

Statistical Applications in Genetics and Molecular Biology

Volume 11, Issue 2

2012

Article 5

COMPUTATIONAL STATISTICAL METHODS FOR GENOMICS AND
SYSTEMS BIOLOGY

Bayesian Sparsity-Path-Analysis of Genetic Association Signal using Generalized t Priors

Anthony Lee, *University of Oxford*

Francois Caron, *INRIA Bordeaux*

Arnaud Doucet, *University of Oxford*

Chris Holmes, *University of Oxford*

Recommended Citation:

Lee, Anthony; Caron, Francois; Doucet, Arnaud; and Holmes, Chris (2012) "Bayesian Sparsity-Path-Analysis of Genetic Association Signal using Generalized t Priors," *Statistical Applications in Genetics and Molecular Biology*: Vol. 11: Iss. 2, Article 5.

DOI: 10.2202/1544-6115.1712

Available at: <http://www.bepress.com/sagmb/vol11/iss2/art5>

©2012 De Gruyter. All rights reserved.

Bayesian Sparsity-Path-Analysis of Genetic Association Signal using Generalized t Priors

Anthony Lee, Francois Caron, Arnaud Doucet, and Chris Holmes

Abstract

We explore the use of generalized t priors on regression coefficients to help understand the nature of association signal within “hit regions” of genome-wide association studies. The particular generalized t distribution we adopt is a Student distribution on the absolute value of its argument. For low degrees of freedom, we show that the generalized t exhibits “sparsity-prior” properties with some attractive features over other common forms of sparse priors and includes the well known double-exponential distribution as the degrees of freedom tends to infinity. We pay particular attention to graphical representations of posterior statistics obtained from sparsity-path-analysis (SPA) where we sweep over the setting of the scale (shrinkage/precision) parameter in the prior to explore the space of posterior models obtained over a range of complexities, from very sparse models with all coefficient distributions heavily concentrated around zero, to models with diffuse priors and coefficients distributed around their maximum likelihood estimates. The SPA plots are akin to LASSO plots of maximum a posteriori (MAP) estimates but they characterise the complete marginal posterior distributions of the coefficients plotted as a function of the precision of the prior. Generating posterior distributions over a range of prior precisions is computationally challenging but naturally amenable to sequential Monte Carlo (SMC) algorithms indexed on the scale parameter. We show how SMC simulation on graphic-processing-units (GPUs) provides very efficient inference for SPA. We also present a scale-mixture representation of the generalized t prior that leads to an expectation-maximization (EM) algorithm to obtain MAP estimates should only these be required.

KEYWORDS: sparsity prior, Bayesian analysis, genome-wide association studies

Author Notes: Anthony Lee, Oxford-Man Institute and Department of Statistics, University of Oxford, UK. François Caron, INRIA Bordeaux Sud-Ouest and Institut de Mathématiques de Bordeaux, University of Bordeaux, France. Arnaud Doucet, Institute of Statistical Mathematics, Japan and Department of Statistics, University of Oxford, UK. Chris Holmes, Department of Statistics and Oxford-Man Institute and Wellcome Trust Centre for Human Genetics, University of Oxford, and MRC Harwell, UK. Anthony Lee would like to acknowledge the support of the University of Oxford Clarendon Fund Scholarship and the Oxford-Man Institute. Chris Holmes would like to acknowledge support for this research by the Medical Research Council, Oxford-Man Institute and Wellcome Trust Centre for Human Genetics, Oxford. The authors would additionally like to thank the organizers of the recent meeting on "Computational statistical methods for genomics and system biology" held at the Centre de Recherches Mathématiques, Montreal.

1 Introduction

Genome-wide association studies (GWAS) have presented a number of interesting challenges to statisticians (Balding, 2006). Conventionally GWAS use single marker univariate tests of association, testing marker by marker in order to highlight “hit regions” of the genome showing evidence of association signal. The motivation for our work concerns methods in Bayesian sparse multiple regression analysis to characterise and decompose the association signal within such regions spanning possibly hundreds of markers. Hit-regions typically cover loci in high linkage disequilibrium (LD) which leads to strong correlation between the markers and makes multiple regression analysis non-trivial. *A priori* we would expect only a small number of markers to be responsible for the association and hence priors that induce sparsity are an important tool to aid understanding of the genotype-phenotype dependence structure; an interesting question being whether the association signal is consistent with a single causal marker or due to multiple effects. We restrict attention to case-control GWAS, i.e. binary phenotypes, these being by far the most prevalent although generalisations to other exponential family likelihoods or non-linear models is relatively straightforward.

In the non-Bayesian literature sparse regression analysis via penalised likelihood has gained increasing popularity since the seminal paper on the LASSO (Tibshirani, 1996, 2011). The LASSO estimates have a Bayesian interpretation as maximum *a posteriori* (MAP) statistics under a double exponential prior on regression coefficients $p(\beta) \propto \exp(-\lambda \sum_j |\beta_j|)$ and it is known that, unlike ridge penalties using a normal prior, the Lasso prior tends to produce ‘sparse’ solutions in that as an increasing function of the regularisation penalty λ more and more of the MAP estimates will be zero. From a Bayesian perspective the use of MAP estimates holds little justification, see the discussion of Tibshirani (2011), and in Park and Casella (2008) they explore full Bayesian posterior analysis with Lasso double-exponential priors using Markov chain Monte Carlo (MCMC) to make inference. While the Lasso penalty / prior is popular there is increasing awareness that the use of identical ‘penalization’ on each coefficient, e.g. $\lambda \sum_{j=1}^p |\beta_j|$, can lead to unacceptable bias in the resulting estimates (Fan and Li, 2001). In particular, coefficients get shrunk towards zero even when there is overwhelming evidence in the likelihood that they are non-zero. This has motivated use of sparsity-inducing non-convex penalties which reduce bias in the estimates of large coefficients at the cost of increased difficulty in computing. Of note we would mention the “adaptive” methods (Zou, 2006, Zou and Li, 2008) in the statistics literature and iteratively reweighted methods (Candès, Wakin, and Boyd, 2008, Chartrand and Yin, 2008) in the signal processing literature; although these papers are only concerned with MAP estima-

tion, there being no published articles to date detailing the use of full Bayesian analysis.

In the context of GWAS there have been recent illustrations of the use of adaptive sparsity priors and MAP estimation (Hoggart, Whittaker, De Iorio, and Balding, 2008), reviewed and compared in Ayers and Cordell (2010). In Hoggart et al. (2008) they use the Normal-Exponential-Gamma sparsity-prior of Griffin and Brown (2007) to obtain sparse MAP estimates from logistic regression. The use of continuous sparse priors can be contrasted with an alternative approach using two component mixture priors such as in George and McCulloch (1993) which adopt a distribution containing a spike at zero and another component with broad scale. Coefficients are then *a posteriori* classified into either state, relating to whether their corresponding variables (genotypes) are deemed predictively irrelevant or relevant. In the GWAS setting the two-component mixture prior has been explored by Wilson, Iversen, Clyde, Schmidler, and Schildkraut (2010) and Fridley (2009). For model exploration the sparsity prior has some benefits in allowing the statisticians to visualize the regression analysis over a range of scales / model complexities.

In this paper we advocate the use of the generalized t prior, first considered in McDonald and Newey (1988), on the regression coefficients and a full Bayesian analysis. The generalized t prior we adopt is a Student distribution on the absolute value of the coefficient. We demonstrate it has a number of attractive features as a 'sparsity prior' in particular due to its simple analytic form and interpretable parameters. In applications the setting of the scale parameter in the prior is the key task that affects the sparsity of the posterior solution. We believe much is to be gained in exploring the continuum of posterior models formed by sweeping through the scale of the prior, even over regions of low posterior probability, in an exploratory approach we term sparsity-path-analysis (SPA). This leads towards providing graphical representations of the posterior densities that arise as we move from the most sparse models with all coefficient densities heavily concentrated around the origin through to models with diffuse priors and coefficients distributed around their maximum likelihood estimates. We stress that even though parts of this model space have low probability, they have utility and are interesting in gaining a fuller understanding of the predictor-response (genotype-phenotype) dependence structure. SPA is an exploratory tool that is computationally challenging but highly suited to sequential Monte Carlo (SMC) algorithms indexed on the scale parameter and simulated in parallel on graphics cards using graphical processing units (GPUs).

The generalized t as a prior for Bayesian regression coefficients was first suggested by McDonald and Newey (1988) and more recently in the context of sparse MAP estimation for normal (Gaussian) likelihoods (Lee, Caron, Doucet, and Holmes, 2010a, Cevher, 2009, Armagan, Dunson, and Lee, 2011), and using MCMC for a full Bayesian linear regression analysis in Armagan et al. (2011).

