

Simultaneous parameter estimation and variable selection via the logit-normal continuous analogue of the spike-and-slab prior

Thomson, W.; Jabbari, S.; Taylor, A. E.; Arlt, W.; Smith, David

DOI:
[10.1098/rsif.2018.0572](https://doi.org/10.1098/rsif.2018.0572)

License:
None: All rights reserved

Document Version
Peer reviewed version

Citation for published version (Harvard):
Thomson, W, Jabbari, S, Taylor, AE, Arlt, W & Smith, D 2019, 'Simultaneous parameter estimation and variable selection via the logit-normal continuous analogue of the spike-and-slab prior', *Journal of The Royal Society Interface*, vol. 16, no. 150, 20180572, pp. 20180572. <https://doi.org/10.1098/rsif.2018.0572>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:
Forthcoming in Journal of the Royal Society
DOI, to follow

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Simultaneous Parameter Estimation and Variable Selection via the LN-CASS Prior

W. Thomson,¹ S. Jabbari,^{1,2} A. E. Taylor,^{2,3} W. Arlt,^{4,3} and D. J. Smith^{1,4,3}

¹*School of Mathematics, University of Birmingham, UK^a*

²*Institute of Microbiology & Infection, University of Birmingham, UK*

³*Centre for Endocrinology, Diabetes and Metabolism, Birmingham Health Partners, Birmingham, B15 2TT, United Kingdom*

⁴*Institute of Metabolism & Systems Research, University of Birmingham, UK*

(Dated: 23 November 2018)

We introduce a Bayesian prior distribution, the Logit-Normal continuous analogue of the spike-and-slab (LN-CASS), which enables flexible parameter estimation and variable/model selection in a variety of settings. We demonstrate its use and efficacy in three case studies – a simulation study and two studies on real biological data from the fields of metabolomics and genomics. The prior allows the use of classical statistical models, which are easily interpretable and well-known to applied scientists, but performs comparably to common machine learning methods in terms of generalisability to previously unseen data.

I. INTRODUCTION

Often in real-world regression problems, we are faced with a situation in which we have a large number of potentially irrelevant predictors, possibly even greater than the number of observations. This so-called $p \gg n$ problem is especially prevalent in the biological and medical sciences with the advent of high-throughput experimental methods and an increasing focus on synthesising knowledge of molecular details into models predicting much lower-dimensional observable outcomes. Regularisation and shrinkage methods aim to reduce the influence of the inherent noise in such problems and provide sparse estimated parameter vectors, essentially performing simultaneous variable selection and parameter fitting. The motivation is two-fold. Firstly, regularisation aims to more robustly distinguish strong from weak effects, i.e. more reliably identify the genuine driving forces of the process of interest. Secondly, we wish to reduce overfitting to improve the generalisability of our models. The performance of our method in both of these respects is demonstrated below.

The most common means of dealing with the $p \gg n$ problem is the LASSO^{1,2}, whose ability to induce genuine sparsity (i.e. estimates of *exactly* zero) and whose computationally efficient implementation make it attractive for general-purpose regularised regression. A

number of Bayesian analogues of the LASSO and other penalised likelihood methods have been proposed in order to more fully account for the uncertainty in parameter estimates, which we contend is particularly important in small n /large p problems, and to tackle the tendency of the LASSO to underestimate large effects^{3–5}.

Some authors have focused on the subset of such problems in which the predictors have a known grouping structure, for example in problems from genetics in which the groups correspond to known regulatory networks⁶. This has led to the development of both penalised likelihood⁷ and Bayesian^{8,9} modifications of common shrinkage methods.

In this paper, we present a new shrinkage prior – the Logit-normal continuous analogue of the spike-and-slab (LN-CASS) – based on a Logit-Normal relaxation of the Bernoulli distribution used in the spike-and-slab prior¹⁰. The spike-and-slab is considered the gold-standard of Bayesian variable selection¹¹, but is computationally intractable in practice due to its combinatorial complexity.

The LN-CASS prior has the advantage that its intuitive formulation allows it to be simply extended to almost any hierarchical situation – two of which are covered below – allowing the modeller to tailor the specifications of common statistical models to favour ‘simpler’ models in a variety of senses. Below we structure our models to favour first homogeneous groups of predictors before allowing within-group heterogeneity (simulation study) and to favour purely linear effects before nonlinear effects (metabolomics study), as well as applying the method in its simplest form to shrink logistic regression coefficients (microarray case study). The Bayesian formalism ‘allows the data to decide’ the appropriate level of complexity through the likelihood function.

