} = \arg \max_w \sum_{i=1}^n \log \left(\sum_{k=1}^8 w_k p_k(y_i) \right) = \arg \max_w \sum_{i=1}^n \log \left(\sum_{k=1}^8 w_k \exp \left(-\frac{(y_k - k)^2}{2} \right) \right).$$

We can now compare the posterior predictive distribution $\hat{p}(\tilde{y}|y) = \hat{w}_k p(\tilde{y}|y, M_k)$, for these three methods of model averaging. Figure 1 shows the predictive distributions in one simulation with the sample size n varies from 3 to 200. We first notice that stacking of means (the middle row of graphs) gives an unappealing predictive distribution, even if its point estimate is reasonable. The broad and oddly spaced distribution here arises from nonidentification of w , and it demonstrates the general point that stacking of means does not even try to match the shape of the predictive distribution. The top and bottom row of graphs show how BMA picks up the single model that is closest in KL divergence, while stacking picks a combination; the benefits of stacking becomes clear for large n .

In this trivial non-parametric case, stacking of predictive distributions is almost the same as fitting a mixture model, except for the absence of the prior. The true model $N(3.4, 1)$ is actually a convolution of single models rather than a mixture, hence no approach can recover the true one from the model list. From Figure 2 we can compare the mean squared error and the mean logarithmic score of this three combination methods. The average log scores and errors are calculated through 500 repeated simulation and 200 test data generating from the true distribution. The left panel shows logarithmic score (or equivalent, expected log predictive density) of the predictive distribution. Stacking of predictive distributions always gives a larger score except for extremely small n . In the middle panel, it shows the main squared error by considering the posterior mean of predictive distribution to be a point estimation, even if it is not our focus. In this case, it is not surprising to see that stacking of predictive distributions gives almost the same optimal mean squared error as the stacking of means, both of which are better than the BMA. Two distributions close in KL divergence are close in each moment, while the reverse doesn't necessarily hold. This illustrates the necessary of matching the *distribution*, rather than matching the *moments* for the predictive distribution.

Finally it is worth pointing out that stacking depends only on the space expanded by all candidate models, while BMA or Pseudo-BMA weighting may be misled by such model expansion. If we add another $N(4, 1)$ as the 9th model in the model list above, stacking will not change at all in theory, even though it becomes non-convex and has infinite same-height mode. For BMA, it is equivalent to putting double prior mass on the original 4th model, which doubles the final weights for it. The right panel of Figure 2 shows such phenomenon: we fix sample size n to be 15 and add more and more $N(4, 1)$ models. As a result, BMA (or Pseudo-BMA weighting) puts more weights on $N(4, 1)$ and behaves worse, while stacking changes nowhere except for numerical fluctuation. It illustrates another benefit of stacking compared to BMA or Pseudo-BMA weights. We would not expect a combination method that would perform even worse as model candidate list expands, which may become a disaster when many similar weak models exist.

3.2 Linear subset regressions

The previous section demonstrates a simple example of combining several different nonparametric models. Now we turn to the parametric case. This example comes from Breiman (1996) who compares stacking to model selection. Here we work in a Bayesian framework.

Suppose the true model is

$$y = \beta_1 x_1 + \cdots + \beta_J x_J + \epsilon.$$

In the model ϵ independently comes from $N(0, 1)$. All the covariates x_j are independently from $N(5, 1)$. The number of predictors J is 15. The coefficient β is generated by

$$\beta_j = \gamma \left((1_{|j-4|<h} (h - |j - 4|)^2 + (1_{|j-8|<h} (h - |j - 8|)^2 + (1_{|j-12|<h} (h - |j - 12|)^2) \right),$$

where γ is determined by fixing the signal-to-noise ratio such that

$$\frac{\text{Var}(\sum_j \beta_j x_j)}{1 + \text{Var}(\sum_j \beta_j x_j)} = \frac{4}{5}.$$

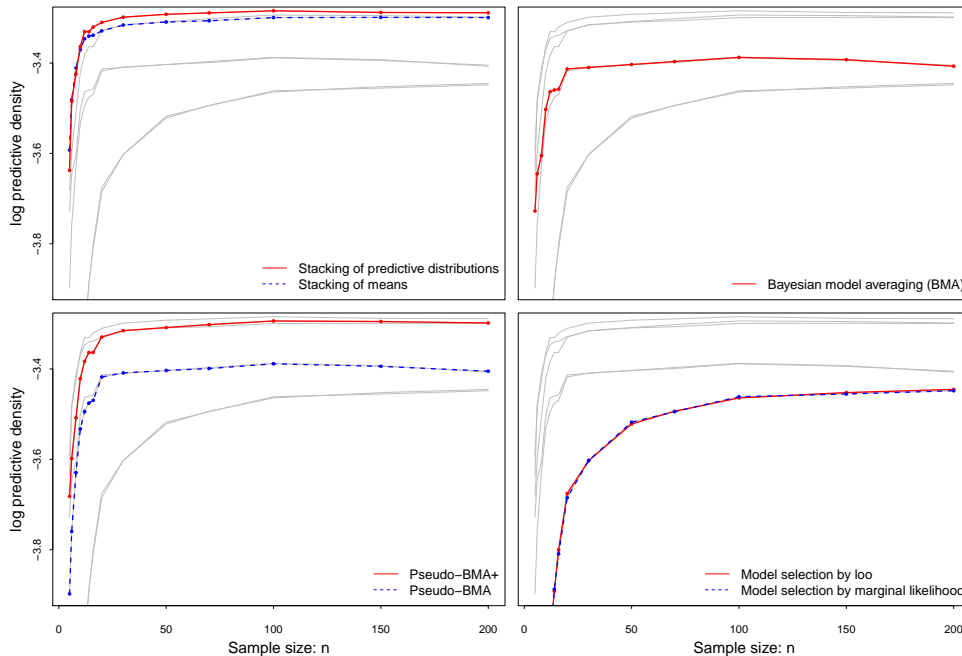


Figure 3: Mean log predictive density of 7 combination methods in the linear regression variable selection example: the k -th model is a univariate linear regression with the k -th variable. The mean log score is evaluated by 100 repeated experiments and 200 test data.