In Cevher (2009), Armagan et al. (2011) the prior is referred to as a generalized double-pareto distribution but we prefer to use the terminology of the generalized t as we feel it is more explicit and easier to interpret as such. Our contribution in this paper is to consider a full Bayesian analysis of logistic regression models with generalized t priors in the context of GWAS applications for which we develop SMC samplers and GPU implementation. The GPU implementation is important in practice; in the GWAS analysis it reduces run-time from over five days to hours. We pay particular attention to developing graphical displays of summary statistics from the joint posterior distribution over a wide range of prior precisions; arguing that parts of the model space with low probability have high utility and are informative to the exploratory understanding of regression signal.

In Section 2 we introduce the generalized t prior and show that it has a scale-mixture representation and exhibits sparsity type properties for regression analysis. Section 3 provides a brief description of the genetic association data and simulated phenotypes used in our methods development. Section 4 concerns computing posterior distributions over a range of prior scale parameters using sequential Monte Carlo algorithms simulated on GPUs. Section 5 deals with graphical display of SPA statistics and plots of posterior summary statistics over the path of scales. We conclude with a discussion in Section 6.

2 Generalized t prior for sparse regression analysis

In Bayesian logistic regression analysis we are given a data set $\{\mathbf{X}, \mathbf{y}\}$, of an $(n \times p)$ predictor matrix \mathbf{X} and an $(n \times 1)$ vector of binary phenotypes (or response variables) \mathbf{y} . We adopt a prior $p(\boldsymbol{\beta})$ on the p regression coefficients from which the log-posterior updates as,

$$\log p(\boldsymbol{\beta}|\mathbf{X}, \mathbf{y}, a, c) = \log f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}) + \log p(\boldsymbol{\beta}) + C$$

where $f(\cdot)$ is the likelihood function, C is a constant independent of $\boldsymbol{\beta}$, and

$$\begin{aligned} f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}) &= \prod_{i=1}^n (p_i)^{y_i} (1 - p_i)^{1-y_i} \\ p_i &= \text{logit}^{-1}(x_i \boldsymbol{\beta}). \end{aligned}$$

In this paper we propose the generalized t distribution as an interesting prior for each component of $\boldsymbol{\beta}$, which has the form of a Student density on the L_q -norm of its argument, (McDonald and Newey, 1988),

$$p(\beta_j|\mu, a, c, q) = \frac{q}{2ca^{1/q}B(1/q, a)} \left(1 + \frac{|\beta_j - \mu|^q}{ac^q}\right)^{-(a+1/q)} \quad (1)$$

where we note that the usual Student's density with ν degrees of freedom is obtained with $\{a = \nu/2, q = 2, c = \sqrt{2}\}$; the exponential power family of distributions arise as the degrees of freedom parameter $a \rightarrow \infty$; and the double-exponential or Laplace distribution occurs with $\{q = 1, a = \infty\}$. The use of this distribution as a prior on regression coefficients was first described in McDonald and Newey (1988).

We shall only consider the L_1 -norm case with central location $\mu = 0$, and independent priors across the p coefficients for which we can simplify the distribution to,

$$p(\beta_j | \mu = 0, a, c, q = 1) = \frac{1}{2c} \left(1 + \frac{|\beta_j|}{ac} \right)^{-(a+1)} \quad (2)$$

where $c > 0$ is a scale parameter and $a > 0$ is related to the degrees of freedom; we will denote this centred L_1 generalized t as $Gt(a, c)$. Note that Cevher (2009) and Armagan et al. (2011) refer to this distribution as a generalized double-pareto but we prefer to think of it as an instantiation of a generalized t following the earlier derivation in the context of priors on regression coefficients in McDonald and Newey (1988). This also allows for easier understanding of the parameters and direct connections with the double-exponential distribution and other standard densities.

It is interesting to note that (2) can be derived from the marginal of a scale-mixture of double-exponential densities. For multivariate $\boldsymbol{\beta} = (\beta_1, \dots, \beta_p)$ where

$$\begin{aligned} p(\boldsymbol{\beta}) &= \prod_{j=1}^p p(\beta_j | \tau_j) \\ p(\beta_j | \tau_j) &= \frac{1}{2\tau_j} \exp\left(-\frac{|\beta_j|}{\tau_j}\right) \\ \tau_j &\sim IG(a, b) \end{aligned} \quad (3)$$

and $IG(\cdot, \cdot)$ denotes the Inverse-Gamma density. This gives

$$p(\beta_j) = \int_0^\infty p(\beta_j | \tau_j) p(\tau_j | a, b) d\tau_j = Gt(a, b/a)$$

where we can set $b = ac$ to generate a marginal $Gt(a, c)$ density. Hence as a prior on the set of regression coefficients, one way to think of (2) is as defining a hierarchical prior whereby each coefficient has a double-exponential prior with a coefficient specific scale parameter distributed *a priori* as $IG(a, ac)$. A plot of the log densities for $Gt(1, 1)$, and the double-exponential, $DE(0, 1)$ is shown in Fig. 1 where we can see the greater kurtosis and heavier tails provided by the generalized t. As mentioned in the Section 1, the relatively light tails of the $DE(\cdot)$ prior is unattractive




Figure 1: Plot of log density under the $Gt(1, 1)$ (solid line) and the Lasso prior $DE(0, 1)$ (dashed line).

as it tends to shrink large values of the coefficients even when there is clear evidence in the likelihood that they are predictively important.

We are interested with inferring and graphing the regression coefficients' posterior distributions over a range of values for c for fixed a , an exploratory procedure we call sparsity-path-analysis. But first it is instructive to examine the form of the MAP estimates obtained under the generalized t prior as this sheds light on their utility as sparsity priors.

2.1 MAP estimates via the EM algorithm

The optimization problem associated with obtaining the MAP estimates for $\boldsymbol{\beta}$ under the generalized t-distribution prior is not concave. However, one can find local modes of the posterior by making use of the scale-mixture representation (3) and using the EM algorithm with the hierarchical scale parameters $\boldsymbol{\tau} = \tau_{1:p}$ as latent variables. Indeed, each iteration of EM takes the form

$$\boldsymbol{\beta}^{(t+1)} = \arg \max_{\boldsymbol{\beta}} \log f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}) + \int \log[p(\boldsymbol{\beta}|\boldsymbol{\tau})]p(\boldsymbol{\tau}|\boldsymbol{\beta}^{(t)}, a, b)d\boldsymbol{\tau}$$

where $b = ac$. The conjugacy of the inverse-gamma distribution with respect to the Laplace distribution gives

$$\tau_j|\beta_j^{(t)}, a, b \sim IG(a + 1, b + |\beta_j^{(t)}|)$$

allowing for different prior parameters for each coefficient, and with $p(\boldsymbol{\beta}|\boldsymbol{\tau}) = \prod_{j=1}^p p(\beta_j|\tau_j) = \prod_{j=1}^p 1/(2\tau_j) \exp(-|\beta_j|/\tau_j)$ yields

$$\boldsymbol{\beta}^{(t+1)} = \arg \max_{\boldsymbol{\beta}} \log f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}) - \sum_{j=1}^p |\beta_j| \int \frac{1}{\tau_j} p(\tau_j|\beta_j^{(t)}, a, b) d\tau_j$$

where the expectation of $1/\tau_j$ given $\tau_j \sim IG(a+1, b+|\beta_j^{(t)}|)$ is $(a+1)/(b+|\beta_j^{(t)}|)$. As such, one can find a local mode of the posterior $p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{X}, a, b)$ by starting at some point $\boldsymbol{\beta}^{(0)}$ and then iteratively solving

$$\boldsymbol{\beta}^{(t+1)} = \arg \max_{\boldsymbol{\beta}} \log f(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta}) - \sum_{j=1}^p w_j^{(t)} |\beta_j| \tag{4}$$

where

$$\begin{aligned} w_j^{(t)} &= \frac{a+1}{b+|\beta_j^{(t)}|} \\ &= \frac{a+1}{ac+|\beta_j^{(t)}|} \end{aligned}$$

For logistic regression the update (4) is a convex optimization problem for which standard techniques exist. We can see the connection to the Lasso as with $a \rightarrow \infty$ we find $w_j^{(t)} \rightarrow \frac{1}{c}$, where $\lambda = 1/c$ is the usual Lasso penalty.

2.2 Oracle properties of the MAP

In the penalized optimization literature, some estimators are justified at least partially by their possession of the oracle property: that for appropriate parameter choices, the method performs just as well as an oracle procedure in terms of selecting the correct predictors and estimating the nonzero coefficients correctly asymptotically in n when the likelihood is Gaussian. Other related properties include asymptotic unbiasedness, sparsity and continuity in the data (Fan and Li, 2001). Penalization schemes with the oracle property include the smoothly clipped absolute deviation method (Fan and Li, 2001), the adaptive lasso (Zou, 2006) and the one-step local linear approximation method (Zou and Li, 2008).