In the simulation study, the LN-CASS prior empirically appears robust to group misspecification, and outperforms the horseshoe prior³, the LASSO¹ and the sparse group LASSO⁷. Additionally, we apply the LN-CASS prior to a real-life classification task, in which we aim to distinguish benign from malignant adrenal tumours. Our method leads to an out-of-sample predictive performance comparable to state-of-the-art machine learning methods, but offers more interpretable

^aElectronic mail: wmt095@bham.ac.uk

results. We also use the method to build a predictive model of colon cancer malignancy using the well-known Colon dataset¹². The code is available as a collection of R functions (see Data Accessibility).

II. RESULTS

To illustrate the utility of the LN-CASS prior, we conducted three comparative studies with the intention of assessing its two primary functions: identification of genuinely non-zero effects and improving out-of-sample performance by reducing overfitting. Additionally, we chose two of the three settings to highlight the flexibility of the approach, in particular its capacity to include known group structure (simulation, section C.1) and to perform nonparametric regression (metabolomics, section C.2). These two extensions are by no means exhaustive, but are illustrative of the myriad possible areas of application (see Discussion).

A. The LN-CASS Prior

We now provide a brief outline of the LN-CASS prior. Mathematical details are available in section 1 of the Electronic Supplementary Material (ESM).

The fundamental motivation for developing the LN-CASS prior is to provide a computationally tractable alternative to the theoretical gold-standard of Bayesian variable selection, the spike-and-slab prior.

The spike-and-slab prior is based on the simple idea that, *a priori*, we believe each parameter has some non-zero probability of being zero, and the rest of the probability mass is assigned to other plausible parameter values (often uniformly). This hard zero/non-zero distinction introduces a discrete component into our prior beliefs and renders the practical use of the prior combinatorially intractable – we need to visit 2^p parameter combinations in order to adequately cover the parameter space where p is the number of parameters in our model. For even moderately sized problems, this complexity renders the spike-and-slab impractical. Indeed, this combinatorial complexity is the same problem faced by the frequentist ‘best subset selection’, in which every possible subset of parameters is compared and the best performing subset is chosen.

By constructing a fully continuous approximation to the mixed spike-and-slab prior, we enable greatly improved sampling efficiency at the cost of relaxing the hard distinction between zero and non-zero parameters. The LN-CASS prior accomplishes this relaxation by replacing the discrete Bernoulli distribution in the mixture formulation of the spike-and-slab with a Logit-normal distribution (see the sections 1 & 2 of the ESM for details). The Logit-normal distribution, with suitable parameter choices, is a U-shaped distribution on $(0, 1)$, assigning most of its mass to values close to the endpoints

(figure 1). The reason for choosing the Logit-Normal distribution for this purpose over the similar and more common Beta distribution is that it can be expressed as a transformation of standard normal random variables, which greatly aids the convergence properties of our sampler. Indeed, models can be specified purely in terms of parameters with (conditionally) standard normal prior distributions.

We interpret the values of the Logit-normal random variable as approximate variable inclusion probabilities, which allows simple propagation of these probabilities through a hierarchical prior structure. For example, in the simulation study below we impose a hierarchical prior structure in which we favour first exclusion of whole groups of variables, then allow inclusion of groups with a shared parameter, and finally allow groups with differing parameters. In the metabolomics case study below, we utilise this prior structure to favour linear effects first, before allowing non-linear effects if the data support such effects strongly enough. This corresponds to imposing a hierarchy on the complexity of the model and allows us to refine exactly how we control model complexity.