The value h determines the number of nonzero coefficient in the true model. For $h = 1$, there are 3 “strong” coefficient. For $h = 5$, there are 15 “weak” coefficients. In the following simulation, we fix $h = 5$.

We consider the following two model spaces:

1. \mathcal{M} -open

Each subset contains only one single variable. Hence, the k -th model is a univariate linear regression with the k -th variable x_k . We have $K = J = 15$ different models in total. One advantage of stacking and Pseudo-BMA weighting is that they are not sensitive to prior, hence even a flat prior will work, while BMA can be sensitive to the prior. For each single model $M_k : y \sim N(\beta_k x_k, \sigma^2)$, we set prior $\beta_k \sim N(0, 10)$, $\sigma \sim \text{Gamma}(0.1, 0.1)$.

2. \mathcal{M} -closed

Let model k be the linear regression with subset (x_1, \dots, x_k) . Then there are still $K = 15$ different models. Similarly, in model $M_k : y \sim N(\sum_{j=1}^k \beta_j x_j, \sigma^2)$, we set prior $\beta_j \sim N(0, 10), j = 1, \dots, k, \sigma \sim \text{Gamma}(0.1, 0.1)$.

In both cases, we have seven methods for combine predictive densities: stacking of predictive distributions, stacking of means, Pseudo-BMA, Pseudo-BMA+, best model selection

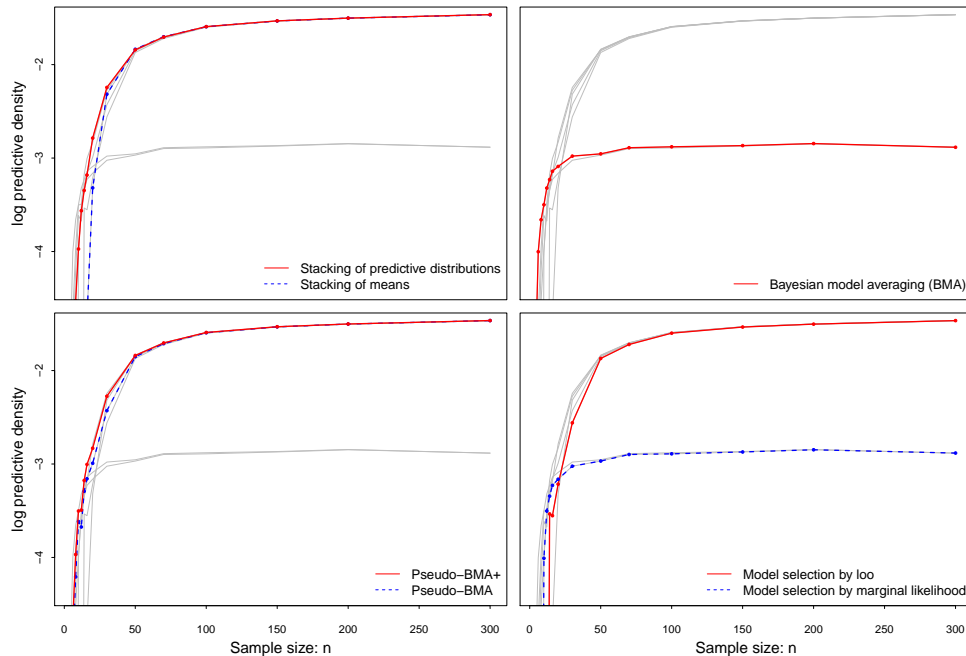


Figure 4: Mean log predictive density of 7 combination methods in the linear regression variable selection example: the k -th model is a linear regression with the first k variables. We evaluate the mean log score using 100 repeated experiments and 200 test data.

by mean LOO value, best model selection by marginal likelihood, and finally BMA. Each of these can be expressed as a linear combination $p(\tilde{y}|y) = \sum_{k=1}^K \hat{w}_k p(\tilde{y}|M_k)$ is what we have in the end as estimation for the posterior density of the new data \tilde{y} . We generate a test dataset $(\tilde{x}_i, \tilde{y}_i)$, $i = 1, \dots, 200$ from the underlying true distribution to calculate the out of sample score for the combined distribution under each method k :

$$\frac{1}{200} \sum_{i=1}^{200} \log \sum_{k=1}^K \hat{w}_k p(\tilde{y}_i | M_k).$$

We loop the test simulation 100 times to get the expected predictive accuracy for each method.

Figure 3 shows the expected out-of-sample log predictive density for the seven methods, for a set of experiments with sample size n ranging from 5 to 200. Stacking seems to outperform all other methods even for small n . Stacking of predictive distributions is asymptotically better than any other combination method. Pseudo-BMA+ weighting dominates naive Pseudo-BMA weighting. Finally, BMA performs similarly to Pseudo-BMA weighting, always better than any kind of model selection, but that advantage vanishes in the limit since BMA picks up one model. In this \mathcal{M} -open setting, we know model selection can never be optimal.