For the generalized t-distribution, the MAP estimate obtained using this prior has been examined in (Lee et al., 2010a, Cevher, 2009, Armagan et al., 2011), all of which derive the expectation-maximisation algorithm for finding a local mode

of the posterior. It is straightforward to establish that the estimate is asymptotically unbiased, sparse when $c < 2\sqrt{(a+1)}/a$ and continuous in the data when $c = \sqrt{(a+1)}/a$ following the conditions laid out in Fan and Li (2001). Armagan et al. (2011) shows that this scheme satisfies the oracle property if simultaneously $a \rightarrow \infty$, $a/\sqrt{n} \rightarrow 0$ and $ac\sqrt{n} \rightarrow C < \infty$ for some constant C .

It is worth remarking that the oracle property requires the prior on β to depend on the number of observations n . This is in conflict with conventional Bayesian analysis, where priors represent beliefs about the data generating mechanism irrespective of the number of observations that will be obtained. In addition, for a given application it may not be the case that n and p are in the realm of asymptotic relevance for this property to be meaningful.

2.3 Comparison to other priors

There have been a number of sparsity prior representations advocated recently in the literature, although it is important to note that most of the papers deal solely with MAP estimation and not full Bayesian analysis. To name a few important contributions we highlight the normal-gamma, normal-exponential-gamma (NEG) and the Horseshoe priors, see (Griffin and Brown, 2010, Caron and Doucet, 2008, Griffin and Brown, 2007, Carvalho, Polson, and Scott, 2010). Like all these priors the $Gt(\cdot, \cdot)$ shares the properties of sparsity and adaptive shrinkage, allowing for less shrinkage on large coefficients as shown in Fig. 2 for a normal observation, and a scale-mixture representation which facilitates an efficient EM algorithm to obtain MAP estimates.

In Table 1 we have listed the log densities for the sparsity-priors mentioned above. From Table 1 we see an immediate benefit of the generalized t density is that it has a simple analytic form compared with other sparsity-priors and is well known and understood by statisticians, which in turn aids explanation to data owners. The generalized t is also easy to compute relative to the NEG, normal-gamma and Horseshoe priors which all involve non-standard computation, making them more challenging to implement and less amenable to fast parallel GPU simulation. Finally, the setting of the parameters of the generalized t is helped as statisticians are used to thinking about degrees of freedom and scale parameters of Student's densities. We believe this will have an impact on the attractiveness of the prior to non-expert users outside of sparsity researchers.




Figure 2: Estimate of the posterior mean under a single normal observation, y , for the $Gt(a = 1, c = 0.1)$ (blue dashed) and $Gt(a = 4, c = 0.1)$ (red dashed). We can see the sparsity by the setting to zero around $y = 0$ and the approximate unbiasedness with reduction in shrinkage as $|y| \gg 0$.

	$\log p(\beta_j)$
double-exponential	$-\lambda \beta_j + C$
normal-gamma	$\left(\frac{1}{2} - \lambda\right) \log \beta_j - \log K_{\lambda-1/2} \left(\frac{ \beta_j }{\lambda}\right) + C$
NEG	$-\frac{\beta^2}{4\lambda^2} - \log D_{-2(\lambda+\frac{1}{2})} \left(\frac{ \beta }{\lambda}\right) + C$
Horseshoe	$\log \left[\int_0^\infty (2\tau^2 \lambda^2)^{-1/2} \exp \left(-\frac{\beta_j^2}{2\tau^2 \lambda^2} \right) (1 + \lambda)^{-2} d\lambda \right] + C$
$Gt(a, c)$	$-(a + 1) \log \left(1 + \frac{ \beta_j }{ac} \right) + C$

Table 1: Log densities for the double-exponential and various sparsity-priors, up to an additive prior-specific constant C . $K(\cdot)$ denotes a Bessel function of the third kind and $D_\nu(z)$ is the parabolic cylinder function. The Horseshoe prior does not have an analytic form.

3 Genotype Data

In the development and testing of the methods in this paper we made use of anonymous genotype data obtained from the Wellcome Trust Centre for Human Genetics (WTCHG), Oxford. The data consists of a subset of single-nucleotide polymorphisms (SNPs) from a genome-wide data set originally gathered in a case-control study to identify colorectal cancer risk alleles. The data set we consider consists of genotype data for 1859 subjects on 184 closely spaced SNPs from a hit-region on chromosome 18q. We standardise the columns of \mathbf{X} to be mean zero and unit standard deviation. A plot of the correlation structure across the first 100 markers in the region is shown in Fig. 3. We can see the markers are in high LD with block like strong correlation structure that makes multiple regression analysis challenging.




Figure 3: Plot of correlation structure across the first 100 adjacent genotype markers from chromosome 18q.

We generated pseudo-phenotype data, \mathbf{y} , by selecting at random five loci and generating coefficients, $\beta \sim N(0, 0.2^2)$, judged to be realistic signal sizes seen in GWAS (Stephens and Balding, 2009). All other coefficients were set to zero. We then generated phenotype, case-control, binary \mathbf{y} data by sampling Bernoulli trials with probability given by the logit of the log-odds, $\mathbf{x}_i\boldsymbol{\beta}$, for an individual's genotype \mathbf{x}_i . We considered two scenarios. The first when we construct a genotype matrix of 500 individuals using the five relevant markers and an additional 45 “null” markers whose corresponding coefficients were zero, i.e. a subset of the data from

the study. The second data set uses all available markers and all subjects to create a (1859×184) design matrix. To clarify, all covariate data used in the following examples is real genotype data from chromosome 18q on individuals genotyped as part of the WTCHG study. However the phenotype (response variable) \mathbf{y} is simulated due to confidentiality issues.

4 Sequential Monte Carlo and graphic card simulation

We propose to explore the posterior distributions of $\boldsymbol{\beta}|\mathbf{X}, \mathbf{y}, a, c$ with varying $c = b/a$ via Sequential Monte Carlo (SMC) sampler methodology (Del Moral, Doucet, and Jasra, 2006). In particular, we let the target distribution at time t be $\pi_t(\boldsymbol{\beta}) = p(\boldsymbol{\beta}|\mathbf{X}, \mathbf{y}, a, b_t/a)$ with $t = \{1, \dots, T\}$ and some specified sequence $\{b_t\}_{t=1}^T$ where $b_i > b_{i+1}$ and b_T is small enough for the prior to dominate the likelihood with all coefficients posterior distributions heavily concentrated around zero, and b_1 large enough for the prior to be diffuse and have limited impact on the likelihood. In order to compute posterior distributions over the range of $\{b_t\}_{t=1}^T$ we make use of sequential Monte Carlo samplers indexed on this sequence of scales.

It is important to note that we cannot evaluate

$$\pi_t(\boldsymbol{\beta}) = \frac{p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta})p(\boldsymbol{\beta}|a, b_t)}{\int p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta})p(\boldsymbol{\beta}|a, b_t)d\boldsymbol{\beta}}$$

since we cannot evaluate the integral in the denominator, the marginal likelihood of \mathbf{y} . However, we can evaluate $\gamma_t(\boldsymbol{\beta}) = Z_t \pi_t(\boldsymbol{\beta})$ where $Z_t = \int p(\mathbf{y}|\mathbf{X}, \boldsymbol{\beta})p(\boldsymbol{\beta}|a, b_t)d\boldsymbol{\beta}$.

At time t we have a collection of N particles $\{\boldsymbol{\beta}_t^{(i)}\}_{i=1}^N$ and accompanying weights $\{W_t^{(i)}\}_{i=1}^N$ such that the empirical distribution

$$\tilde{\pi}_t^N(d\boldsymbol{\beta}) = \sum_{i=1}^N W_t^{(i)} \delta_{\boldsymbol{\beta}_t^{(i)}}(d\boldsymbol{\beta})$$

is an approximation of π_t . The weights are given by

$$W_t^{(i)} = \frac{W_{t-1}^{(i)} w_t(\boldsymbol{\beta}_{t-1}^{(i)}, \boldsymbol{\beta}_t^{(i)})}{\sum_{j=1}^N W_{t-1}^{(j)} w_t(\boldsymbol{\beta}_{t-1}^{(j)}, \boldsymbol{\beta}_t^{(j)})} \quad (5)$$

where

$$w_t(\boldsymbol{\beta}_{t-1}^{(i)}, \boldsymbol{\beta}_t^{(i)}) = \frac{\gamma_t(\boldsymbol{\beta}_t^{(i)})L_{t-1}(\boldsymbol{\beta}_t^{(i)}, \boldsymbol{\beta}_{t-1}^{(i)})}{\gamma_{t-1}(\boldsymbol{\beta}_{t-1}^{(i)})K_t(\boldsymbol{\beta}_{t-1}^{(i)}, \boldsymbol{\beta}_t^{(i)})} \quad (6)$$

We use an Markov chain Monte Carlo (MCMC) kernel K_t with invariant distribution π_t to sample each particle $\boldsymbol{\beta}_t^{(i)}$, i.e. $\boldsymbol{\beta}_t^{(i)} \sim K_t(\boldsymbol{\beta}_{t-1}^{(i)}, \cdot)$. The backwards kernel L_{t-1} used to weight the particles is essentially arbitrary but we use the reversal of K_t

$$L_{t-1}(\boldsymbol{\beta}_t, \boldsymbol{\beta}_{t-1}) = \frac{\pi_t(\boldsymbol{\beta}_{t-1})K_t(\boldsymbol{\beta}_{t-1}, \boldsymbol{\beta}_t)}{\pi_t(\boldsymbol{\beta}_t)}$$

so that the weights in (6) are given by

$$w_t(\boldsymbol{\beta}_{t-1}^{(i)}, \boldsymbol{\beta}_t^{(i)}) = \frac{\gamma_t(\boldsymbol{\beta}_{t-1}^{(i)})}{\gamma_{t-1}(\boldsymbol{\beta}_{t-1}^{(i)})}$$

as suggested in (Crooks, 1998, Neal, 2001, Del Moral et al., 2006).