B. Performance measures

The main measure of performance we employ is the area under the Receiver Operating Characteristic (ROC) curve (AUC). The ROC curve is a plot of the false-positive rate (specificity) against the true-positive rate (sensitivity) as the probability threshold for classifying a prediction as positive or negative is varied. The AUC is interpretable as the probability of successfully distinguishing a positive result from a negative result, i.e. the probability of correctly assigning a larger predicted value to a positive case than a negative case. An AUC of 0.5 corresponds to a model that simply uses the class proportions as a prediction, while an AUC of 1 corresponds to a classifier which perfectly distinguishes positive and negative cases at some threshold. We use the AUC to quantify the trade-off between false- and true-positives in two settings. Firstly, in the simulation study the AUC is used to quantify the degree to which each method uncovers the correct ordering of ground-truth parameter values – the degree to which genuinely small parameters are estimated to be small, and large parameters to be large. In this first case, positive and negative results are related to the sparsity pattern of the ground-truth parameter values: if a parameter has a non-zero ground truth value, it is assigned a positive result, otherwise it is assigned a negative result. The AUC, therefore, measures the probability that truly non-zero parameters are estimated to be larger (in absolute value) than truly zero parameters. Secondly, in the metabolomics and microarray case studies, we use the AUC in the more conventional setting of quantifying the out-of-sample performance of a classifier.




FIG. 1: (a) The Logit-Normal distribution with $\mu = 0$ and σ given by (blue, orange, green) = (2.5, 5, 50). (b) The LN-CASS priors induced by the Logit-Normal distributions of (a).

To quantify the agreement between the estimated and true parameters in the simulation study, we use the mean absolute error (MAE). The MAE is simply the average distance of the estimated from the true parameters.

C. Applications

We now present the results of three case studies to evaluate the comparative ability of the LN-CASS prior to perform its two main duties – sparse parameter estimation and improving out of sample performance. In the first case study, we attempt to recover ground-truth parameters in a simulation study in which we impose a known grouping structure in the predictors. In the second case study, we use real-world metabolomics data¹³ to build a predictive model of adrenal tumour malignancy. In the third, we apply the LN-CASS prior in the context of Bayesian logistic regression to the well-known colon cancer dataset¹².

1. Simulation Study (grouped predictors)

The motivation for this case study is to illustrate the ability of the LN-CASS prior to penalise not only model complexity in terms of the number of parameters, but also the granularity of the model. Such a formulation might be applied when there is some ‘subset’ or ‘tree-like’ structure in the predictors. For example, in immunological applications, cell subsets are often nested – T-cells are subdivided into $CD4^+$ and $CD8^+$ T-cells, which in turn are subdivided into naive and memory

subsets. The grouped LN-CASS prior favours within group homogeneity, essentially favouring less granular models, i.e. a model using total T-cell counts would be favoured over a model using $CD4^+$ and $CD8^+$ subsets as predictors.

We generated a simulated dataset of $n = 100$ observations from the linear regression model

$$y_i = \beta_0 + \mathbf{X}_i\boldsymbol{\beta} + \varepsilon_i, \quad (1)$$

for three different settings with grouped predictors, i.e. where pre-specified groups share mostly the same or similar parameter values (table 1, ESM). The matrix \mathbf{X} was sampled from a Unit Latin Hypercube. The ε_i were chosen to be i.i.d. zero-mean Gaussian.

We then fit the model in R with the following methods for each of the three settings ($p = 20, 70, 120$): Group LN-CASS, LASSO, Horseshoe, Sparse Group LASSO and Ordinary Least Squares. Ordinary least squares was tested only for the $p < n$ cases because the problem is not well-defined when $p > n$. For code see Data Accessibility. Details of the grouped LN-CASS prior are available in section 2.1 of the ESM, and full details of all methods can be found in the code (see Data Accessibility).

LN-CASS substantially outperforms all of the other methods in recovering the ground-truth parameters and correctly identifying zero parameters (figure 2).

2. Steroid metabolomics and adrenal tumour malignancy (hierarchical GAM)

We applied a hierarchical version of the LN-CASS prior to clinical data regarding the concentrations of