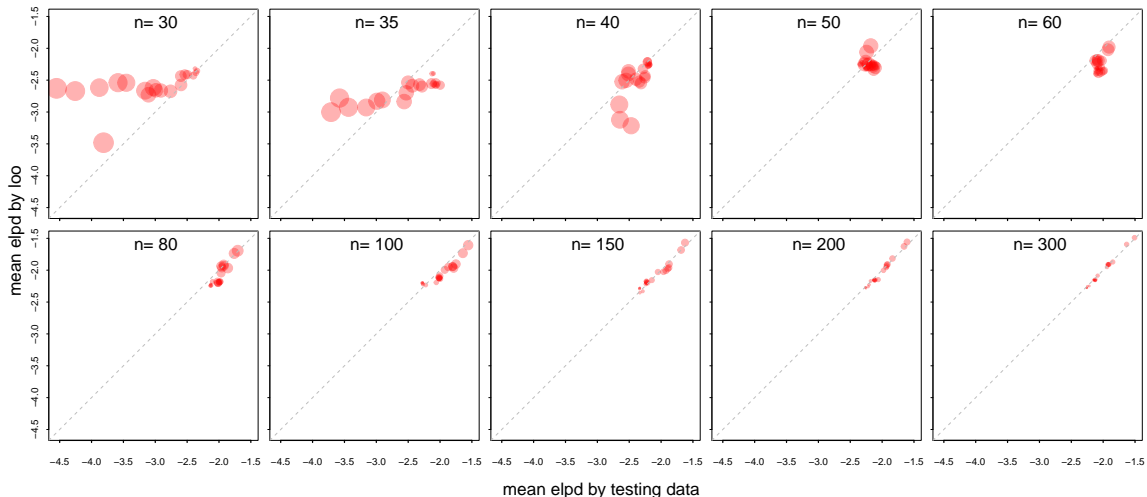


Figure 5: Comparison of the expected log predictive density (elpd) estimated by LOO, $\text{elpd}_{\text{loo}} = \frac{1}{n} \sum_{i=1}^n \log \hat{w}_k p_k(y_i | y_{-i})$, and the elpd calculated from test data, $\text{elpd}_{\text{test}} = \frac{1}{200} \sum_{i=1}^{200} \log \hat{w}_k p_k(\tilde{y}_i | y)$. The area of each dot represents the relative complexity of the model as measured by effective number of parameter divided by sample size.

The results change when we move to the second case, in which the k -th model contains variables x_1, \dots, x_k so that we are comparing models of differing dimensionality. The problem is \mathcal{M} -closed because the largest subset contains all the variables, and we have simulated data from this model. Figure 4 shows the mean log predictive density of the seven combination methods in this case. For a large sample size n , almost all methods recover the true answer (putting weight 1 on the full model), except BMA and model selection based on marginal likelihood. The poor performance of BMA comes from the parameter priors: recall that the optimality of BMA arises when averaging over the priors and not necessarily conditional on any particular chosen set of parameter values. There is no general no way to obtain a “correct” prior that accounts for the complexity for BMA in an arbitrary model space. Model selection by LOO can recover the true model, while selection by marginal likelihood cannot due to the same prior problems. Once again, we see that BMA eventually become the same as model selection by marginal likelihood, which is much worse than any other methods asymptotically.

In this example, stacking is unstable for extremely small n . In fact, our computations for stacking of predictive distributions and Pseudo-BMA depend on the the PSIS approximation $\log p(y_i | y_{-i})$. If this approximation is crude, then the second step optimization cannot be accurate. It is known that the parameter \hat{k} in the generalized Pareto distribution can be used to diagnose the accuracy of PSIS approximation. In our method, we replace PSIS approximation by running exact LOO for any data points with estimated $\hat{k} > 0.7$ (see Vehtari et al., 2016).

Figure 5 shows the comparison of the elpd estimated by LOO,

$$\text{elpd}_{\text{loo}} = \frac{1}{n} \sum_{i=1}^n \log \hat{w}_k p_k(y_i | y_{-i}),$$

and the elpd calculated using 200 independent test data,

$$\text{elpd}_{\text{test}} = \frac{1}{200} \sum_{i=1}^{200} \log \hat{w}_k p_k(\tilde{y}_i | y).$$

The area of each dot in Figure 5 represents the relative complexity of the model as measured by effective number of parameters in the model divided by sample size. We evaluate the effective number of parameters using LOO (Vehtari et al., 2016). The sample size n varies from 30 to 200 and variables size is fixed to be 20. Clearly, the relationship is far from the the line $y = x$ for extremely small sample size, and the relative bias ratio ($\text{elpd}_{\text{loo}}/\text{elpd}_{\text{test}}$) depends on the complexity of the model. Empirically, we have found the approximation to be poor when the sample size is less than 5 times the number of parameters.

As a result, in the small sample case, LOO can lead to relatively large variance, which makes the stacking of predictive distributions and Pseudo-BMA/BMA+ unstable, with performance improving quickly as n grows.

3.3 Comparison with mixture models

Stacking is inherently a two-step procedure. In contrast, when fitting a mixture model, one estimates the model weights and the status within parameters in the same step. One can go further with meta-learning algorithms such as boosting and bagging (Monteith et al., 2011). In BMA, (x_i, y_i) is an element of data set D , H is a finite set of all individual models. Then the posterior probability of h given D can be calculated as

$$p(y|x_i, e, E, H) = \sum_{h \in H} p(y_i|x_i, h)p(h|D).$$

In the mixture model, this is modified as

$$p(y|x_i, D, E, H) = \sum_{e \in E} p(y_i|x_i, H, e)p(e|D),$$

where e is an element in the space of all possible model combinations.

The mixture model seems to be the most straightforward continuous model expansion. Nevertheless, there are several reasons why we may prefer stacking to the mixture model, though the latter one is a full Bayesian approach. First, the computation cost of mixture models can be relatively large. If the true model is a mixture model and the estimation of each model depends a lot on others, then it is worth paying the extra computation cost. However, it is not quite possible to combine several components in real application. It is more likely that the researcher is running combination among several mixture models with different pre-specified number of components. In other words, E , the space of model combination can be always extended so it is infeasible to make such kind of full Bayesian inference.




Figure 6: *Predictive density of the combined posterior distribution obtained by stacking of predictive distributions, BMA, Pseudo-BMA, Pseudo-BMA+, model selection by marginal likelihood, or mixture models. In each case, we evaluate the predictive density by taking mean of 100 testing data and 100 repeated simulations experiments. The correlation of variables range from -0.3 to 0.9 , and sample sizes range from 4 to 50. Stacking on predictive distribution and Pseudo-BMA+ outperform mixture models in all cases.*

Second, if every single model is close to one another and sample size is small, the mixture model can face non-identification or instability problem, unless a strong prior is added. Since the mixture model is relatively complex, this leads to a poor small sample behavior.

