
On the Hyperprior Choice for the Global Shrinkage Parameter in the Horseshoe Prior

Juho Piironen
juho.piironen@aalto.fi

Aki Vehtari
aki.vehtari@aalto.fi

Helsinki Institute for Information Technology, HIIT
Department of Computer Science, Aalto University

Abstract

The horseshoe prior has proven to be a noteworthy alternative for sparse Bayesian estimation, but as shown in this paper, the results can be sensitive to the prior choice for the global shrinkage hyperparameter. We argue that the previous default choices are dubious due to their tendency to favor solutions with more unshrunk coefficients than we typically expect a priori. This can lead to bad results if this parameter is not strongly identified by data. We derive the relationship between the global parameter and the effective number of nonzeros in the coefficient vector, and show an easy and intuitive way of setting up the prior for the global parameter based on our prior beliefs about the number of nonzero coefficients in the model. The results on real world data show that one can benefit greatly – in terms of improved parameter estimates, prediction accuracy, and reduced computation time – from transforming even a crude guess for the number of nonzero coefficients into the prior for the global parameter using our framework.

1 INTRODUCTION

With modern rich datasets, statistical models with a large number of parameters are nowadays commonplace in many application areas. A typical example is a regression or classification problem with a large number of predictor variables. In such problems, a careful formulation of the prior distribution – or regularization

in the frequentist framework – plays a key role.

Often it is reasonable to assume that only some of the model parameters $\theta = (\theta_1, \dots, \theta_D)$ (such as the regression coefficients) are far from zero. In the frequentist literature, these problems are typically handled by LASSO (Tibshirani, 1996) or one of its close cousins, such as the elastic net (Zou and Hastie, 2005). We focus on the probabilistic approach and carry out the full Bayesian inference on the problem.

Two prior choices dominate the Bayesian literature: two component discrete mixture priors known as the spike-and-slab (Mitchell and Beauchamp, 1988; George and McCulloch, 1993), and a variety of continuous shrinkage priors (see e.g., Polson and Scott, 2011, and references therein). The spike-and-slab prior is intuitively appealing as it is equivalent to Bayesian model averaging (BMA) (Hoeting et al., 1999) over the variable combinations, and often has good performance in practice. The disadvantages are that the results can be sensitive to prior choices (slab width and prior inclusion probability) and that the posterior inference can be computationally demanding with a large number of variables, due to the huge model space. The inference could be sped up by analytical approximations using either expectation propagation (EP) (Hernández-Lobato et al., 2010, 2015) or variational inference (VI) (Titsias and Lázaro-Gredilla, 2011), but this comes at the cost of a substantial increase in the amount of analytical work needed to derive the equations separately for each model and a more complex implementation.

The continuous shrinkage priors on the other hand are easy to implement, provide convenient computation using generic sampling tools such as Stan (Stan Development Team, 2016), and can yield as good or better results. A particularly interesting example is the horseshoe prior (Carvalho et al., 2009, 2010)

$$\begin{aligned} \theta_j | \lambda_j, \tau &\sim N(0, \lambda_j^2 \tau^2), \\ \lambda_j &\sim C^+(0, 1), \quad j = 1, \dots, D, \end{aligned} \quad (1)$$

which has shown comparable performance to the spike-and-slab prior in a variety of examples where a sparsifying prior on the model parameters θ_j is desirable (Carvalho et al., 2009, 2010; Polson and Scott, 2011). The horseshoe is one of the so called global-local shrinkage priors, meaning that there is a global parameter τ that shrinks all the parameters towards zero, while the heavy-tailed half-Cauchy priors allow some parameters θ_j to escape the shrinkage (see Section 2 for more thorough discussion).

So far there has been no consensus on how to carry out inference for the global hyperparameter τ which determines the overall sparsity in the parameter vector $\boldsymbol{\theta}$ and therefore has a large impact on the results. For reasons discussed in Section 3.1, we prefer full Bayesian inference. For many interesting datasets, τ may not be well identified by data, and in such situations the hyperprior choice $p(\tau)$ becomes crucial. This is the question we address in this paper.

The novelty of the paper is summarized as follows. We derive analytically the relationship between τ and m_{eff} – the effective number of nonzero components in $\boldsymbol{\theta}$ (to be defined later) – and show an easy and intuitive way of formulating the prior for τ based on our prior beliefs about the sparsity of $\boldsymbol{\theta}$. We focus on regression and classification, but the methodology is applicable also to other generalized linear models and to other shrinkage priors than the horseshoe. We argue that the previously proposed default priors are dubious based on the prior they impose on m_{eff} , and that they yield good results only when τ (and therefore m_{eff}) is strongly identified by the data. Moreover, we show with several real world examples that in those cases where τ is only weakly identified by data, one can substantially improve inferences by transforming even a crude guess of the sparsity level into $p(\tau)$ using our method.

We first briefly review the key properties of the horseshoe prior in Section 2, and then proceed to discuss the prior choice $p(\tau)$ in Section 3. The importance of the concept is illustrated in Section 4 with real world data. Section 5 concludes the paper by giving some recommendations on the prior choice based on the theoretical considerations and numerical experiments.