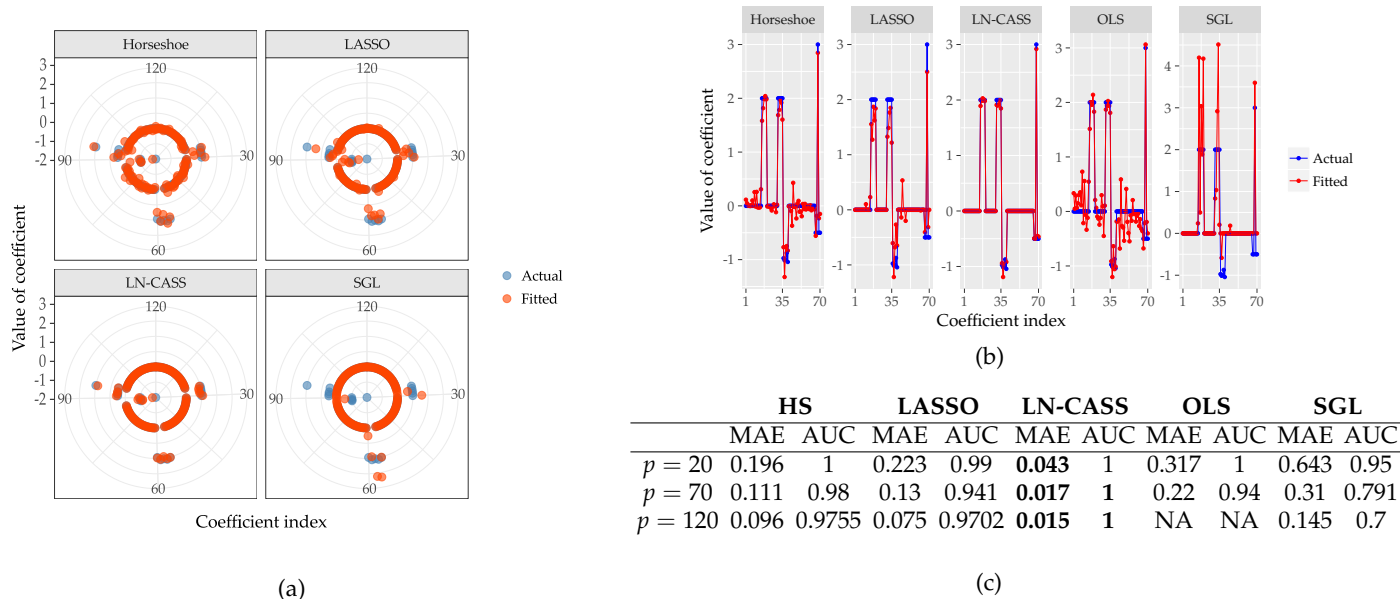


FIG. 2: Agreement between ground truth and estimated parameters for the simulation study in the (a) $p = 120$ case (b) $p = 70$ case; (c) performance measures for each method. Acronyms: HS – Horseshoe; OLS – Ordinary Least Squares; SGL – Sparse Group LASSO

metabolites in the urine of patients with two different adrenal tumours. The task was to predict the tumour type based on the metabolites, and to do this we used a generalised additive model (GAM) with logit link. The implementation of the prior in a hierarchical fashion here was strongly inspired by a recent paper by Griffin and Brown¹⁴.

The GAM we implemented represented the effect of each covariate as the sum of linear basis functions. We imposed a hierarchy through the LN-CASS prior which favoured firstly the complete removal of a covariate from the model, then inclusion of a purely linear effect, and finally allowed each of the basis functions to be used to construct a non-linear effect (for details, see the ESM, subsection 2.2).

The dataset consisted of 158 measurements of 32 covariates¹³ collected as part of the EURINE-ACT study, with 45 positive cases (malignant adrenal tumours). All of the covariates are measurements of steroid concentrations in urine samples taken from each of the patients. There is a small proportion of missing data (up to 7% of a covariate’s measurements), which we imputed via the `mice()` function in R¹⁵. We then $\log(1+x)$ transformed all of the data because many of the predictors spanned several orders of magnitude. We subsequently scaled all covariates to lie in the interval $[0, 1]$.

We compared the classification performance of our hierarchical GAM with the performance of the following methods: Support Vector Machine (SVM), neural network (NN), random forest (RF) and elastic net (a modified version of the LASSO). Classification performance was measured using the mean AUC over 16×10 -fold

cross-validated runs. The results are presented in fig. 3 (a).

All of the methods perform comparably in terms of out-of-sample predictive performance, with the neural network performing the best and LN-CASS second in terms of both the mean and variability (inter-quartile range) of cross-validated AUCs. The authors are not aware of an appropriate and well-established statistical test to formalise the comparative performances of each method given the unequal variances, clear non-normality, and obvious dependency between samples for a given method. However, the Kruskal-Wallis test with a *post hoc* Dunn test (and appropriate multiplicity correction) provides a non-parametric test for stochastic dominance (i.e. the tendency of values from one distribution to be larger than values from the other). We used two multiplicity corrections, both of which account for positive dependency (i.e. the tendency of large AUCs to be correlated within cross-validation folds). Using the Benjamini-Hochberg¹⁶ correction, the only null hypotheses to be rejected at 95% significance levels were that the distribution of AUCs for the neural network stochastically dominates those for the Elastic Net and the Random Forest (adjusted p -values 0.0344 and 0.0203, respectively). Using the Benjamini-Yekutieli¹⁷ correction, no null hypotheses were rejected; that is, no significant differences were found between the distributions in terms of stochastic dominance. Note that the Benjamini-Yekutieli correction allows for arbitrary dependencies.