Figure 6 shows a comparison of mixture model and our model averaging methods in a numerical experiment, in which the true model is

$$y \sim N(\beta_1 x_1 + \beta_2 x_2 + \beta_3 x_2, 1), \quad \beta_k \text{ is generated from } N(0, 1),$$

and there are 3 candidate models, each containing one predictor:

$$M_k : y \sim N(\beta_k x_k, \sigma_k^2), \text{ with a prior } \beta_k \sim N(0, 1), \quad k = 1, 2, 3.$$

In the simulation, we generate the design matrix by $\text{Var}(x_i) = 1$ and $\text{Cor}(x_i, x_j) = \rho$. ρ determines how correlated these models are and it ranges from -0.3 to 0.9 .

Figure 6 shows that both the performance of mixture models and single model selection are worse than any other model averaging method we suggest, even though the mixture model takes much longer time to run (about 30 more times) than stacking or Pseudo-BMA+. When the sample size is small, the mixture model is too complex to fit. On the other hand, stacking of predictive distributions and Pseudo-BMA+ outperform all other methods with a moderate sample size.

Clarke (2003) argues that the effect of (point estimation) stacking only depends on the space spanned by the model list, hence he suggests putting those “independent” models in the list. Figure 6 shows high correlations do not hurt stacking and Pseudo-BMA+ in this example.

3.4 Variational inference with different initial values

In variational inference, if the posterior distribution is multimodal, then the optimization procedure may stop at a local mode, hence the posterior draws will depend on initial values.

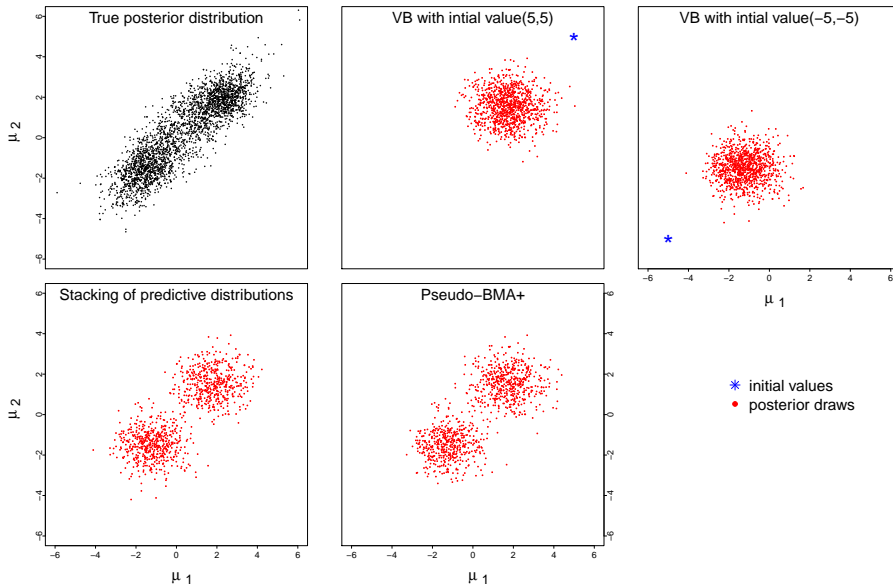


Figure 7: (1) A multimodal posterior distribution of (μ_1, μ_2) . (2–3) Posterior draws from variational inference with different initial values. (4–5) Averaged posterior distribution using stacking of predictive distributions and Pseudo-BMA+ weighting.

Besides averaging several models, stacking of predictive distributions and Pseudo-BMA+ weighting can also average several sets of posterior draws coming from different initial values. In such setting, BMA is not applicable.

Figure 7 gives a simple numerical example that the averaging strategy helps overcome the uncertainty of initial values. The model is

$$y_1 \sim \text{Cauchy}(\mu_1, 1), \quad y_2 \sim \text{Cauchy}(\mu_2, 1).$$

A prior is put on $\mu_1 - \mu_2$:

$$\mu_1 - \mu_2 \sim N(0, 1).$$

We generate two observations $y_1 = (3, -2), y_2 = (2, -2)$. The first panel shows the true posterior distribution of (μ_1, μ_2) , which is bimodal. With the initial value set to be $(\mu_1, \mu_2) = (5, 5)$ and $(-5, -5)$, the variation inference can recover one of the two modes separately. The lower 2 panels show the combined posterior distribution using stacking and Pseudo-BMA+. Though neither can recover the true distribution, the averaged version is closer to it, measured in KL divergence.

In general, variational inference tries to minimize $\text{KL}(q(\theta), p(\theta|y)) = \mathbb{E}_q \log(p) - \mathbb{E}_q \log(q)$, while stacking of predictive distributions minimizes $\text{KL}(p(\theta), q(\theta|y)) = C + \mathbb{E}_q \log(p)$. Since stacking uses LOO cross-validation in the first step to avoid overfitting, it seems unnecessary to put another extra penalization term on the entropy $\mathbb{E}_q \log(q)$.

3.5 Proximity and directional models of voting

Adams et al. (2004) use Senate voting data from 1988 to 1992 to study voters' preference on the candidates who propose policies that are similar to them. They introduce two similar variables that indicate the distance between voters and candidates. *Proximity voting comparison* represents the i -th voter's comparison between the candidates' ideological positions:

$$U_i(D) - U_i(R) = (x_R - x_i)^2 - (x_D - x_i)^2,$$

where x_i represents the i -th voter's preferred ideological position, and x_D and x_R represent the ideological positions of the Democratic and Republican candidates, respectively. In contrast, the i -th voter's *directional comparison* is defined by

$$U_i(D) - U_i(R) = (x_D - X_N)(x_i - X_N) - (x_R - X_N)(x_i - X_N),$$

where X_N is the neutral point of the ideology scale.

Finally, all these comparison is aggregated in the party level, leading to two party-level variable *Democratic proximity advantage* and *Democratic directional advantage*. The sample size is $n = 94$.