2 THE HORSESHOE PRIOR FOR LINEAR REGRESSION

Consider the single output linear Gaussian regression model with several input variables, given by

$$y_i = \boldsymbol{\beta}^\top \mathbf{x}_i + \varepsilon_i, \quad \varepsilon_i \sim \text{N}(0, \sigma^2), \quad i = 1, \dots, n, \quad (2)$$

where \mathbf{x} is the D -dimensional vector of inputs, $\boldsymbol{\beta}$ contains the corresponding weights and σ^2 is the noise

variance. The horseshoe prior is set for the regression coefficients $\boldsymbol{\beta} = (\beta_1, \dots, \beta_D)$

$$\begin{aligned} \beta_j | \lambda_j, \tau &\sim \text{N}(0, \lambda_j^2 \tau^2), \\ \lambda_j &\sim \text{C}^+(0, 1), \quad j = 1, \dots, D. \end{aligned} \quad (3)$$

If an intercept term β_0 is included in the model (2), we give it a relatively flat prior, because there is usually no reason to shrink it towards zero. As discussed in the introduction, the horseshoe prior has been shown to possess several desirable theoretical properties and good performance in practice (Carvalho et al., 2009, 2010; Polson and Scott, 2011; Datta and Ghosh, 2013; van der Pas et al., 2014). The intuition is the following: the global parameter τ pulls all the weights globally towards zero, while the thick half-Cauchy tails for the local scales λ_j allow some of the weights to escape the shrinkage. Different levels of sparsity can be accommodated by changing the value of τ : with large τ all the variables have very diffuse priors with very little shrinkage towards zero, but letting $\tau \rightarrow 0$ will shrink all the weights β_j to zero.

The above can be formulated more formally as follows. Let \mathbf{X} denote the n -by- D matrix of observed inputs and \mathbf{y} the observed targets. The conditional posterior for the coefficients $\boldsymbol{\beta}$ given the hyperparameters and data $\mathcal{D} = (\mathbf{X}, \mathbf{y})$ can be written as

$$\begin{aligned} p(\boldsymbol{\beta} | \boldsymbol{\Lambda}, \tau, \sigma^2, \mathcal{D}) &= \text{N}(\boldsymbol{\beta} | \bar{\boldsymbol{\beta}}, \boldsymbol{\Sigma}), \\ \bar{\boldsymbol{\beta}} &= \tau^2 \boldsymbol{\Lambda} (\tau^2 \boldsymbol{\Lambda} + \sigma^2 (\mathbf{X}^\top \mathbf{X})^{-1})^{-1} \hat{\boldsymbol{\beta}}, \\ \boldsymbol{\Sigma} &= (\tau^{-2} \boldsymbol{\Lambda}^{-1} + \frac{1}{\sigma^2} \mathbf{X}^\top \mathbf{X})^{-1}, \end{aligned}$$

where $\boldsymbol{\Lambda} = \text{diag}(\lambda_1^2, \dots, \lambda_D^2)$ and $\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}^\top \mathbf{y}$ is the maximum likelihood solution (assuming $(\mathbf{X}^\top \mathbf{X})^{-1}$ exists). If the predictors are uncorrelated with zero mean and unit variance, then $\mathbf{X}^\top \mathbf{X} \approx n\mathbf{I}$, and we can approximate

$$\bar{\beta}_j = (1 - \kappa_j) \hat{\beta}_j, \quad (4)$$

where

$$\kappa_j = \frac{1}{1 + n\sigma^{-2}\tau^2\lambda_j^2} \quad (5)$$

is the *shrinkage factor* for coefficient β_j . The shrinkage factor describes how much coefficient β_j is shrunk towards zero from the maximum likelihood solution ($\kappa_j = 1$ meaning complete shrinkage and $\kappa_j = 0$ no shrinkage). From (4) and (5) it is easy to verify that $\bar{\boldsymbol{\beta}} \rightarrow 0$ as $\tau \rightarrow 0$, and $\bar{\boldsymbol{\beta}} \rightarrow \hat{\boldsymbol{\beta}}$ as $\tau \rightarrow \infty$.

The result (5) holds for any prior that can be written as a scale mixture of Gaussians like (3), regardless of the prior for λ_j . The horseshoe employs independent half-Cauchy priors for all λ_j , and for this choice one can




Figure 1: Density for the shrinkage factor (5) for the horseshoe prior (3) when $n\sigma^{-2}\tau^2 = 1$ (solid) and when $n\sigma^{-2}\tau^2 = 0.1$ (dashed).

show that, for fixed τ and σ , the shrinkage factor (5) follows the prior

$$p(\kappa_j | \tau, \sigma) = \frac{1}{\pi} \frac{\sigma^{-1}\tau\sqrt{n}}{(n\sigma^{-2}\tau^2 - 1)\kappa_j + 1} \frac{1}{\sqrt{\kappa_j}\sqrt{1 - \kappa_j}}. \quad (6)$$

When $n\sigma^{-2}\tau^2 = 1$, this reduces to $\text{Beta}(\frac{1}{2}, \frac{1}{2})$ which looks like a horseshoe, see Figure 1. Thus, a priori, we expect to see both relevant ($\kappa = 0$, no shrinkage) and irrelevant ($\kappa = 1$, shrinkage) variables. By changing the value of τ , the prior for κ places more mass either close to 0 or 1. For instance, choosing τ so that $n\sigma^{-2}\tau^2 = 0.1$ favors complete shrinkage ($\kappa = 1$) and thus we expect more variables to be shrunk a priori. In the next section we discuss an intuitive way of designing a prior distribution for τ based on assumptions about the number of nonzero components in β .

3 THE GLOBAL SHRINKAGE PARAMETER

This section discusses the prior choice for the global hyperparameter τ . We begin with a short note on why we prefer full Bayesian inference for τ over point estimation, and then go on to discuss how we propose to set up the prior $p(\tau)$.