The results suggest that the out-of-sample performance of the hierarchical GAM with LN-CASS prior is comparable with that of state-of-the-art machine learn-

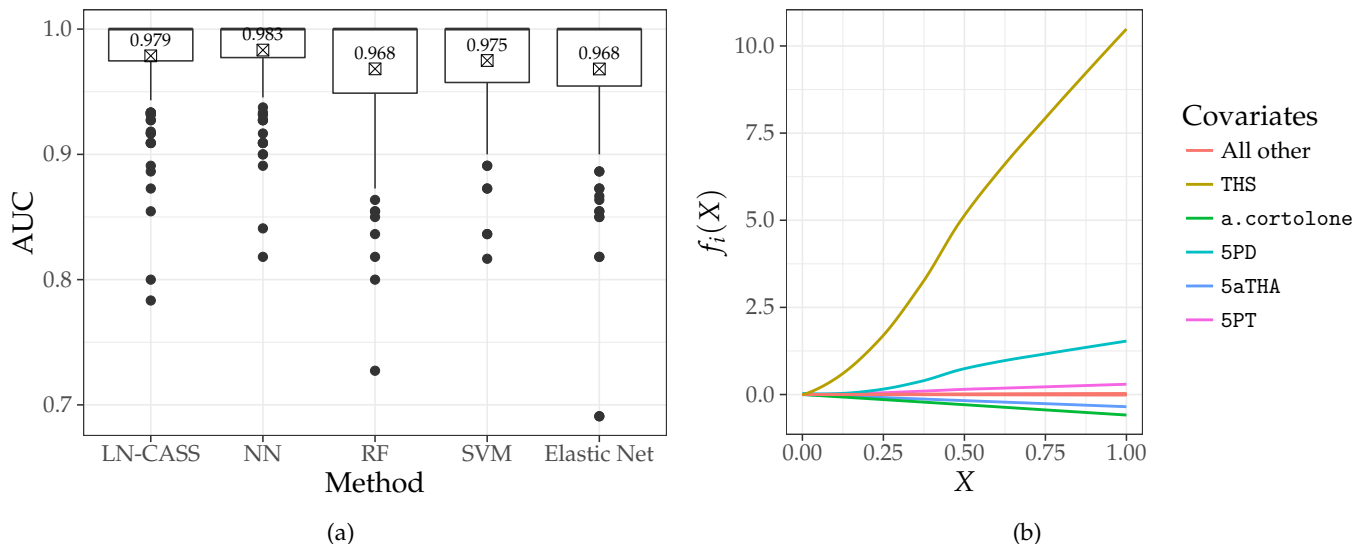


FIG. 3: Metabolomics case study. (a) Boxplots of AUCs for each method computed via 16×10 -fold cross-validation; (b) estimated mean functions f_i from the LN-CASS hierarchical GAM. Functions have been smoothed for presentation purposes with a LOESS smoother using a small span.

ing methods. We argue that this performance, in conjunction with the accuracy with which LN-CASS recovers ‘true’ parameters and offers more classically interpretable results make it a valuable addition to the shrinkage and regularisation toolbox for applied scientists.

The recovered effects for each of the metabolites are presented in fig. 3 (b), as estimated from the full dataset. Clearly, the dominant predictor is THS which is in agreement with the original study, as are the influential roles of both 5PD and 5PT. We believe that the ability of the hierarchical GAM to produce plots such as these constitutes a considerable advantage over the machine learning methods tested and highlights the ability of LN-CASS to generate not only strong predictive models, but to be used as an exploratory tool for the generation of hypotheses for future study.