For both of these two variables, there are two ways to measure candidates' ideological positions x_D and x_R , which leads to two different datasets. In the *Mean candidate* dataset, they are calculated by taking the average of all respondents' answers in the relevant state and year. In the *Voter-specific* dataset, they are calculate by using respondents' own placements of the two candidates. In both datasets, there are 4 other party-level variables.

The two variables *Democratic proximity advantage* and *Democratic directional advantage* are highly correlated. Montgomery and Nyhan (2010) point out that Bayesian model averaging is an approach to helping arbitrate between competing predictors in a linear regression model. They average over all 2^6 linear subset models excluding those that contain both variables *Democratic proximity advantage* and *Democratic directional advantage*, (i.e., 48 models in total). Each subset regression is with the form

$$M_\gamma : y|X, \beta_0, \beta \sim N(\beta_0 + X_\gamma \beta_\gamma, \sigma^2).$$

Accounting for the different complexity, they used the hyper- g prior (Liang et al., 2008). Let ϕ to be the inverse of the variance $\phi = \frac{1}{\sigma^2}$. The hyper- g prior with a prespecified hyperparameter α is,

$$\begin{aligned} p(\phi) &\propto \frac{1}{\phi} \\ \beta|g, \phi, X &\sim N\left(0, \frac{g}{\phi}(X^T X)^{-1}\right) \\ p(g|\alpha) &= \frac{\alpha - 2}{2}(1 + g)^{-\alpha/2}, \quad g > 0. \end{aligned}$$

The first two columns of Figure 8 show the linear regression coefficients as estimated using least squares. The remaining columns show the posterior mean and standard deviation of the regression coefficients using BMA, stacking of predictive distributions, and

	<i>Full model</i>		<i>BMA</i>		<i>Stacking of predictive distributions</i>		<i>Pseudo-BMA+ weighting</i>	
	<i>Mean Candidate</i>	<i>Voter-specific</i>	<i>Mean Candidate</i>	<i>Voter-specific</i>	<i>Mean Candidate</i>	<i>Voter-specific</i>	<i>Mean Candidate</i>	<i>Voter-specific</i>
Dem. prox. adv.	-3.05 (1.32)	-2.01 (1.06)	-0.22 (0.95)	0.75 (0.68)	0.00 (0.00)	0.00 (0.00)	-0.02 (0.08)	0.04(0.24)
Dem. direct. adv.	7.95 (2.85)	4.18 (1.36)	3.58 (2.02)	2.36 (0.84)	2.56 (2.32)	1.93 (1.16)	1.60 (4.91)	1.78 (1.22)
Dem. incumb. adv.	1.06 (1.20)	1.14 (1.19)	1.61 (1.24)	1.30 (1.24)	0.48 (1.70)	0.34 (0.89)	0.66 (1.13)	0.54 (1.03)
Dem. quality adv.	3.12 (1.24)	2.38 (1.22)	2.96 (1.25)	2.74 (1.22)	2.20 (1.71)	2.30 (1.52)	2.05 (2.86)	1.89 (1.61)
Dem. spend adv.	0.27 (0.04)	0.27 (0.04)	0.32 (0.04)	0.31 (0.04)	0.31 (0.07)	0.31 (0.03)	0.31 (0.04)	0.30 (0.04)
Dem. partisan adv.	0.06 (0.05)	0.06 (0.05)	0.08 (0.06)	0.07 (0.06)	0.01 (0.04)	0.00 (0.00)	0.03 (0.05)	0.03 (0.05)
Const.	53.3 (1.2)	52.0 (0.8)	51.4 (1.0)	51.6 (0.8)	51.9 (1.1)	51.6 (0.7)	51.5 (1.2)	51.4 (0.8)

Figure 8: *Regression coefficient and standard error in the voting example, from the full model (columns 1–2), the subset regression model averaging using BMA (columns 3–4), stacking of predictive distributions (columns 5–6) and Pseudo-BMA+ (columns 7–8). Democratic proximity advantage and Democratic directional advantage are two highly correlated variables. Mean candidate and Voter-specific are two datasets that provide different measurements on candidates’ ideological placement.*

Pseudo-BMA+, respectively. Under all three averaging strategies, the coefficient of *proximity advantage* is no longer statistically significantly negative, and the coefficient of *directional advantage* is shrunk. As fit to these data, stacking puts near-zero weights on all subset models containing *proximity advantage*, whereas Pseudo-BMA+ weighting always gives some weight to each model. In this example, averaging subset models by stacking or Pseudo-BMA+ weighting gives a way to deal with competing variables, which should be more reliable than BMA according to our previous argument.

3.6 Predicting well-switching behavior in Bangladesh

Many wells in Bangladesh and other South Asian countries are contaminated with natural arsenic. People whose wells have arsenic levels that exceed a certain threshold are encouraged to switch to nearby safe wells (for background details, see Chapter 5 in Gelman and Hill, 2006). We are analyzing a dataset including 3020 respondents to find factors predictive of the well switching among all people with unsafe wells. The outcome variable is

$$y_i = \begin{cases} 1 & \text{if household } i \text{ switched to a new well} \\ 0 & \text{if household } i \text{ continued using its own well.} \end{cases}$$

And we consider the following inputs:

- The distance (in meters) to the closest known safe well,

- The arsenic level of the respondent’s well,
- Whether any member of the household is active in the community association,
- The education level of the head of the household.