3.1 Full Bayes versus Point Estimation

In principle, one could use a plug-in estimate for τ , obtained either by cross-validation or maximum marginal likelihood (sometimes referred to as “empirical Bayes”). The maximum marginal likelihood estimate has the drawback that it is always in danger of collapsing to $\hat{\tau} = 0$ if the parameter vector happens to be very sparse. Moreover, rather than being computationally convenient, this approach might actually complicate matters as the marginal likelihood is not analytically available for non-Gaussian likelihoods. While cross-validation avoids the latter problem and possibly also the first one, it is computationally less efficient than the full

Bayesian solution and fails to account for the posterior uncertainty. For these reasons we recommend full Bayesian inference for τ , and focus on how to specify the prior distribution.

3.2 Earlier Approaches

Carvalho et al. (2009) also recommend full Bayesian inference for τ , and following Gelman (2006), they propose prior

$$\tau \sim C^+(0, 1), \quad (7)$$

whereas Polson and Scott (2011) recommend

$$\tau | \sigma \sim C^+(0, \sigma^2). \quad (8)$$

Here $C^+(0, a^2)$ denotes the half-Cauchy distribution with location 0 and scale a . If the target variable y is scaled to have marginal variance of 1, unless the noise level σ is very small, both of these priors typically lead to quite similar posteriors. However, as we argue in Section 3.3, there is a theoretical justification for letting τ scale with σ . The main motivation for using a half-Cauchy prior for τ is that it evaluates to a finite positive value at the origin, yielding a proper posterior and allowing even complete shrinkage $\tau \rightarrow 0$, while still having a thick tail which can accommodate a wide range of values. For these reasons, $C^+(0, a^2)$ is a desirable choice when there are enough observations to let τ be identified by data. Still, we show that in several cases one can clearly benefit by choosing the scale a in a more careful manner than simply $a = 1$ or $a = \sigma$, because for most applications these choices place far too much mass for implausibly large values of τ . This point is discussed in Section 3.3. Moreover, in some cases, τ can be so weakly identified by the data that one can obtain even better results by using a more informative and tighter prior, such as half-normal, in place of half-Cauchy (see the experiments in Section 4). van der Pas et al. (2014) study the optimal selection of τ in the model

$$y_i \sim \beta_i + \varepsilon_i, \quad \varepsilon_i \sim N(0, \sigma^2), \quad i = 1, \dots, n. \quad (9)$$

They prove that in such a setting, the optimal value (up to a log factor) in terms of mean squared error and posterior contraction rates in comparison to the true β^* is

$$\tau^* = \frac{p^*}{n}, \quad (10)$$

where p^* denotes the number of nonzeros in the true coefficient vector β^* (assuming such exists). Their proofs assume that $n, p^* \rightarrow \infty$ and $p^* = o(n)$. Model (9) corresponds to setting $\mathbf{X} = \mathbf{I}$ and $D = n$ in the usual

regression model (2). It is unclear whether and how this result could be extended to a more general \mathbf{X} , and how one should utilize this result when p^* is unknown (as it usually is in practice). In the next section, we formulate our method of constructing the prior $p(\tau)$ based on prior beliefs about p^* , and show that if p^* was known, our method would also give rise to result (10), but is more generally applicable.

3.3 Effective Number of Nonzero Coefficients

Consider the prior distribution for the shrinkage factor of the j th regression coefficient, Eq. (6). The mean and variance can be shown to be

$$\mathbb{E}[\kappa_j | \tau, \sigma] = \frac{1}{1 + \sigma^{-1}\tau\sqrt{n}}, \quad (11)$$

$$\text{Var}[\kappa_j | \tau, \sigma] = \frac{\sigma^{-1}\tau\sqrt{n}}{2(1 + \sigma^{-1}\tau\sqrt{n})^2}. \quad (12)$$

A given value for the global parameter τ can be understood intuitively via the prior distribution that it imposes on the effective number of coefficients distinguishable from zero (or effective number of nonzero coefficients, for short)

$$m_{\text{eff}} = \sum_{j=1}^D (1 - \kappa_j). \quad (13)$$

When the shrinkage factors κ_j are close to 0 and 1 (as they typically are for the horseshoe prior), this quantity describes essentially how many active or unshrunk variables we have in the model. It serves therefore as a useful indicator of the effective model size.

Using results (11) and (12), the mean and variance of m_{eff} given τ and σ are given by

$$\mathbb{E}[m_{\text{eff}} | \tau, \sigma] = \frac{\sigma^{-1}\tau\sqrt{n}}{1 + \sigma^{-1}\tau\sqrt{n}}D, \quad (14)$$

$$\text{Var}[m_{\text{eff}} | \tau, \sigma] = \frac{\sigma^{-1}\tau\sqrt{n}}{2(1 + \sigma^{-1}\tau\sqrt{n})^2}D. \quad (15)$$

The expression for the mean (14) is helpful. First, from this expression it is evident that to keep our prior beliefs about m_{eff} consistent, τ must scale as σ/\sqrt{n} . Priors that fail to do so, such as (7), favor models of varying size depending on the noise level σ and the number of data points n . Second, if our prior guess for the number of relevant variables is p_0 , it is reasonable to choose the prior so that most of the prior mass is located near the value

$$\tau_0 = \frac{p_0}{D - p_0} \frac{\sigma}{\sqrt{n}}, \quad (16)$$

which is obtained by solving equation $\mathbb{E}[m_{\text{eff}} | \tau, \sigma] = p_0$. Note that this is typically quite far from 1 or σ , which

are used as scales for priors (7) and (8). For instance, if $D = 1000$ and $n = 200$, then prior guess $p_0 = 5$ gives about $\tau_0 = 3.6 \cdot 10^{-4}\sigma$.