D. Microarray data

The final case study we conducted focused on the well-known Colon dataset of Alon et al.¹². The dataset consists of measurements of the expression levels of 2000 genes in 62 subjects, with the response variable being an indicator of Colon cancer incidence, representing a typical $p \gg n$ problem in the biological/medical sciences. We compared the performance of logistic regression, with LN-CASS priors on the coefficients, to LASSO, Random Forest and Neural Network classifiers. We performed leave-one-out cross-validation (LOOCV) and computed the AUC across the left out samples in order to compare the estimated out-of-sample predictive accuracy of each method. In order to reduce the bias of

the AUC estimates, we randomly removed an observation of the opposite class in each fold so that the class proportions were identical across folds. Code to reproduce the results of this section is available in Data Accessibility and contains details of the particular implementations of all algorithms used.

We preprocessed the data by first log-transforming and subsequently standardising (i.e. subtracting the mean and dividing by the standard deviation) the expression level of each gene. We then screened the genes via a preliminary Wald test and selected the 500 genes with the largest Z-scores in absolute value, leaving us with a predictor matrix consisting of the expression levels of 500 genes in 62 tissues which acted as the input to all subsequent models.

The pooled LOOCV AUCs for each method were as follows: LN-CASS, 0.904; Neural Network, 0.8898; Random Forest, 0.8892; LASSO, 0.858. LN-CASS performs the best, but again all of the methods perform well and there is not a substantial difference between the estimated out-of-sample performance of each method.

Interestingly, there is some biological evidence for class-mislabelling, i.e. samples being incorrectly marked as either tumour or healthy^{12,18} in the colon dataset. According to Bootkrajang & Kabán¹⁸, there are nine such samples. Figure 4 shows the mean posterior prediction for each subject with these ‘suspicious’ subjects circled. Clearly, there is reasonable agreement based on a visual inspection of the plot between the potentially mislabelled samples and those suggested by visual inspection of the LN-CASS model predictions. This suggests a possible secondary function of the LN-CASS prior in identifying mislabelled samples, the details however are left to future work.

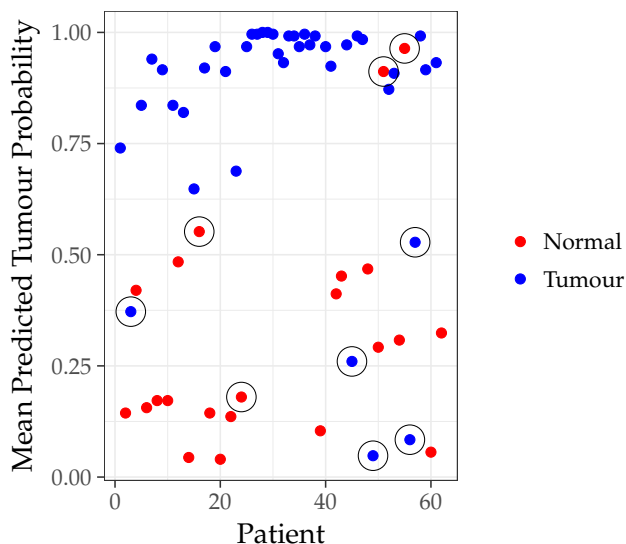


FIG. 4: Mean predictions and observed outcomes from the LN-CASS model for the microarray data. Circled points have been identified as potentially mislabelled by^{12,18}

III. DISCUSSION

We have presented a new prior distribution for performing regularisation/shrinkage in a Bayesian framework. We have shown that its ability to produce generalisable predictive models is comparable to state-of-the-art machine learning methods on two datasets of biological interest. Additionally, we have demonstrated with a simulation study the ability of our method to recover ground-truth parameters, even when the number of parameters is larger than the number of datapoints. In this regard, the performance of the LN-CASS prior is considerably better than other regularisation/shrinkage methods which aim to estimate the parameters of classical, generative probability models (linear regression, logistic regression etc.).

We believe that, combined, these two properties of the LN-CASS prior make it a worthwhile addition to the toolboxes of applied scientists working with typical biological datasets.