Since the weights do not depend on the value of $\boldsymbol{\beta}_t$, it is better to resample the particles before moving them according to K_t , when resampling is necessary. We resample the particles when the effective sample size (ESS) (Liu and Chen, 1995) is below $(3/4)N$. Also note that even if we knew the normalizing constants Z_{t-1} and Z_t , they would have no impact since the weights are normalized in (5). We can however obtain an estimate of Z_t/Z_{t-1} at time t via

$$\frac{Z_t}{Z_{t-1}} = \int \frac{\gamma_t(\boldsymbol{\beta}_{t-1})}{\gamma_{t-1}(\boldsymbol{\beta}_{t-1})} \pi_{t-1}(\boldsymbol{\beta}_{t-1}) d\boldsymbol{\beta}_{t-1} \quad (7)$$

$$\approx \sum_{i=1}^N W_{t-1}^{(i)} w_t(\boldsymbol{\beta}_{t-1}^{(i)}, \boldsymbol{\beta}_t^{(i)}) = \frac{\widehat{Z}_t}{Z_{t-1}} \quad (8)$$

which allows one to estimate the unnormalized marginal density Z_t/Z_1 for each value of b_t via $\widehat{\frac{Z_t}{Z_1}} = \prod_{t=2}^T \widehat{\frac{Z_t}{Z_{t-1}}}$.

If we additionally give $\log b$ a uniform prior on $[\log b_T, \log b_1]$, these unnormalized marginal densities can also be used to weight the collections of particles at each time, allowing us to approximate the distribution of $\boldsymbol{\beta}$ with b marginalized out.

$$\tilde{\pi}_t^N(d\boldsymbol{\beta}) = \frac{1}{\sum_{t=1}^T \widehat{\frac{Z_t}{Z_1}}} \sum_{t=1}^T \frac{\widehat{Z}_t}{Z_1} \sum_{i=1}^N W_t^{(i)} \delta_{\boldsymbol{\beta}_t^{(i)}}(d\boldsymbol{\beta}) \quad (9)$$

Note that this is equivalent to systematic sampling of the values of b , and that an improper uniform prior for $\log b$ on $(-\infty, \infty)$ would induce an improper posterior for $\boldsymbol{\beta}$.

4.1 Further Details

For analysis of the data with $n = 500$ and $p = 50$, we use $N = 8192$ particles and we consider two settings of $a = 1$ and $a = 4$ with $T = 450$ and $T = 350$ respectively, with the lower T for $a = 4$ explained by the low marginal likelihood associated with small values of b when a is large. We set $b_t = 2(0.98)^{t-1}$ for $t \in \{1, \dots, T\}$. Given the nature of the posterior, a simple p -dimensional random walk Metropolis-Hastings kernel does not seem to converge quickly. We found that using a cycle of kernels that each update one element of β using a random walk proposal with variance 0.25 mixed better for a given amount of computation than a p -dimensional random walk proposal, which is consistent with Neal and Roberts (2006). To ensure that our kernels mix fast enough to get good results we construct K_t by cycling this cycle of kernels 5 times. This is computationally demanding since each step of the algorithm consists of $5Np$ MCMC steps, each of which has a complexity in $O(n)$, since the calculation of the likelihood requires us to update $\mathbf{X}\beta$ for a change in one element of β .

In fact, computation on a 2.67GHz Intel Xeon requires approximately 20 hours of computation time with these settings and $T = 450$. However, we implemented the algorithm to run on an NVIDIA 8800 GT graphics card using the Compute Unified Device Architecture (CUDA) parallel computing architecture, using the SMC sampler framework in Lee, Yau, Giles, Doucet, and Holmes (2010b) to take advantage of the ability to compute the likelihood in parallel for many particles. On this hardware, the time to run the algorithm is approximately 30 minutes for $T = 450$, giving around a 40 fold speedup.

The initial particles for $t = 1$ are obtained by running an MCMC algorithm targeting π_1 for a long time and picking suitably thinned samples to initialize the particles. The weights of each particle start off equal at $1/N$. While not ideal, trivial importance sampling estimates have a very low ESS for reasonable computational power and, in any case, the SMC sampler methodology requires us to have a rapidly mixing MCMC kernel for our resulting estimates to be good.

The particle with the highest posterior density at each time is used as an initial value for the MAP algorithm described in Section 2.1. The density of the estimated MAP from this initial value is compared with the MAP obtained by using the previous MAP as an initial value and the estimate with the higher posterior density is chosen as the estimated MAP. This has the effect of both removing variables that the original algorithm has not removed and keeping variables in that the original algorithm has not, as it is not possible for the MAP algorithm to explore the multimodal posterior density fully.

While the independence of the weights of the particles makes it possible to adapt the selection of b_t to allow aggressive changes in b while maintaining

a given ESS (Del Moral, Doucet, and Jasra, 2011, Bornn, Doucet, and Gottardo, 2010), we found that this scheme had a poor effect on the quality of our samples. This is primarily because we discovered that our cycle length for K_t was not large enough to get the particles to approximate π_t after a large change in b , despite a small change in ESS. Increasing the cycle length has an essentially linear effect on the computational resources required, so this appears to be a poor option for reducing the computational complexity of the method. In addition, having b_t on a logarithmic schedule enables us to use the simple approximation to the posterior of $\beta|\mathbf{X}, \mathbf{y}, a$ given by (9).

For the larger experiment with $n = 1859, p = 184$ and $a = 4$, we use the same parameters except that we used the schedule $b_t = (0.98)^{t-1}$ with $T = 250$. This was sufficient to capture a wide range of model complexities. Nevertheless, the running time for this experiment on our GPU was 7.5 hours. This demonstrates the key practical benefit to using GPUs, as a similar run on a single CPU would take upwards of 5 days.

4.2 Tuning the SMC sampler

Following exploratory analysis we found that $b_t = b_1(0.98)^{t-1}$ worked well, and that setting b_1 equal to one or two produced a diffuse enough prior when sample size was in the 100s or 1000s, as is the situation in GWAS.

In practice, one will have to experiment with the parameters of the SMC sampler in order to ensure satisfactory results. In particular, given an MCMC kernel K_t , one must decide how many cycles should be performed per time step to mix adequately in relation to how quickly one wishes to decrease b . In an ideal scenario, K_t produces essentially independent samples from π_t regardless of the current set of particles. This would allow one to set $\{b_t\}_{t=1}^T$ such that computational power is only spent in areas of b that are of interest. In practice, however, K_t will often only mix well locally and so small steps in b are desirable to ensure that the Monte Carlo error in the scheme does not dominate the results. It is difficult to characterize the relationship between the number of cycles of K_t to perform and how aggressively one can decrease b in general. However, in this context, one might not want to decrease b too quickly anyway, so that a smooth picture of the relationship between successive posteriors can be visualized.

The number of particles, N , is a crucial parameter that should be set at such a level that some fraction of N can provide a suitable approximation of each intermediate target density. In practice, the number of effective samples is lower than N and resampling when the ESS is low is an attempt to ensure that this number is controlled. It is standard practice to choose to resample particles when the ESS

is below CN for some $C \geq 1/2$. We choose to resample the particles when the ESS fell below $(3/4)N$.

A recommended way to test whether the sampler is performing acceptably is to run an MCMC scheme with fixed $b = b_t$ for some intermediate value of t and compare the samples one obtains after convergence with those produced by the SMC sampler. This is possible in this context because the posterior is not particularly troublesome to sample from for fixed b with a sensible kernel.

Finally it is worth noting that if the aim was to infer a single distribution given a single setting of $c = b/a$, or with a prior on c , one could achieve better results via one long MCMC run. The benefit of the SMC sampler is that it provides a unified method to obtain all of the exploratory statistics and estimates of the marginal likelihood in a single run.

5 Sparsity-Path-Analysis

Lasso plots of MAP estimates as functions of the penalty parameter λ are a highly informative visual tool for statisticians to gain an understanding of the relative predictive importance of covariates over a range of model complexities (Hastie, Tibshirani, and Friedman, 2009). However, MAP statistics are single point estimates that fail to convey the full extent of information contained within the joint posterior distribution. One key aim of this report is to develop full Bayesian exploratory tools that utilize the joint posterior information.