We start with what we call Model 1, a simple logistic regression with all variables above as well as a constant term:

$$y \sim \text{Bernoulli}(\theta)$$

$$\theta = \text{logit}^{-1}(\beta_0 + \beta_1 \text{dist} + \beta_2 \text{arsenic} + \beta_3 \text{assoc} + \beta_4 \text{edu})$$

Model 2 contains the interaction between distance and arsenic level.

$$\theta = \text{logit}^{-1}(\beta_0 + \beta_1 \text{dist} + \beta_2 \text{arsenic} + \beta_3 \text{assoc} + \beta_4 \text{edu} + \beta_5 \text{dist} \times \text{arsenic})$$

Furthermore, it makes sense to us a nonlinear model for logit switching probability as a function of distance and arsenic level. We can use spline to capture that. Our Model 3 contains the B-splines for distance and arsenic level with polynomial degree 2,

$$\theta = \text{logit}^{-1}(\beta_0 + \beta_1 \text{dist} + \beta_2 \text{arsenic} + \beta_3 \text{assoc} + \beta_4 \text{edu} + \alpha_{dis} B_{dis} + \alpha_{ars} B_{ars})$$

where B_{dis} is the B-spline basis of distance with the form $(B_{dis,1}(\text{dist}), \dots, B_{dis,q}(\text{dist}))$ and $\alpha_{dis}, \alpha_{ars}$ are vectors. We also fix the number of knots to be 10 for both distance and arsenic level. Model 4 and 5 are the similar models with 3-degree and 5-degree B-splines, respectively.

Next, we can add a bivariate spline to capture a nonlinear interaction,

$$\theta = \text{logit}^{-1}(\beta_0 + \beta_1 \text{dist} + \beta_2 \text{arsenic} + \beta_3 \text{assoc} + \beta_4 \text{edu} + \beta_5 \text{dist} \times \text{arsenic} + \alpha B_{dis,ars})$$

where $B_{dis,ars}$ is the bivariate spline basis with the degree to be $2 \times 2, 3 \times 3, 5 \times 5$ in Model 6, 7 and 8 respectively.

Figure 9 shows the inference results in all 8 models, which are summarized by the posterior mean, 50% confidence interval and 95% confidence interval of the probability of switching from an unsafe well as a function of distance or arsenic level. Any other variables such as `assoc` and `edu` are fixed at their means. It is not obvious to pick one best model. Spline models give a more flexible shape, but also introduce more variance for posterior estimation.

Finally, we run stacking of predictive distributions and Pseudo-BMA+ to combine these 8 models. The calculated model weights are listed above each panel in Figure 9. For both combination methods, Model 5 (univariate splines with 5th degree) accounts for the majority share. It is also worth pointing out that Model 8 is the most complicated one, but both stacking and Pseudo-BMA+ avoid overfitting by assigning a very small weight on it.

Figure 10 shows the posterior mean, 50% confidence interval, and 95% confidence interval of the switching probability in the stacking-combined model. Pseudo-BMA+ weighting gives a similar combination result for this example. At first glance, the combination looks quite similar to Model 5, while it may not seem necessary to put an extra 0.09 weight on Model 1 in stacking combination since Model 1 is completely contained in Model 5 if

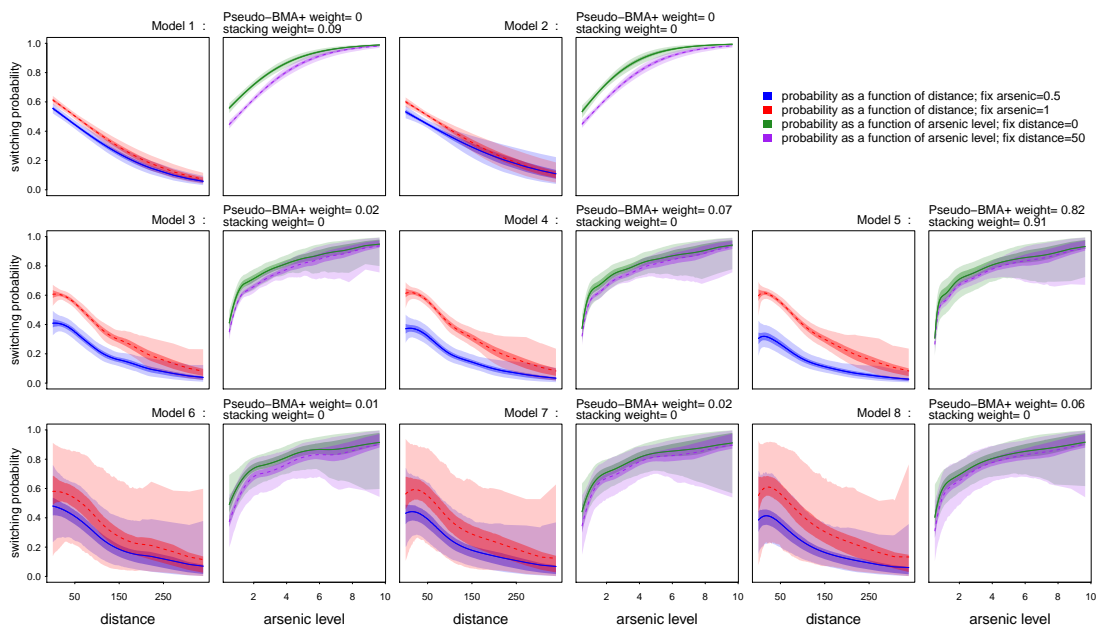


Figure 9: *Posterior mean, 50% confidence interval, and 95% confidence interval of the probability of switching from an unsafe well in Models 1–8. For each model, the switching probability is shown as a function of (a) the distance to the nearest safe well or (b) the arsenic level of the existing well. For each plot, the other input variable is held constant at different representative values. The model weights by stacking of predictive distributions and Pseudo-BMA+ are listed above each panel.*




Figure 10: *Posterior mean, 50% confidence interval, and 95% confidence interval of the probability of switching from an unsafe well in the combined model via stacking of predictive distributions. Pseudo-BMA+ weighting gives a similar result for the combination.*

setting $\alpha_{dis} = \alpha_{ars} = 0$. However, Model 5 is not perfect since it predicts that the posterior mean of switching probability will decrease as a function of distance to the nearest safe well, for very small distances. In fact, without further control, it is not surprising to find boundary fluctuation as a main drawback for higher order splines. Fortunately, we notice this decrease trend around the left boundary is a little bit flatter in the combined distribution since the combination contains part of straightforward logistic regression (in stacking weights) or lower order splines (in Pseudo-BMA+ weights). In this example the sample size $n = 3020$ is large, hence we have reasons to believe stacking of predictive distributions gives the optimal combination.