To further develop the intuition about the connection between τ and m_{eff} , it is helpful to visualize the prior imposed on m_{eff} for different prior choices for τ . This is most conveniently done by drawing samples for m_{eff} ; we first draw $\tau \sim p(\tau)$ and $\lambda_1, \dots, \lambda_D \sim C^+(0, 1)$, then compute the shrinkage factors $\kappa_1, \dots, \kappa_D$ from (5), and finally m_{eff} from (13).

Figure 2 shows histograms of prior draws for m_{eff} for some different prior choices for τ , with total number of inputs $D = 10$ and $D = 1000$, assuming $n = 100$ observations with $\sigma = 1$. The first three priors utilize the value τ_0 which is computed from (16) using $p_0 = 5$ as our hypothetical prior guess for the number of relevant variables. Fixing $\tau = \tau_0$ results in a nearly symmetric distribution around p_0 , while a half-normal prior with scale τ_0 yields a skewed distribution favoring solutions with $m_{\text{eff}} < p_0$ but allowing larger value to also be accommodated. The half-Cauchy prior behaves similarly to the half-normal, but results in a distribution with a much thicker tail giving substantial mass also to values much larger than p_0 when D is large. Figure 2 also illustrates why the prior $\tau \sim C^+(0, 1)$ is often a dubious choice: it places far too much mass on large values of τ , consequently favoring solutions with most of the coefficients unshrunk. Thus when only a small number of the variables are relevant – as we typically assume – this prior results in sensible inference only when τ is strongly identified by data. Note also that, if we changed the value of σ or n , the first three priors for τ would still impose the same prior for m_{eff} , but this is not true for $\tau \sim C^+(0, 1)$.

This way, by studying the prior for m_{eff} , one can easily choose the prior for τ based on the beliefs about the number of nonzero parameters. Because the prior beliefs can vary substantially for different problems and the results depend on the information carried by the data, there is no globally optimal prior for τ that works for every single problem. Some recommendations, however, will be given in Section 5 based on these theoretical considerations and experiments presented in Section 4.

We shall conclude this section by pointing out a connection between our reference value (16) and the oracle result (10) for the simplified model (9). As pointed out in the last section, model (9) corresponds to setting $\mathbf{X} = \mathbf{I}$ (which implies $n = D$ and $\mathbf{X}^T\mathbf{X} = \mathbf{I}$) in the usual regression model (2). Using this fact and repeating the steps needed to arrive at (16), we get

$$\tau_0 = \frac{p_0}{D - p_0} \sigma. \quad (17)$$

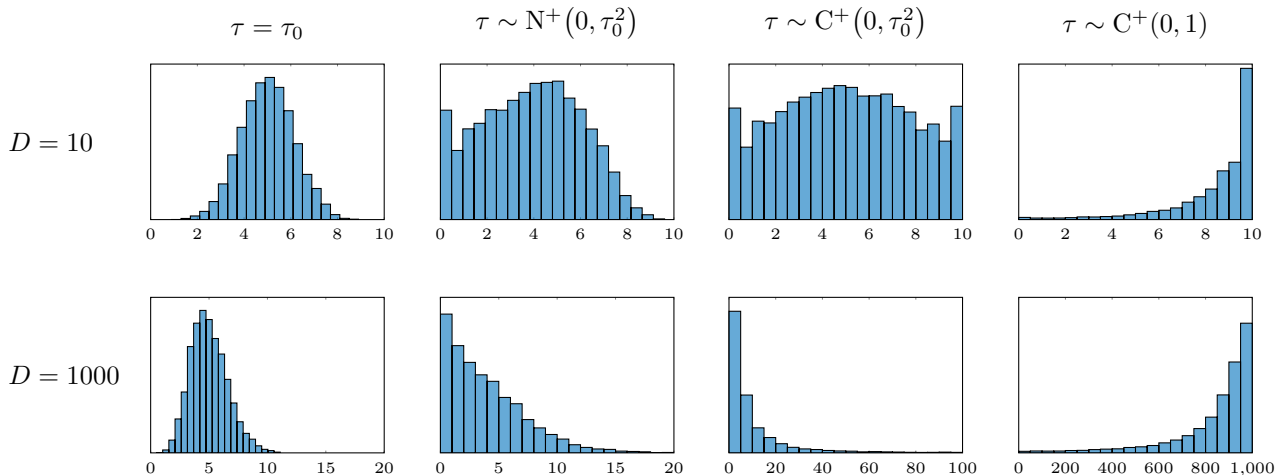


Figure 2: Histograms of prior draws for m_{eff} (effective number of nonzero regression coefficients, Eq. (13)) imposed by different prior choices for τ , when the total number of input variables is $D = 10$ and $D = 1000$. τ_0 is computed from formula (16) assuming $n = 100$ observations with $\sigma = 1$ and $p_0 = 5$ as the prior guess for the number of relevant variables. Note the varying scales on the horizontal axes in the bottom row plots.

Suppose now that we select $p_0 = p^*$, that is, our prior guess is oracle. Using the same assumptions as van der Pas et al. (2014), namely that $n, p^* \rightarrow \infty$ and $p^* = o(n)$, and additionally that $\sigma = 1$, we get $\tau_0 \rightarrow p^*/D = \tau^*$. This result is natural, as it means it is optimal to choose τ so that the imposed prior for the effective number of nonzero coefficients m_{eff} is centered at the true number of nonzeros p^* . This further motivates why m_{eff} is a useful quantity.