In this section we describe graphical displays of summary statistics contained in SPA results by first considering the $n = 500$, $p = 50$ data set described above. The indices of the variables with non-zero β 's are $\mathcal{J} = \{10, 14, 24, 31, 37\}$ with corresponding $\beta_{\mathcal{J}} = \{-0.2538, 0.4578, -0.1873, -0.1498, 0.0996\}$. In the $Gt(a, c)$ prior we found the $Gt(a = 1, c)$, or $Gt(a = 4, c)$, to be good default settings the degrees of freedom parameter, with both leading to heavy tails. We ran the SMC-GPU sampler described above and then generated a number of plots using the set of particles and their weights.

5.1 Path plots

In Fig. 4 we plot four graphs of summary statistics obtained from the SMC sampler across a range of scales $\log(c) \in (-8, 0)$.

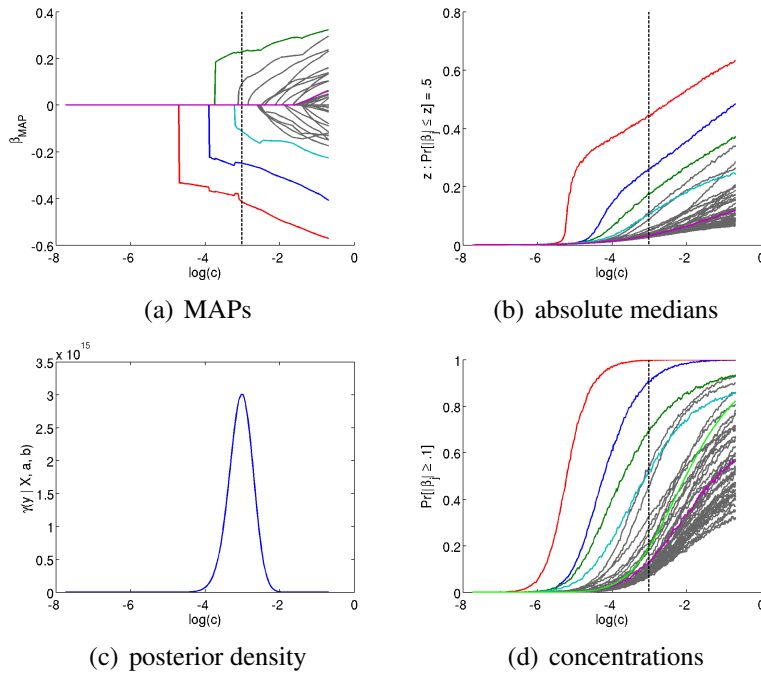


Figure 4: SPA plots with $a = 4$ degrees of freedom; using 50 markers around 18q. The 5 non-zero coefficients are indicated by non-grey lines. The vertical dashed line indicates the mode of the marginal likelihood, $p(c|\mathbf{X}, \mathbf{y}, a)$, for c as shown in plot (c).




Figure 5: Plot of marginal posterior of β_{24} as a function of c , c.f. the red curve in Fig. 4(a). We can see the bimodality and hence the jumping of the MAP mode in Fig. 4(a) as $\log c$ changes from $(-5, -4)$.

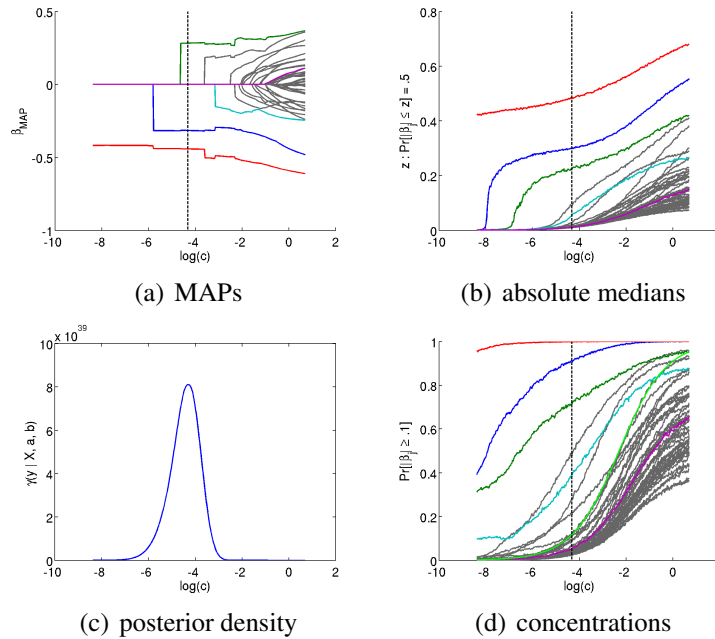


Figure 6: SPA plots with $a = 1$ degrees of freedom; using 50 markers around 18q. This Figure should be compared with Fig. 4 using $a = 4$ degrees of freedom.

- (a) MAP plots: In Fig. 4(a) we plot the MAP paths of all 50 coefficients moving from the most sparse models with $c = e^{-8}$ where all $\hat{\beta}_{MAP} = 0$ through to $c = 1$ when all $\hat{\beta}_{MAP} \neq 0$. The true non-zero coefficients are plotted in colors other than gray. One can observe that the MAP paths are non-smooth in $\log(c)$ such that the modes suddenly jump away from zero. This is a property of the non-log-concavity of the posterior distributions; such that the marginal posterior densities can be bimodal with one mode at $\beta_j = 0$ and another away from zero. As c increases the MAP mode at zero decreases and at some point the global mode switches to the mode at $\beta_j \neq 0$. This can be seen clearly in Fig. 5 where we plot the marginal posterior distribution of β_{24} over the range $\log c \in (-5, -4)$; c.f. β_{24} is shown as the red curve in Fig. 4(a). We can see in Fig. 5 the non-concavity of the posterior density and how the global mode jumps from 0 to $\approx -.04$ as $\log c$ is around -4.75 . As mentioned above the MAP is found by iterating the EM algorithm beginning from the particle with largest posterior density.
- (b) Median plots: In Fig. 4(b) we plot the absolute value of the median of the marginal distribution of β_j 's. This is a plot of $z_j(\log[c])$ vs. $\log c$, for $j = 1, \dots, p$, where,

$$z_j(c) = |\hat{F}_{\beta_j}^{-1}(0.5|c)|$$

where $F_{\beta_j}^{-1}(\cdot)$ is the inverse cumulative posterior distribution of β_j ,

$$F_{\beta_j}(x|c) = \int_{-\infty}^x p(\beta_j|\mathbf{X}, \mathbf{y}, a, c) d\beta_j$$

and where $p(\beta_j|\mathbf{X}, \mathbf{y}, a, c)$ is the marginal posterior distribution of β_j given c ; $p(\beta_j|\mathbf{X}, \mathbf{y}, a, c) = \int_{\mathbb{R}^{p-1}} p(\beta_j, \beta_{-j}|\mathbf{X}, \mathbf{y}, a, c) d\beta_{-j}$ where $\{-j\}$ indicates all indices other than the j th. The plot of absolute medians gives an indication of how quickly the posterior distributions are moving outward away from the origin as the precision of the prior decreases. By plotting on the absolute scale we are better able to compare the coefficients with one another and we also see that, unlike the MAPs, the medians are smooth functions of c .

- (c) Posterior for scale c : In Fig. 4(c) we show the marginal posterior distribution $p(c|\mathbf{X}, \mathbf{y}, a)$,

$$p(c|\mathbf{X}, \mathbf{y}, a) = \int_{\mathbb{R}^p} p(c, \boldsymbol{\beta}|\mathbf{X}, \mathbf{y}, a) d\boldsymbol{\beta}.$$

The posterior on c graphs the relative evidence for particular prior scale. The mode of Fig. 4(c),

$$\tilde{c} = \arg \max_c p(c|\mathbf{X}, \mathbf{y}, a)$$

is indicated on plots (a),(b),(d) by a vertical dashed line.

- (d) Concentration plots: In Fig. 4(d) we plot the concentration of the marginal posteriors of β_j 's around the origin, as well as the *a priori* concentration (light green). In particular, for user specified tolerance Δ , this is a plot of $V(c)$ vs. c where,

$$V(c) = 1 - \int_{-\Delta}^{\Delta} p(\beta_j|\mathbf{X}, \mathbf{y}, a, c) d\beta_j = 1 - Pr(\beta_j \in (-\Delta, \Delta)|\mathbf{X}, \mathbf{y}, a, c)$$

This is a direct measure of the predictive relevance of the corresponding covariate (genotype). In Fig. 4(d) we have set $\Delta = 0.1$ although this is a user set parameter that can be specified according to the appropriate level below which a variable is not deemed important.

Taken together the SPA plots highlight the influential genotypes and the relative evidence for their predictive strength. We can also gain an understanding of the support in the data for differing values of $\log(c)$. Having generated the SPA plots it is interesting to compare the results of changing the degrees of freedom from $Gt(a = 4, c)$ to $Gt(a = 1, c)$. In Fig. 6 we show plots corresponding to Fig. 4 but using $a = 1$. We can see that, as expected, $Gt(a = 1, c)$ produces sparser solutions with only three non-zero MAP estimates at the mode of the posterior, \tilde{c} . Moreover,

comparing the concentration plots Fig. 6(d) and Fig. 4(d) at the marginal mode of c we see that for $a = 1$ we see much greater dispersion in the concentration plot.