4. Discussion

4.1 Sparse structure and high dimensions

Yang and Dunson (2014) propose to estimate a linear combination of point forecasts, $f = \sum_{k=1}^K w_k f_k$, using a Dirichlet aggregation prior, $w \sim \text{Dirichlet}(\frac{\alpha}{K^\gamma}, \dots, \frac{\alpha}{K^\gamma})$, to pull toward structure, and estimating the weights w_k using adaptive regression rather than cross-validation. They show that the combination under this setting can achieve the mini-max squared risk among all convex combinations,

$$\sup_{f_1, \dots, f_K \in F_0} \inf_{\hat{f}} \sup_{f_\lambda^* \in F_\Gamma} E \|\hat{f} - f_\lambda^*\|,$$

where $F_0 = \{f : \|f\|_\infty \leq 1\}$.

Similar to our problem, when the dimension of model space is high, it can make sense to assign a strong prior can to the weights in estimation equation (4) to improve the regularization, using a hierarchical prior to pull toward sparsity if that is desired.

4.2 Constraints and regularity

In point estimation stacking, the simplex constraint is the most widely used regularization so as to overcome potential problems with multicollinearity. Clarke (2003) suggests relaxing the constraint to make it more flexible.

When combining distributions, there is no need to worry about multicollinearity except in degenerate cases. But in order to guarantee a meaningful posterior predictive density, the simplex constraint becomes natural, which is satisfied automatically in BMA and Pseudo-BMA weighting. As mentioned in the previous section, stronger priors can be added.

Another assumption is that the separate posterior distributions are combined linearly and with weights that are positive and sum to 1. There could be gains from going beyond convex linear combinations. For instance, in the subset regression example when each individual model is a univariate regression, the true model distribution is a convolution instead of a mixture of each possible models distribution. Both of them lead to the additive model in the point estimation, so stacking of the means is always valid, while stacking of predictive distributions is not possible to recover the true model in the convolution case.

Our explanation is that when the model list is large, the convex span should be large enough to approximate the true model. And this is the reason why we prefer adding stronger priors to make the estimation of weights stable in high dimensions.

4.3 General recommendations

The methods discussed in this paper are all based on the idea of fitting models separately and then combining the estimated predictive distributions. This approach is limited in that it does not pool information between the different model fits: as such, it is only ideal when the K different models being fit have nothing in common. But in that case we would prefer to fit a larger super-model that includes the separate models as special cases, perhaps using an informative prior distribution to ensure stability in inferences.

That said, in practice it is common for different sorts of models to be set up without any easy way to combine them, and in such cases it is necessary from a Bayesian perspective to somehow aggregate their predictive distributions. The often-recommended approach of Bayesian model averaging can fail catastrophically in that the required Bayes factors can depend entirely on arbitrary specifications of noninformative prior distributions. Stacking is a more promising general method in that it is directly focused on performance of the combined predictive distribution. Based on our theory, simulations, and examples, we recommend stacking (of predictive distributions) for the task of combining separately-fit Bayesian posterior predictive distributions. As an alternative, Pseudo-BMA+ is computationally cheaper and can serve as an initial guess for stacking. The computations can be done in R and Stan, and the optimization required to compute the weights connects directly to the predictive task.

Appendix A. Implementation in Stan and R

The $n \times p$ matrix of cross-validated log likelihood values, $p(y_i|y_{-i}, M_k)$, can be computed in the generated quantities block in a Stan program, following the approach of Vehtari et al. (2016). For the example in Section 3.2, the k -th model is a linear regression with the k -th covariates. We put the corresponding Stan code in the file `regression.stan`:

```
data {
  int n;
  int p;
  vector[n] y;
  matrix[n, p] X;
}
parameters {
  vector[p] beta;
  real<lower=0> sigma;
}
transformed parameters {
  vector[n] theta;
  theta = X * beta;
}
model {
  y ~ normal(theta, sigma);
  beta ~ normal(0, 10);
  sigma ~ gamma(0.1, 0.1);
}
```

```

}
generated quantities {
  vector[n] log_lik;
  for (i in 1:n)
    log_lik[i] = normal_lpdf(y[i] | theta[i], sigma);
}

```

In R we can simulate the likelihood matrices from all K models and save them as a list:

```

library("rstan")
log_lik_list <- list()
for (k in 1:K){
  # Fit the k-th model with Stan
  fit <- stan("regression.stan", data=list(y=y, X=X[,k], n=length(y), p=1))
  log_lik_list[[k]] <- extract(fit)[["log_lik"]]
}

```

The function `model_weights()` in the `loo` package in R can give model combination weights according to stacking of predictive distributions, Pseudo-BMA and Pseudo-BMA+ weighting. We can choose whether to use the Bayesian bootstrap to make further regularization for Pseudo-BMA, if computation time is not a concern.

```

model_weights_1 <- model_weights(log_lik_list, method="stacking")
model_weights_2 <- model_weights(log_lik_list, method="pseudobma", BB=TRUE)

```

in one simulation with six models to combine, the output gives us the computed weights under each approach:

The stacking weights are:

```

[1,] "Model 1" "Model 2" "Model 3" "Model 4" "Model 5" "Model 6"
[2,] "0.25"    "0.06"    "0.09"    "0.25"    "0.35"    "0.00"

```

The Pseudo-BMA+ weights using Bayesian Bootstrap are:

```

[1,] "Model 1" "Model 2" "Model 3" "Model 4" "Model 5" "Model 6"
[2,] "0.28"    "0.05"    "0.08"    "0.30"    "0.28"    "0.00"

```

For reasons discussed in the paper, we generally recommend stacking for combining separate Bayesian predictive distributions.

Acknowledgments

We thank the U.S. National Science Foundation, Institute for Education Sciences, Office of Naval Research, and Defense Advanced Research Projects Administration for partial support of this work.