3.4 Non-Gaussian Likelihood

When the observation model is non-Gaussian, the exact analysis from Section 3.3 is analytically intractable. We can, however, perform the analysis using a Gaussian approximation to the likelihood. Using the second order Taylor expansion for the log likelihood, the approximate posterior for the regression coefficients given the hyperparameters becomes (see the supplementary material for details)

$$p(\boldsymbol{\beta} | \boldsymbol{\Lambda}, \tau, \phi, \mathcal{D}) \approx \mathcal{N}(\boldsymbol{\beta} | \bar{\boldsymbol{\beta}}, \boldsymbol{\Sigma}),$$

$$\bar{\boldsymbol{\beta}} = \tau^2 \boldsymbol{\Lambda} \left(\tau^2 \boldsymbol{\Lambda} + (\mathbf{X}^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{X})^{-1} \right)^{-1} \hat{\boldsymbol{\beta}},$$

$$\boldsymbol{\Sigma} = (\tau^{-2} \boldsymbol{\Lambda}^{-1} + \mathbf{X}^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{X})^{-1},$$

where $\tilde{\mathbf{z}} = (\tilde{z}_1, \dots, \tilde{z}_n)$, $\tilde{\boldsymbol{\Sigma}} = \text{diag}(\tilde{\sigma}_1^2, \dots, \tilde{\sigma}_n^2)$ and $\hat{\boldsymbol{\beta}} = (\mathbf{X}^\top \tilde{\boldsymbol{\Sigma}}^{-1} \mathbf{X})^{-1} \mathbf{X}^\top \tilde{\boldsymbol{\Sigma}}^{-1} \tilde{\mathbf{z}}$ (assuming the first inverse exists). Here ϕ denotes the possible dispersion parameter and $(\tilde{z}_i, \tilde{\sigma}_i^2)$ the location and variance for the i th Gaussian pseudo-observation. The fact that some of the observations are more informative than others ($\tilde{\sigma}_i^2$ vary) makes further simplification somewhat difficult.

To proceed, we make the rough assumption that we can replace each $\tilde{\sigma}_i^2$ by a single variance term $\tilde{\sigma}^2$. Assuming further that the covariates are uncorrelated with zero mean and unit variance (as in Sec. 3.3), the posterior mean for the j th coefficient satisfies $\bar{\beta}_j = (1 - \kappa_j) \hat{\beta}_j$ with shrinkage factor given by

$$\kappa_j = \frac{1}{1 + n \tilde{\sigma}^{-2} \tau^2 \lambda_j^2}. \quad (18)$$

The discussion in Section 3.3 therefore also approximately holds for the non-Gaussian observation model, except that σ^2 is replaced by $\tilde{\sigma}^2$. Still, this leaves us with the question, which value to choose for $\tilde{\sigma}^2$ when using this result in practice?

We consider binary classification here as an example. It can be shown (see the supplementary material) that for the logistic regression

$$p(y_i = 1 | f_i) = s(f_i) = \frac{1}{1 + \exp(-f_i)}, \quad f_i = \boldsymbol{\beta}^\top \mathbf{x}_i,$$

for those points that lie on the classification boundary, we have $\tilde{\sigma}_i^2 = 4$, and for others $\tilde{\sigma}_i^2 > 4$. For this reason we propose to use the results of Section 3.3 as they are, by plugging in $\sigma = 2$. In practice this introduces some error, but the good thing is that we know in which way the results are biased. For instance, because the true (unknown) effective noise deviation would be $\sigma > 2$, the prior mean (14) using $\sigma = 2$ is actually an *overestimate* of the true value. Thus also when using the result (16), we tend to favor slightly too small values of τ and thus also solutions with slightly less than p_0 nonzero coefficients. In practice we observe that this approach,

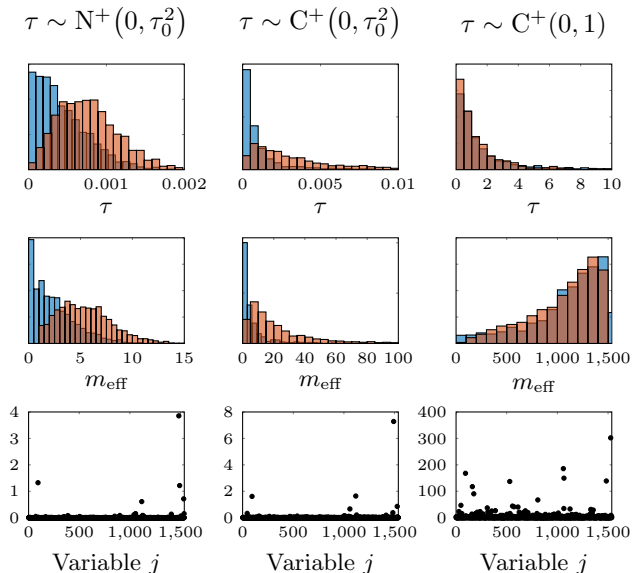


Figure 3: *Ovarian dataset*: Histograms of prior (blue) and posterior (red) samples for τ (top row) and m_{eff} (middle row), and absolute values of the posterior means for the regression coefficients $|\hat{\beta}_j|$ (bottom row) imposed by three different choices of prior for τ . τ_0 is computed from (16) using $p_0 = 3$ and $\sigma = 2$ (for reasons explained in Sec. 3.4). Note the changing scales on the horizontal axes in the top and middle row plots.

though relying on several crude approximations, is still useful and gives reasonably accurate results.

Finally, we note that similar approximate substitute values for σ to be used in equations of Section 3.3 could also be derived for other link functions and observation models, but due to limited space, we do not focus on them in this paper.