5.2 Individual coefficient plots

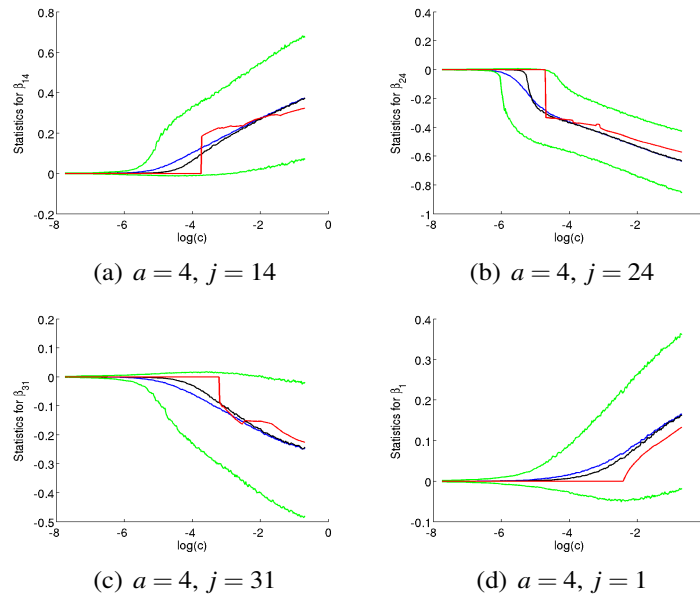


Figure 7: Stats for individual coefficients from Fig. 4 with $a = 4$. We plot 90% credible intervals (green), median (black), mean (blue) and MAP (red).

Examination of the plots in Fig. 4 may highlight to the statistician some interesting predictors to explore in greater detail. Individual coefficient plots of summary statistics can be produced to provide greater information on the posterior distributions. In Fig. 7 we show summary plots for four representative coefficients with their 90% credible intervals (green), median (black), mean (blue) and MAP (red). These are obtained from the set of weighted SMC particles. We can see that as expected the mean and median are smooth functions of the prior scale, which the MAP can exhibit the characteristic jumping for bimodal densities. In Fig. 8 we show corresponding density estimates. These coefficients were chosen to represent markers with strong association Fig. 7(a), weaker association Fig. 7(b),(c) and no association Fig. 7(d). We can see in the plots for Fig. 7(d) and Fig. 8(d) that for a null marker with no association signal the MAP appears to be much smoother in $\log(c)$.

Equivalent plots but for $a = 1$ are shown in Fig. 9 and Fig. 10 where we see the greater sparsity induced by the heavier tails of the $Gt(a = 1, c)$ relative to $Gt(a = 4, c)$ at the corresponding marginal mode of c .




Figure 8: Posterior density plots corresponding to Fig. 7, $a = 4$

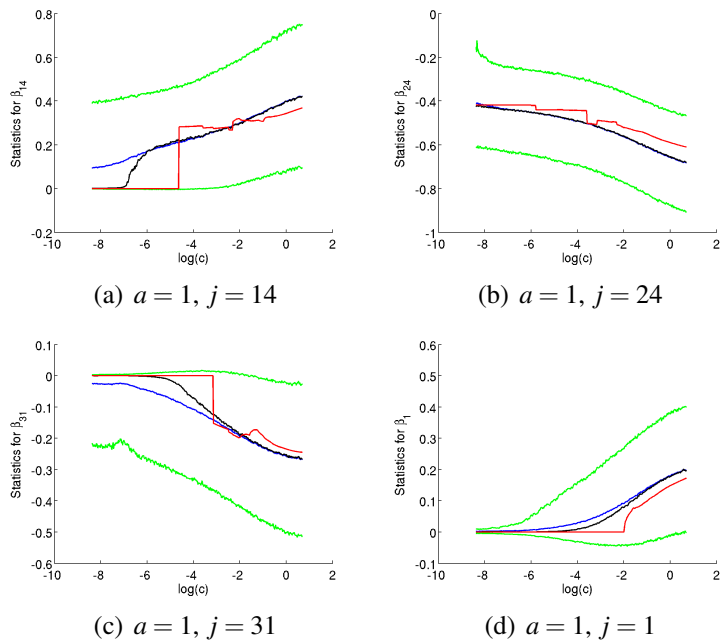


Figure 9: Stats for individual coefficients from Fig. 6 with $a = 1$. Compare with Fig. 7 where $a = 4$.




Figure 10: Posterior density plots corresponding to Fig. 9, $a = 1$

5.3 Marginal plots

The SMC sampler also allows us to estimate the marginal posterior probability of β , using (9), integrating over the uncertainty in c ,

$$p(\beta|\mathbf{X}, \mathbf{y}, a) = \int_0^\infty p(\beta, c|\mathbf{X}, \mathbf{y}, a)dc$$

Moreover we can also calculate the marginal posterior concentration away from zero, for given tolerance Δ as,

$$V = 1 - \int_0^\infty \int_{-\Delta}^\Delta p(\beta_j, c|\mathbf{X}, \mathbf{y}, a)d\beta_j dc$$

In Fig. 11 we plot summaries for the marginal posterior distributions of β as well as the marginal concentrations, for $a = 4$ and $a = 1$. We can see in Fig. 11 that the marginal plots provide a useful overview of the relative importance of each variable. It may be difficult to distinguish the medians from the means in the visualization of the marginal posterior distributions on paper, but it is easy to see the difference on a computer screen.

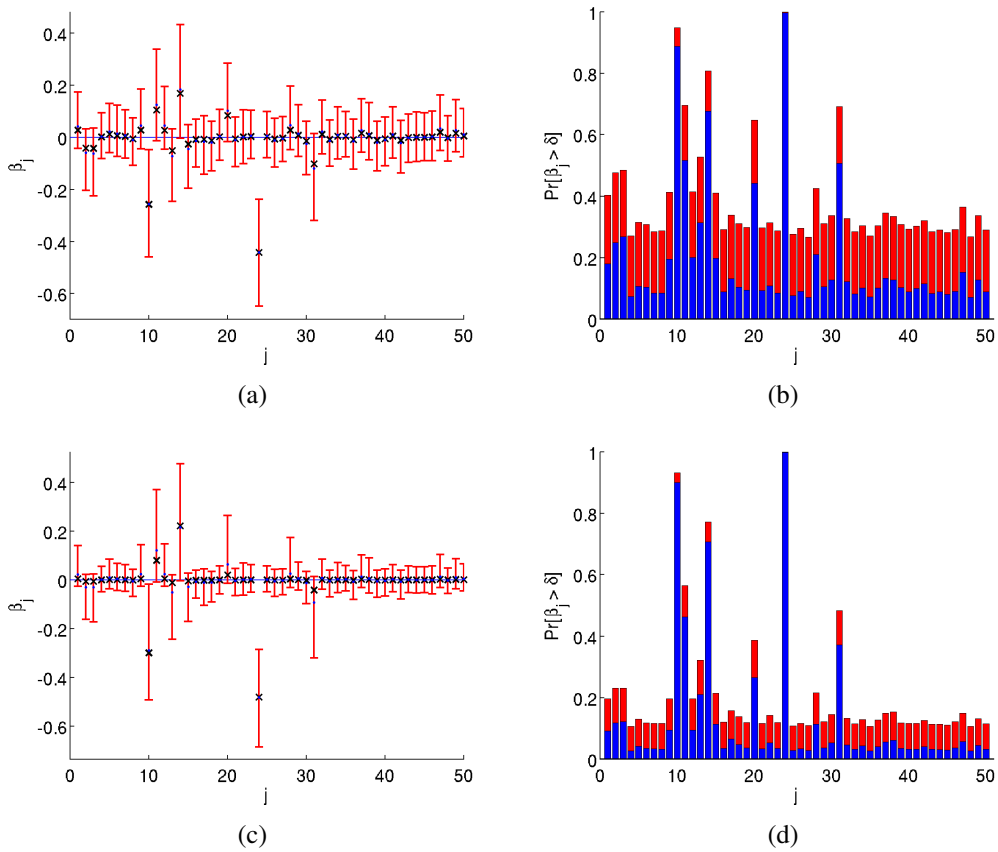


Figure 11: Marginal plots: (a) the summary stats from marginal posterior distributions showing medians (crosses), means (blue dots), and 90% credible intervals (bars) for $a = 4$ degrees of freedom; (b) the marginal concentrations $\Delta = 0.05$ red bars, $\Delta = 0.1$ (blue bars) for $a = 4$; (c) marginal posteriors as in (a) but with $a = 1$ prior; (d) marginal concentrations as in (b) but for $a = 1$, prior.