Our prior requires the choices of three hyperparameters, although we contend that they are much more interpretable than those required for other Bayesian shrinkage methods (see the ESM sections 1 & 2 for further detail on the roles of the hyperparameters). The three hyperparameters required correspond to, firstly, the standard deviation of the ‘slab’ component, which we refer to as τ ; for standardised predictors, a default value of $\tau = 5$ has been sufficient for all of our applications because it essentially provides a vague Gaussian prior for non-zero coefficients. Secondly, the parameters of the Logit-Normal distribution (fig. 1 (a)) must be specified; we refer to these parameters as μ_λ and

σ_λ . μ_λ can be chosen based on our prior beliefs about the probability of a zero coefficient, and in our experience does not require much tuning; the median of the Logit-Normal distribution is given by $\text{sigm}(\mu_\lambda)$, where $\text{sigm}(\cdot)$ is the logistic sigmoid function. Thus, if we believe *a priori* that each coefficient has a probability a of being non-zero, we simply set $\mu_\lambda = \text{logit}(a)$. σ_λ simply controls the quality of the approximation to the spike-and-slab prior, with larger values corresponding to better approximations. We have used a default value of $\sigma_\lambda = 10$ throughout the paper; results are not sensitive to increases in this value.

The final key advantage of the LN-CASS prior is the ease with which it generalises to problems with a hierarchical complexity structure. This allows finer control of what exactly we mean by a ‘complex’ model, and what we mean by a desirable model – our example of using a generalised additive model for studying the metabolomics data above illustrates this point. In that case, we imposed a hierarchical complexity structure: no effect \rightarrow linear effect \rightarrow nonlinear effect. In the simulation study, we favoured a complexity structure: no effect \rightarrow shared group effect \rightarrow individual effect. These hierarchies are accomplished simply by propagating the value of the Logit-Normal random variable through each layer and taking its product with a new Logit-Normal random variable.

We note that the prior is particularly amenable to problems in which a hierarchical complexity structure is desired, by which we mean problems in which simpler models are nested within more complex models. The simplest case is the domain of the majority of the shrinkage/regularisation literature: models with fewer parameters are nested within models with more parameters. However, there are other problems with similar properties; linear models are nested within nonlinear models, models with some predictors sharing coefficients are nested within models in which each predictor has its own coefficient. One possible area of application is in multi-state survival modelling. Multi-state models describe transitions between disease states by distinct hazard functions, which may be difficult to fit with a small sample size. One might expect that the effects of many covariates remain fairly similar regardless of the state, for example age. Thus the LN-CASS prior could be used to introduce a ‘soft’ constraint, encouraging but not enforcing covariates to share a parameter across hazard functions. This would essentially involve placing a grouped LN-CASS prior on the regression coefficients (as in the simulation study), with the groups corresponding to covariate effects.

As with most Bayesian methods, the main obstacle to the implementation of this methodology is the computational burden of MCMC sampling. Recent developments have made this procedure much more straightforward to implement and much faster^{19,20}. However, for large problems this computational burden is likely to be too large to compete with the much faster frequentist

and machine learning methods available. Approximate Bayesian methods offer more computationally tractable alternatives to MCMC sampling, and would be an interesting avenue of future research for this problem and allow its scalability to very large problems. In particular, nonparametric variational inference²¹ appears to be the most reasonable direction, since it is able to deal both with multimodal posterior distributions and non-conjugate prior distributions.

One concern with spike-and-slab type inference procedures is the presence of multi-modal posterior distributions and the subsequent difficulty of some samplers to sample effectively from the posterior distribution, due to them becoming ‘stuck’ in local modes. We checked that the sampler we employed was able to effectively explore multi-modal posterior distributions in a linear model with interactions, a problem that is particularly prone to multi-modal posteriors. We found that whenever multi-modal posteriors appeared, they were effectively explored by multiple MCMC chains, and the Gelman-Rubin statistic²² revealed that samples were consistent across chains. Section 4 of the ESM and the accompanying code provide more details.

The LN-CASS method does not inherently provide ‘hard’ variable selection, i.e. completely removing variables from the model, in the ilk of the LASSO. We advocate using the full model (i.e. including all predictors) for making predictions wherever possible, and using the absolute values of estimated parameters as variable importance measures for identifying the most important predictors for the purposes of hypothesis generation and/or obtaining biological insight. However, particularly in clinical/diagnostic circumstances, hard variable selection is useful to reduce the burden on clinicians/diagnosticians in collecting relevant data for utilising the model at the point of care.

A variety of applicable procedures for hard variable selection in Bayesian shrinkage models are available in an excellent review by Vehtari et al.²³. One particularly simple method is to specify a threshold on the absolute values of the median parameters, i.e. discard all predictors whose absolute value is below some threshold. This threshold could be chosen based on the predictive performance of submodels containing only the predictors corresponding to the largest k coefficients in absolute value – one simply evaluates the predictive performance of each submodel and specifies a percentage of the maximum (i.e. the model including all variables) to retain.