3.5 Other Shrinkage Priors

Our approach could also be used with shrinkage priors other than the horseshoe, as long as the prior can be written as scale mixtures of Gaussians like (3) with some prior for the local scales $\lambda_j \sim p(\lambda_j)$. An example of such is the hierarchical shrinkage prior – a computationally more robust generalization of the horseshoe obtained by replacing the half-Cauchy priors in (3) by half- t distributions with some small degrees of freedom (Piironen and Vehtari, 2015). However, the closed form equations in Section 3.3 apply only to the horseshoe. Depending on the choice of $p(\lambda_j)$, corresponding analytical results may or may not be available, but as long as one is able to sample both from $p(\tau)$ and $p(\lambda_j)$, it is always easy to draw samples from the prior distribution for m_{eff} .

Table 1: Summary of the datasets, number of predictor variables D and dataset size n . Classification datasets are all binary.

Dataset	Type	D	n
Ovarian	Classification	1536	54
Colon	Classification	2000	62
Prostate	Classification	5966	102
ALLAML	Classification	7129	72
Corn (4 targets)	Regression	700	80

4 EXPERIMENTS

We illustrate the importance of the prior choice for τ with real world examples (see the supplementary material for an additional experiment on synthetic data). The datasets are summarized in Table 1 and can be found online.¹ The first four are microarray cancer datasets, some of which have been used as benchmark datasets by several authors (see e.g., Hernández-Lobato et al., 2010, and references therein). The Corn data (Chen and Martin, 2009) consists of three sets of predictors (‘m5’, ‘mp5’, ‘mp6’) and four responses. We use the ‘mp5’ input data and consider prediction for all four targets.

A Gaussian linear model was used for the regression tasks and logistic regression for classification. The horseshoe prior was employed for the regression coefficients and a weakly informative prior $\beta_0 \sim N(0, 10^2)$ for the intercept. The noise variance was given the standard prior $p(\sigma^2) \propto 1/\sigma^2$ in the regression problems. All considered models were fitted using Stan (version 2.12.0, codes in the supplementary material), running 4 chains, 1000 samples each, first halves discarded as warmup.

Effect on parameter estimates We first consider the Ovarian dataset as a representative example of how much the prior choice $p(\tau)$ can affect the parameter estimates. We fitted the model to the data using three different priors for the global parameter; $\tau \sim N^+(0, \tau_0^2)$, $\tau \sim C^+(0, \tau_0^2)$, and $\tau \sim C^+(0, 1)$, where τ_0 is computed from (16) using $p_0 = 3$ as our prior guess for the number of relevant variables and $\sigma = 2$ (for reasons discussed in Section 3.4).

Figure 3 shows prior and posterior samples for τ and m_{eff} , and the absolute values of the posterior means for the regression coefficients, for the three prior choices. The results for $\tau \sim C^+(0, 1)$ illustrate how weakly τ

¹Colon, Prostate and ALLAML: <http://featureselection.asu.edu/datasets.php>; Ovarian data: request from the first author if needed; Corn data: <http://software.eigenvector.com/Data/Corn/index.html>

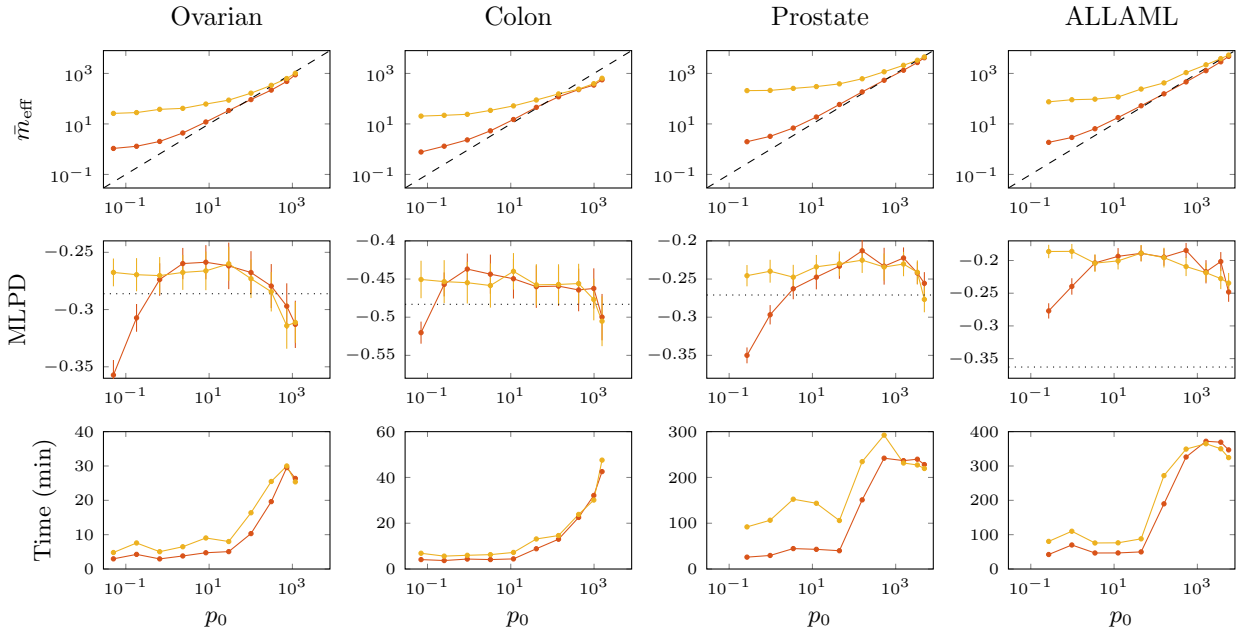


Figure 4: *Classification datasets*: Posterior mean for m_{eff} , mean log predictive density (MLPD) on test data (\pm one standard error), and computation time for two priors for the global hyperparameter: $\tau \sim N^+(0, \tau_0^2)$ (red), and $\tau \sim C^+(0, \tau_0^2)$ (yellow), where τ_0 is computed from (16) varying p_0 (horizontal axis). For each curve, the largest p_0 corresponds to $\tau_0 = 1$. Dotted line in the middle row plots denotes the MLPD for LASSO. All the results are averaged over 50 random splits into training and test sets.

is identified by the data: there is very little difference between the prior and posterior samples for τ and consequently for m_{eff} , and thus this “non-informative” prior actually has a strong influence on the posterior inference. The prior results in a severe under-regularization and implausibly large absolute values for the logistic regression coefficients (magnitude in the hundreds).