5.4 Comparison to double-exponential prior, $a \rightarrow \infty$

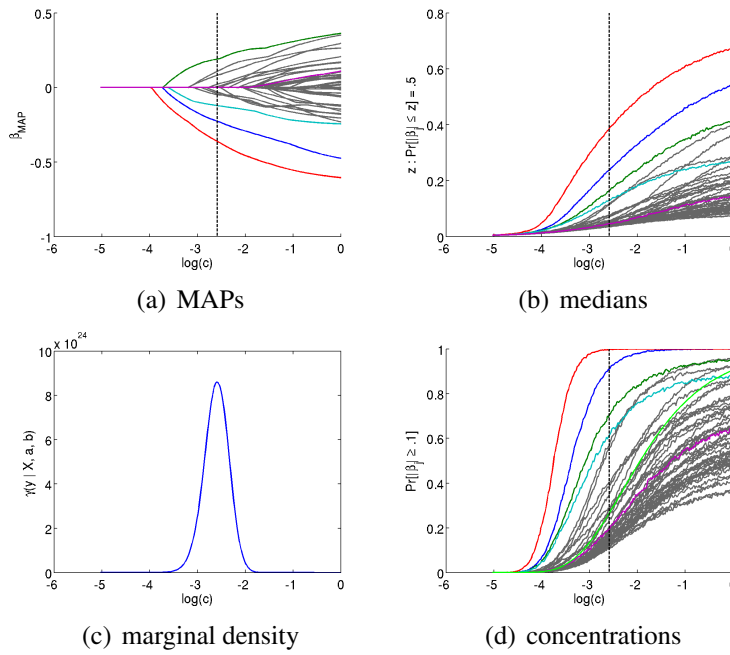


Figure 12: SPA plots for the double-exponential prior, $Gt(a \rightarrow \infty, c)$.

It is informative to compare the results for the $Gt(a, c)$ prior above with that of the double-exponential prior $p(\beta) \propto \exp(-|\beta|/c)$ which is obtained as a generalized t prior with $q = 1$ and $a \rightarrow \infty$. In Fig. 12 we show SPA plots for this case. We can see the much smoother paths of the MAP compared with Fig. 4. This can also be seen in the individual coefficient plots shown in Fig. 13 and Fig. 14. It is interesting to investigate in more detail the posterior density of β_{24} the coefficient with strongest evidence of association. This is shown in Fig. 15 and should be compared to Fig. 5 for $a = 4$. The concavity of the double-exponential prior causes the posterior distribution to move smoothly towards being concentrated around 0 as c decreases, unlike the behaviour seen with the generalized t prior.

5.5 Large Data Set

We next analysed the larger data set with $n = 1859$ and $p = 184$. The indices of the non-zero coefficients are $\mathcal{J} = \{108, 22, 5, 117, 162\}$ with the same values retained for $\beta_{\mathcal{J}} = \{-0.2538, 0.4578, -0.1873, -0.1498, 0.0996\}$. The SPA plots are shown in Fig. 16 with corresponding individual coefficient plots in Fig. 17 and

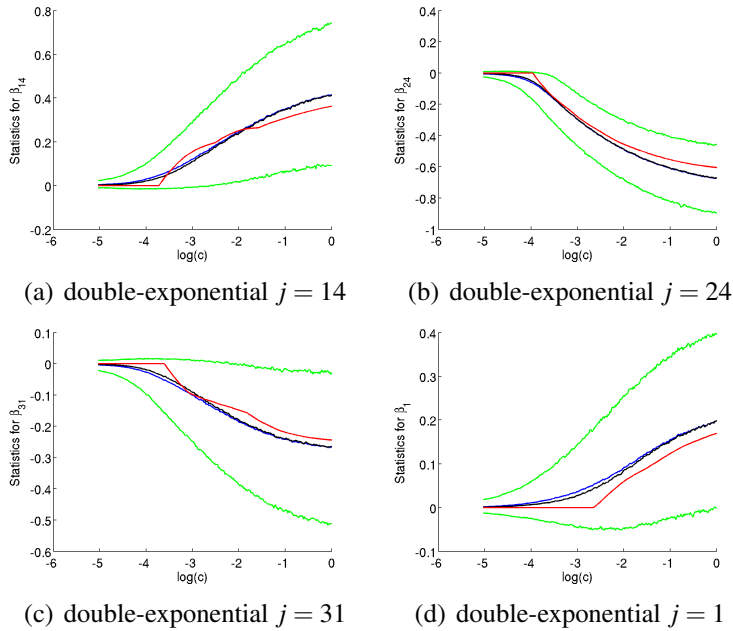


Figure 13: Stats for individual coefficients from Fig. 12 with double-exponential prior. Compare with Fig. 7 where $a = 4$.

Fig. 18. We can see in Fig. 16 that there is much stronger evidence that coefficients $\{108, 22, 5, 117\}$ are important. Interestingly variable 162 is still masked by other predictors. In comparison with Fig. 4(c) we observe that a smaller value of c is supported by the larger data set.

From Fig. 16(d) we can see that one null-coefficient, β_{107} , (with true $\beta_{107} = 0$) appears to have some evidence of predictive importance. In Fig. 19 we show summaries of the marginal distributions and concentration plots having integrated over the uncertainty in c . The null-coefficient SNP turns out to be situated adjacent to a predictive SNP $\beta_{108} = -0.2538$ shown as the blue line in Fig. 16(a), and the two markers are in high LD with correlation coefficient $C(X_{107}, X_{108}) = 0.87$, leaving some ambiguity in the source of the association signal. The plot of absolute medians Fig. 16(b) helps partly resolve this issue pointing to only four clear association signals around the marginal mode of c . We can drill down a little further and plot summary statistics for the posterior distribution for β_{107} , shown in Fig. 20. In Fig. 20 we can see that the MAP (red line) holds at zero for almost all values of c , adding support that this is a “null” marker, although the credible intervals (green lines) show there is considerable uncertainty in the actual value. In such a way we can see how complex association signal can be explored, gaining a better understanding of the source of association signal.

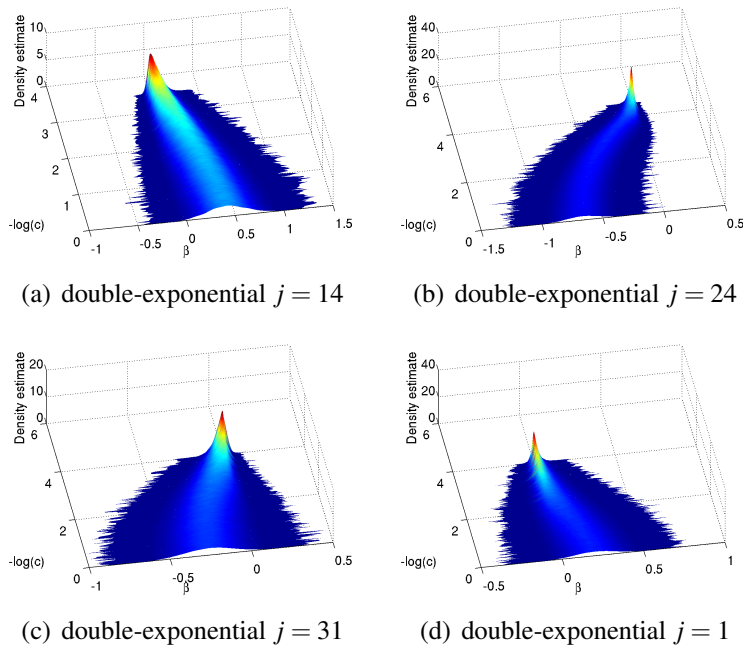


Figure 14: Posterior density plots corresponding to Fig. 13 with double-exponential prior.




Figure 15: Plot of marginal posterior of β_{24} as a function of c for double-exponential prior $a \rightarrow \infty$. Under the double-exponential prior, $a \rightarrow \infty$, we see the posterior remains log-concave and hence the MAP is a smooth function of $\log c$; as compared to Fig. 5.

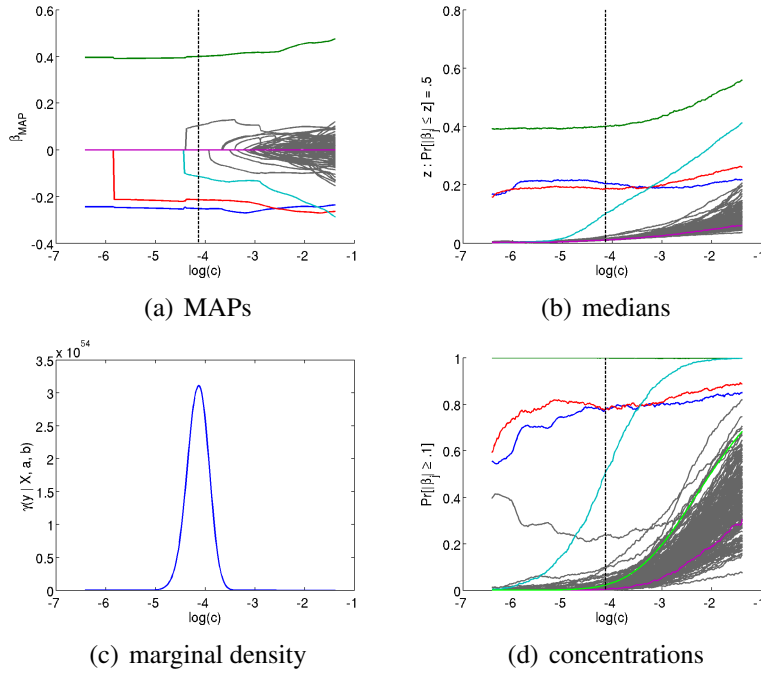


Figure 16: SPA plots using the larger data set $n = 1859$, $p = 184$

6 Discussion

We have presented an exploratory approach using generalized t priors that we call Bayesian sparsity-path-analysis to aid in the understanding of genetic association data. The approach involves graphical summaries of posterior densities of coefficients obtained by sweeping over a range of prior precisions, including values with low posterior probability, in order to characterize the space of models from the sparse to the most complex.