References

- James Adams, Benjamin G. Bishin, and Jay K. Dow. Representation in congressional campaigns: Evidence for discounting/directional voting in U.S. Senate elections. *Journal of Politics*, 66(2):348–373, 2004.
- Hirotsugu Akaike. On the likelihood of a time series model. *The Statistician*, pages 217–235, 1978.
- José M. Bernardo and Adrian F. M. Smith. *Bayesian Theory*. John Wiley & Sons, 1994.
- Leo Breiman. Stacked regressions. *Machine Learning*, 24(1):49–64, 1996.
- Kenneth P. Burnham and David R. Anderson. *Model Selection and Multi-Model Inference: A Practical Information-Theoretic Approach*. Springer, 2nd edition, 2002.
- Bertrand Clarke. Comparing Bayes model averaging and stacking when model approximation error cannot be ignored. *Journal of Machine Learning Research*, 4(Oct):683–712, 2003.
- Merlise Clyde and Edwin Severin Iversen. Bayesian model averaging in the M-open framework. In Paul Damien, Petros Dellaportas, Nicholas G. Polson, and David A. Stephens, editors, *Bayesian Theory and Applications*, pages 483–498. Oxford University Press, 2013.
- Seymour Geisser and William F. Eddy. A predictive approach to model selection. *Journal of the American Statistical Association*, 74(365):153–160, March 1979.
- Alan E. Gelfand. Model determination using sampling-based methods. In W. R. Gilks, S. Richardson, and D. J. Spiegelhalter, editors, *Markov Chain Monte Carlo in Practice*, pages 145–162. Chapman & Hall, 1996.
- Andrew Gelman. Parameterization and Bayesian modeling. *Journal of the American Statistical Association*, 99(466):537–545, 2004.
- Andrew Gelman and Jennifer Hill. *Data Analysis Using Regression and Multi-level/Hierarchical Models*. Cambridge University Press, 2006.
- Andrew Gelman, Jessica Hwang, and Aki Vehtari. Understanding predictive information criteria for Bayesian models. *Statistics and Computing*, 24(6):997–1016, 2014.
- John Geweke and Gianni Amisano. Optimal prediction pools. *Journal of Econometrics*, 164(1):130–141, 2011.
- John Geweke and Gianni Amisano. Prediction with misspecified models. *American Economic Review*, 102(3):482–486, 2012.
- Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- Eduardo Gutiérrez-Peña and Stephen G. Walker. Statistical decision problems and Bayesian nonparametric methods. *International Statistical Review*, 73(3):309–330, 2005.

- Jennifer A. Hoeting, David Madigan, Adrian E. Raftery, and Chris T. Volinsky. Bayesian model averaging: A tutorial. *Statistical Science*, 14(4):382–401, 1999.
- Jane T. Key, Luis R. Pericchi, and Adrian F. M. Smith. Bayesian model choice: What and why. *Bayesian Statistics*, 6:343–370, 1999.
- Tri Le and Bertrand Clarke. A Bayes interpretation of stacking for M-complete and M-open settings. *arXiv:1602.05162*, 2016.
- Michael LeBlanc and Robert Tibshirani. Combining estimates in regression and classification. *Journal of the American Statistical Association*, 91(436):1641–1650, 1996.
- Meng Li and David B. Dunson. A framework for probabilistic inferences from imperfect models. *arXiv:1611.01241v1*, 2016.
- Feng Liang, Rui Paulo, German Molina, Merlise A. Clyde, and James O. Berger. Mixtures of g priors for Bayesian variable selection. *Journal of the American Statistical Association*, 103(481):410–423, 2008.
- David Madigan, Adrian E. Raftery, C. Volinsky, and J. Hoeting. Bayesian model averaging. In *Proceedings of the AAAI Workshop on Integrating Multiple Learned Models*, pages 77–83, 1996.
- Christopher J. Merz and Michael J. Pazzani. A principal components approach to combining regression estimates. *Machine Learning*, 36(1-2):9–32, 1999.
- Kristine Monteith, James L. Carroll, Kevin Seppi, and Tony Martinez. Turning Bayesian model averaging into Bayesian model combination. In *Proceedings of the 2011 International Joint Conference on Neural Networks (IJCNN)*, pages 2657–2663. IEEE, 2011.
- Jacob M. Montgomery and Brendan Nyhan. Bayesian model averaging: Theoretical developments and practical applications. *Political Analysis*, 18(2):245–270, 2010.
- Juho Piironen and Aki Vehtari. Comparison of Bayesian predictive methods for model selection. *Statistics and Computing*, pages 1–25, 2015.
- Donald B. Rubin. The Bayesian bootstrap. *Annals of Statistics*, 9(1):130–134, 1981.
- M. Stone. An asymptotic equivalence of choice of model by cross-validation and Akaike’s criterion. *Journal of the Royal Statistical Society. Series B (Methodological)*, 39(1):44–47, 1977.
- Kai Ming Ting and Ian H. Witten. Issues in stacked generalization. *Journal of Artificial Intelligence Research*, 10:271–289, 1999.
- Aki Vehtari and Jouko Lampinen. Bayesian model assessment and comparison using cross-validation predictive densities. *Neural Computation*, 14(10):2439–2468, 2002.
- Aki Vehtari and Janne Ojanen. A survey of Bayesian predictive methods for model assessment, selection and comparison. *Statistics Surveys*, 6:142–228, 2012.

Aki Vehtari, Andrew Gelman, and Jonah Gabry. Practical Bayesian model evaluation using leave-one-out cross-validation and WAIC. *Statistics and Computing*, pages doi:10.1007/s11222-016-9696-4, 2016.

Eric-Jan Wagenmakers and Simon Farrell. Aic model selection using akaike weights. *Psychonomic bulletin & review*, 11(1):192–196, 2004.

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

David H. Wolpert. Stacked generalization. *Neural Networks*, 5(2):241–259, 1992.

Yun Yang and David B. Dunson. Minimax optimal Bayesian aggregation. *arXiv:1403.1345*, 2014.