Replacing the scale of the half-Cauchy with τ_0 , reflecting a more sensible guess for the number of relevant variables, has a substantial effect on the posterior inference: the posterior mass for m_{eff} becomes concentrated on more reasonable values and the magnitude of the regression coefficients more sensible. Replacing the half-Cauchy by half-normal places a tighter constraint on τ and consequently also on m_{eff} . Either of these two choices is a substantial improvement over $\tau \sim C^+(0, 1)$, and this is also seen in the predictive accuracy (to be discussed in a moment).

A potential explanation for why τ and therefore m_{eff} are not strongly identified here is that there are a lot of correlations in the data. For instance, the predictor $j = 1491$ which appears relevant based on its regression coefficient, has an absolute correlation of at least $|\rho| = 0.5$ with 314 other variables, out of which 65 correlations exceed $|\rho| = 0.7$. This indicates that there are a lot of potentially relevant but redundant predictors in the data. Note also that even though our theoretical

treatment for m_{eff} assumes uncorrelated variables, the correlations do not seem to be a marked problem in this sense.

Prediction accuracy and computation time To investigate the effect of the prior choice $p(\tau)$ on the prediction accuracy, we splitted each dataset into two halves, using one fifth of the data as a test set. All the results were then averaged over 50 such random splits into training and test sets. For the regression problems, we carried out the tests for priors $\tau | \sigma \sim N^+(0, \tau_0^2)$ and $\tau | \sigma \sim C^+(0, \tau_0^2)$ with τ_0 given by Eq. (16) for various p_0 . The experiments for the classification datasets were done using the same priors by plugging in $\sigma = 2$. We also compared the prediction accuracies to LASSO with the regularization parameter tuned by 10-fold cross-validation, and noise variance in regression computed by dividing the sum of squared residuals by $n - p_{\text{lasso}}$, where p_{lasso} is the size of the LASSO active set (Reid et al., 2016).

Figures 4 and 5 show the effect of the prior choice on the posterior mean for m_{eff} , test prediction accuracy, and computation time (wall time). The classification datasets illustrate a clear benefit from using even a crude prior guess for the number of relevant variables p_0 : there is a relatively wide range of values for p_0 which yield simpler models (smaller \bar{m}_{eff}), better predictive accuracy, and substantially reduced computa-

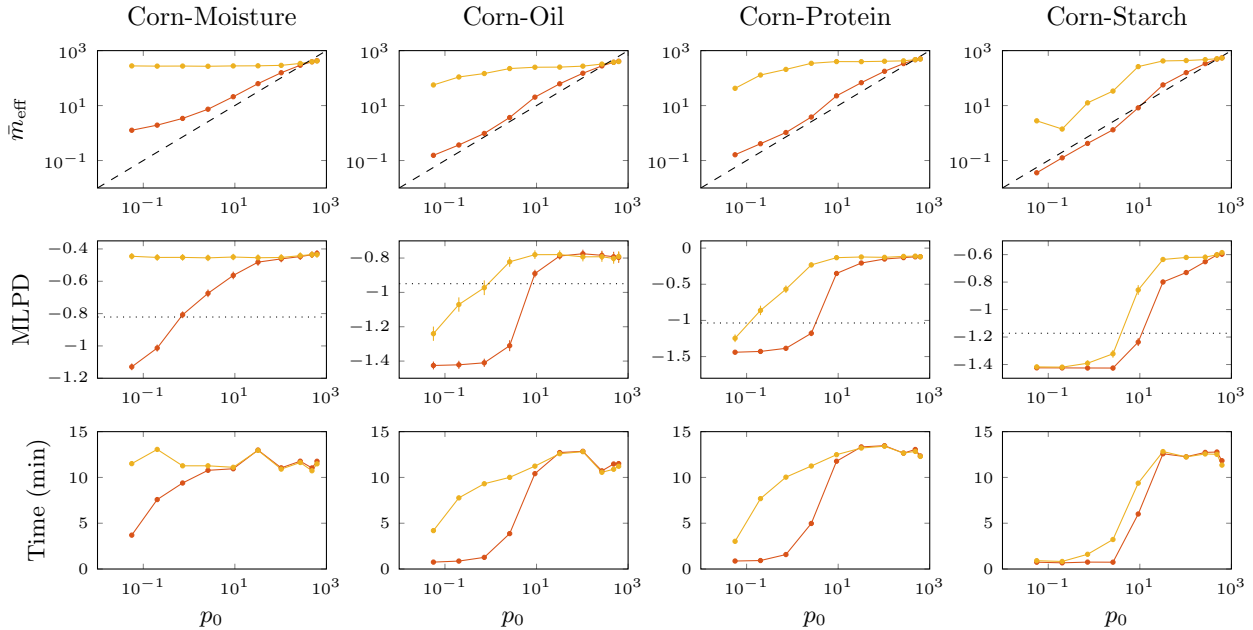


Figure 5: *Regression datasets*: Same as in Figure 4 but for the regression datasets. For each curve, the largest p_0 corresponds to $\tau_0 = \sigma$.

tion time when compared to using $\tau_0 = 1$ (the largest p_0 for each curve), regardless of whether half-normal or half-Cauchy prior is used. The reduced computation time is due to tighter concentration of the posterior mass which aids the sampling greatly. The half-Cauchy prior performs better than the half-normal prior with small p_0 values (both for classification and regression datasets), because the thick tails allow τ to get much larger values than τ_0 , but for too small values of p_0 , even the half-Cauchy can lead to bad results (Corn-Oil – Corn-Starch). On the other hand, when p_0 happens to be well chosen, the tighter half-normal prior can yield a simpler model (smaller m_{eff}) with equally good or even better predictive accuracy (classification datasets).