To summarise, we have presented a flexible tool for performing regularised Bayesian regression in a variety of settings, which allows one to construct (with relative ease) problem-specific penalties on model complexity. The performance on out-of-sample data is typically at least as good as state-of-the-art methods, but the prior allows the use of classical statistical models which can be interpreted simply by applied biomedical scientists.

Acknowledgements

WT would like to thank the EPSRC for funding his doctoral research. SJ would like to acknowledge the BB-SRC for funding (BB/M021386/1). SJ, DJS & WT would like to thank the Wellcome Trust for funding the University of Birmingham’s Model Parameterisation for Predictive Medicine workshop (1516ISSFFEL9).

Data Accessibility

Code and data used (with the exception of data for the Metabolomics case study) are available on GitHub via <https://github.com/willthomson1/RS-Interface-code>.

Author Contributions

WT: designed the research, methodological development, coding, data analysis and write up; DJS and SJ: designed the research, contributed to the write up and data interpretation, supervised the project; WA and AT: biomedical motivation and interpretation, data contribution

References

- ¹R. Tibshirani, *Journal of the Royal Statistical Society. Series B (Methodological)* **58**, 267 (1996).
- ²J. Lever, M. Krzywinski, and N. Altman, *Nature Methods* **13**, 803 (2016).
- ³C. M. Carvalho, N. G. Polson, and J. G. Scott, *Biometrika* **97**, 465 (2010).
- ⁴T. Park and G. Casella, *Journal of the American Statistical Association* **103**, 681 (2008).
- ⁵J. E. Griffin and P. J. Brown, *Bayesian Analysis* **5**, 171 (2010).
- ⁶L. Jacob, G. Obozinski, and J.-P. Vert, in *Proceedings of the 26th annual international conference on machine learning (ACM, 2009)* pp. 433–440.
- ⁷N. Simon, J. Friedman, T. Hastie, and R. Tibshirani, *Journal of Computational and Graphical Statistics* **22**, 231 (2013).
- ⁸R.-B. Chen, C.-H. Chu, S. Yuan, and Y. N. Wu, *Journal of Computational and Graphical Statistics* **25**, 665 (2016).
- ⁹X. Xu, M. Ghosh, *et al.*, *Bayesian Analysis* **10**, 909 (2015).
- ¹⁰E. I. George and R. E. McCulloch, *Journal of the American Statistical Association* **88**, 881 (1993).
- ¹¹N. G. Polson and J. G. Scott, *Bayesian Statistics* **9**, 501 (2010).
- ¹²U. Alon, N. Barkai, D. A. Notterman, K. Gish, S. Ybarra, D. Mack, and A. J. Levine, *Proceedings of the National Academy of Sciences* **96**, 6745 (1999).
- ¹³W. Arlt, M. Biehl, A. E. Taylor, S. Hahner, R. Libe, B. A. Hughes, P. Schneider, D. J. Smith, H. Stiekema, N. Krone, *et al.*, *The Journal of Clinical Endocrinology & Metabolism* **96**, 3775 (2011).
- ¹⁴J. Griffin, P. Brown, *et al.*, *Bayesian Analysis* **12**, 135 (2017).
- ¹⁵S. Buuren and K. Groothuis-Oudshoorn, *Journal of statistical software* **45** (2011).
- ¹⁶Y. Benjamini and Y. Hochberg, *Journal of the royal statistical society. Series B (Methodological)* , 289 (1995).
- ¹⁷Y. Benjamini and D. Yekutieli, *Annals of statistics* , 1165 (2001).
- ¹⁸J. Bootkrajang and A. Kabán, *Bioinformatics* **29**, 870 (2013).
- ¹⁹Stan Development Team, “RStan: the R interface to Stan,” (2016), r package version 2.14.1.
- ²⁰M. D. Hoffman and A. Gelman, *Journal of Machine Learning Research* **15**, 1593 (2014).
- ²¹S. J. Gershman, M. D. Hoffman, and D. M. Blei, in *Proceedings of the 29th International Conference on International Conference on Machine Learning, ICML/12* (2012).
- ²²A. Gelman, D. B. Rubin, *et al.*, *Statistical science* **7**, 457 (1992).
- ²³A. Vehtari, J. Ojanen, *et al.*, *Statistics Surveys* **6**, 142 (2012).