This use of SMC methodology is ideally suited to the inference task, by indexing the SMC sampler on the scale of the prior. The resulting collection of weighted particles provides us with approximations for the coefficient posterior distributions for each scale in addition to estimates of the marginal likelihood and allows for improved robustness in MAP computation. The simulations are computationally demanding and would take days worth of run-time using conventional single-threaded CPU processing. To alleviate this we make use of many-core GPU parallel processing producing around a 20-40 fold improvement in run-times. This has real benefit in allowing for data analysis within a working day for the larger data set we considered.

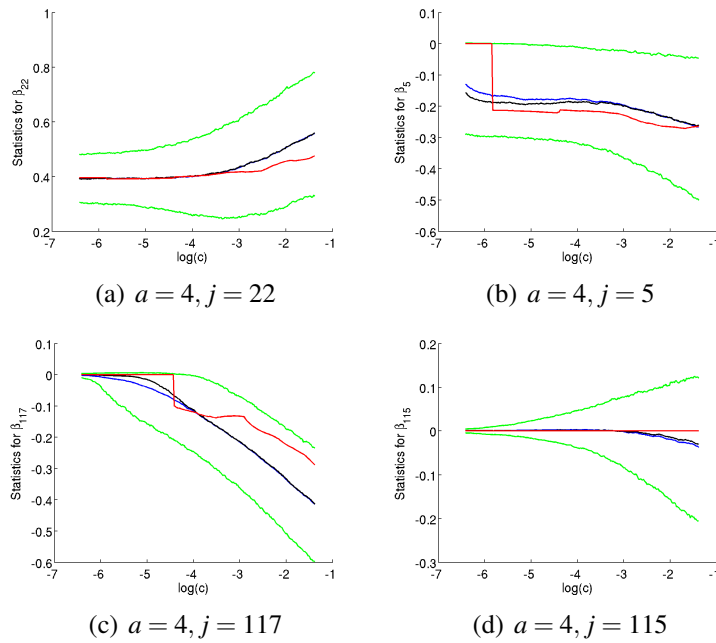


Figure 17: Summary statistics for posterior distributions for individual coefficients, 90% credible intervals (green), median (black), mean (blue) and MAP (red).




Figure 18: Posterior density plots of coefficients in Fig. 17




Figure 19: Marginal plots: (a) the summary stats from marginal posterior distributions showing medians (crosses), means (blue dots), and 90% credible intervals (bars) for $a = 4$ degrees of freedom; (b) the marginal concentrations $\Delta = 0.05$ (red bars), $\Delta = 0.1$ (blue bars) for $a = 4$




Figure 20: Statistics for individual coefficient β_{107} .

References

- Armagan, A., D. Dunson, and J. Lee (2011): “Generalized double Pareto shrinkage,” ArXiv:1104.0861.
- Ayers, K. L. and H. J. Cordell (2010): “SNP selection in genome-wide and candidate gene studies via penalized logistic regression,” *Genetic Epidemiology*, 34, 879–891.
- Balding, D. J. (2006): “A tutorial on statistical methods for population association studies,” *Nature Reviews Genetics*, 7, 81–791.
- Bornn, L., A. Doucet, and R. Gottardo (2010): “An efficient computational approach for prior sensitivity analysis and cross-validation,” *Canadian Journal of Statistics*, 38, 47–64.
- Candès, E. J., M. B. Wakin, and S. P. Boyd (2008): “Enhancing sparsity by reweighted ℓ_1 minimization,” *Journal of Fourier Analysis and Applications*, 14, 877–905.
- Caron, F. and A. Doucet (2008): “Sparse bayesian nonparametric regression,” in *Proceedings of the 25th international conference on Machine learning*, ACM, 88–95.
- Carvalho, C., N. Polson, and J. Scott (2010): “The horseshoe estimator for sparse signals,” *Biometrika*, 97, 465.
- Cevher, V. (2009): “Learning with compressible priors,” in Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, eds., *Advances in Neural Information Processing Systems* 22, 261–269.
- Chartrand, R. and W. Yin (2008): “Iteratively reweighted algorithms for compressive sensing,” in *33rd International Conference on Acoustics, Speech, and Signal Processing (ICASSP)*.
- Crooks, G. E. (1998): “Nonequilibrium measurements of free energy differences for microscopically reversible Markovian systems,” *Journal of Statistical Physics*, 90, 1481–1487.
- Del Moral, P., A. Doucet, and A. Jasra (2006): “Sequential Monte Carlo samplers,” *Journal of the Royal Statistical Society B*, 68, 411–436.
- Del Moral, P., A. Doucet, and A. Jasra (2011): “An adaptive sequential Monte Carlo method for approximate Bayesian computation,” *Statistics and Computing*, to appear.
- Fan, J. and R. Li (2001): “Variable selection via nonconcave penalized likelihood and its oracle properties,” *Journal of the American Statistical Association*, 96, 1348–1360.
- Fridley, B. (2009): “Bayesian variable and model selection methods for genetic association studies,” *Genetic Epidemiology*, 33, 27–37.

- George, E. I. and R. E. McCulloch (1993): “Variable selection via Gibbs sampling,” *Journal of the American Statistical Association*, 88, 881–889.
- Griffin, J. and P. Brown (2010): “Inference with normal-gamma prior distributions in regression problems,” *Bayesian Analysis*, 5, 171–188.
- Griffin, J. E. and P. J. Brown (2007): “Bayesian adaptive lassos with non-convex penalization,” Technical report, IMSAS, University of Kent.
- Hastie, T., R. Tibshirani, and J. Friedman (2009): *The elements of statistical learning: data mining, inference, and prediction*, Springer Verlag.
- Hoggart, C. J., J. C. Whittaker, M. De Iorio, and D. J. Balding (2008): “Simultaneous analysis of all snps in genome-wide and re-sequencing association studies,” *PLOS Genetics*, 4, e1000130.
- Lee, A., F. Caron, A. Doucet, and C. Holmes (2010a): “A hierarchical Bayesian framework for constructing sparsity-inducing priors,” ArXiv:1009.1914.
- Lee, A., C. Yau, M. B. Giles, A. Doucet, and C. C. Holmes (2010b): “On the utility of graphics cards to perform massively parallel simulation of advanced Monte Carlo methods,” *Journal of Computational and Graphical Statistics*, 19, 769–789.
- Liu, J. S. and R. Chen (1995): “Blind deconvolution via sequential imputations,” *Journal of the American Statistical Association*, 90, 567–576.
- McDonald, J. and W. Newey (1988): “Partially adaptive estimation of regression models via the generalized t distribution,” *Econometric Theory*, 4, 428–457.
- Neal, P. and G. Roberts (2006): “Optimal scaling for partially updating MCMC algorithms,” *Annals of Applied Probability*, 16, 475–515.
- Neal, R. M. (2001): “Annealed importance sampling,” *Statistics and Computing*, 11, 125–139.
- Park, T. and G. Casella (2008): “The Bayesian Lasso,” *Journal of the American Statistical Association*, 103, 681–686.
- Stephens, M. and D. Balding (2009): “Bayesian statistical methods for genetic association studies,” *Nature Reviews Genetics*, 10, 681–690.
- Tibshirani, R. (1996): “Regression shrinkage and selection via the lasso,” *Journal of the Royal Statistical Society B*, 58, 267–288.
- Tibshirani, R. (2011): “Regression shrinkage and selection via the lasso: a retrospective,” *Journal of the Royal Statistical Society B*, 73, 273–282.
- Wilson, M. A., E. S. Iversen, M. A. Clyde, S. C. Schmidler, and J. M. Schildkraut (2010): “Bayesian model search and multilevel inference for snp association studies,” *Annals of Applied Statistics*, 4, 1342–1364.
- Zou, H. (2006): “The adaptive lasso and its oracle properties,” *Journal of the American Statistical Association*, 101, 1418–1429.
- Zou, H. and R. Li (2008): “One-step sparse estimates in nonconcave penalized likelihood models,” *Annals of Statistics*, 36, 1509–1533.