For any reasonably selected prior, the horseshoe consistently outperforms LASSO in a pairwise comparison, but note that for Ovarian and Colon, $\tau_0 = 1$ (largest p_0) actually yields worse results. In terms of classification accuracy, the differences are smaller, but in regression, the horseshoe also performs better when measured by the mean squared error (not shown). A clear advantage for LASSO, on the other hand, is that it is hugely faster, with computation time of less than a second for these problems.

5 CONCLUSION AND RECOMMENDATIONS

This paper has discussed the prior choice for the global shrinkage parameter in the horseshoe prior for sparse

Bayesian regression and classification. We have shown that the previous default choices are often dubious based on their tendency to favor solutions with too many parameters unshrunk. The experiments show that for many datasets, one can obtain clear improvements – in terms of better parameter estimates, prediction accuracy and computation time – by coming up even with a crude guess for the number of relevant variables and transforming this knowledge into a prior for τ using our proposed framework. The results also show that there is no globally optimal prior choice that would perform best for all problems, which emphasizes the relevance of the prior choice.

As a new default choice for regression, we recommend $\tau | \sigma \sim C^+(0, \tau_0^2)$, where τ_0 is computed from (16) using the prior guess p_0 for the number of relevant variables. For logistic regression, an approximately equivalent choice is obtained by plugging in $\sigma = 2$ (see Sec. 3.4). This choice seems to perform well unless p_0 is very far from the optimal. If the results still indicate that more regularization is needed, we recommend investigating the imposed prior on m_{eff} as discussed in Section 3.3, and changing $p(\tau)$ so that the prior for m_{eff} corresponds to our beliefs as closely as possible.

Acknowledgements

We thank Andrew Gelman, Jonah Gabry and Eero Siivola for helpful comments to improve the manuscript. We also acknowledge the computational resources provided by the Aalto Science-IT project.

References

- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2009). Handling sparsity via the horseshoe. In *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, pages 73–80. 1, 2, 3
- Carvalho, C. M., Polson, N. G., and Scott, J. G. (2010). The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480. 1, 2
- Chen, T. and Martin, E. (2009). Bayesian linear regression and variable selection for spectroscopic calibration. *Analytica Chimica Acta*, 631:13–21. 6
- Datta, J. and Ghosh, J. K. (2013). Asymptotic properties of Bayes risk for the horseshoe prior. *Bayesian Analysis*, 8(1):111–132. 2
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models. *Bayesian Analysis*, 1(3):515–533. 3
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B. (2013). *Bayesian Data Analysis*. Chapman & Hall, Third edition. 10
- George, E. I. and McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889. 1
- Hernández-Lobato, D., Hernández-Lobato, J. M., and Suárez, A. (2010). Expectation propagation for microarray data classification. *Pattern Recognition Letters*, 31(12):1618–1626. 1, 6
- Hernández-Lobato, J. M., Hernández-Lobato, D., and Suárez, A. (2015). Expectation propagation in linear regression models with spike-and-slab priors. *Machine Learning*, 99:437–487. 1
- Hoeting, J. A., Madigan, D., Raftery, A. E., and Volinsky, C. T. (1999). Bayesian model averaging: a tutorial. *Statistical Science*, 14(4):382–417. 1
- Mitchell, T. J. and Beauchamp, J. J. (1988). Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1036. 1
- Peltola, T., Havulinna, A. S., Salomaa, V., and Vehtari, A. (2014). Hierarchical Bayesian survival analysis and projective covariate selection in cardiovascular event risk prediction. In *Proceedings of the Eleventh UAI Bayesian Modeling Applications Workshop*, volume 1218 of *CEUR Workshop Proceedings*, pages 79–88. 12
- Piironen, J. and Vehtari, A. (2015). Projection predictive variable selection using Stan+R. *arXiv:1508.02502*. 6, 12
- Polson, N. G. and Scott, J. G. (2011). Shrink globally, act locally: sparse Bayesian regularization and prediction. In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian statistics 9*, pages 501–538. Oxford University Press, Oxford. 1, 2, 3
- Reid, S., Tibshirani, R., and Friedman, J. (2016). A study of error variance estimation in Lasso regression. *Statistica Sinica*, 26(1):35–67. 7
- Stan Development Team (2016). Stan modeling language users guide and reference manual, version 2.12.0. <http://mc-stan.org>. 1
- Tibshirani, R. (1996). Regression shrinkage and selection via the Lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, 58(1):267–288. 1
- Titsias, M. K. and Lázaro-Gredilla, M. (2011). Spike and slab variational inference for multi-task and multiple kernel learning. In *Advances in Neural Information Processing Systems 24*, pages 2339–2347. 1
- van der Pas, S. L., Kleijn, B. J. K., and van der Vaart, A. W. (2014). The horseshoe estimator: posterior concentration around nearly black vectors. *Electronic Journal of Statistics*, 8(2):2585–2618. 2, 3, 5, 10, 11
- Zou, H. and Hastie, T. (2005). Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society. Series B (Methodological)*, 67(2):301–320. 1