



HHS Public Access

Author manuscript

Biometrika. Author manuscript; available in PMC 2017 April 21.

Published in final edited form as:

Biometrika. 2016 December ; 103(4): 985–991. doi:10.1093/biomet/asw042.

Fast sampling with Gaussian scale-mixture priors in high-dimensional regression

Anirban Bhattacharya, Antik Chakraborty, and Bani K. Mallick

Department of Statistics, Texas A&M University, College Station, Texas, 77843, USA

Summary

We propose an efficient way to sample from a class of structured multivariate Gaussian distributions. The proposed algorithm only requires matrix multiplications and linear system solutions. Its computational complexity grows linearly with the dimension, unlike existing algorithms that rely on Cholesky factorizations with cubic complexity. The algorithm is broadly applicable in settings where Gaussian scale mixture priors are used on high-dimensional parameters. Its effectiveness is illustrated through a high-dimensional regression problem with a horseshoe prior on the regression coefficients. Other potential applications are outlined.

Keywords

Confidence interval; Gaussian scale mixture; Global-local prior; Shrinkage; Sparsity

1. Introduction

Continuous shrinkage priors have recently received significant attention as a mechanism to induce approximate sparsity in high-dimensional parameters, and can mostly be expressed as global-local scale mixtures of Gaussian distributions (Polson & Scott, 2010; Bhattacharya et al., 2015). These global-local priors (Polson & Scott, 2010) aim to shrink noise coefficients while retaining any signal, thereby providing an approximation to the operating characteristics of discrete mixture priors (George & McCulloch, 1997; Scott & Berger, 2010), which allow a subset of the parameters to be exactly zero.

A major attraction of global-local priors has been computational efficiency and simplicity. Posterior inference poses a stiff challenge for discrete mixture priors in moderate to high-dimensional settings, but the scale-mixture representation of global-local priors allows parameters to be updated in blocks via a fairly automatic Gibbs sampler in a wide variety of problems. These include regression (Caron & Doucet, 2008; Armagan et al., 2013), variable selection (Hahn & Carvalho, 2015), wavelet denoising (Polson & Scott, 2010), factor models and covariance estimation (Bhattacharya & Dunson, 2011; Pati et al., 2014), and time series (Durante et al., 2014). Rapid mixing and convergence of the resulting Gibbs sampler for specific classes of priors has been recently established in the high-dimensional regression context by Khare & Hobert (2013) and Pal & Khare (2014). Moreover, recent results suggest that a subclass of global-local priors can achieve the same minimax rates of posterior concentration as the discrete mixture priors in high-dimensional estimation problems (Bhattacharya et al., 2015; van der Pas et al., 2014; Pati et al., 2014).

In this article, we focus on computational aspects of global-local priors in the high-dimensional linear regression setting

$$y = X\beta + \varepsilon, \quad \varepsilon \sim N(0, \sigma^2 I_n), \quad (1)$$

where $X \in \mathfrak{R}^{n \times p}$ is a $n \times p$ matrix of covariates with the number of variables p potentially much larger than the sample size n . A global-local prior on β assumes that

$$\beta_j | \lambda_j, \tau, \sigma \sim N(0, \lambda_j^2 \tau^2 \sigma^2) \quad (j=1, \dots, p), \quad (2)$$

$$\lambda_j \sim f \quad (j=1, \dots, p), \quad (3)$$

$$\tau \sim g, \quad \sigma \sim h, \quad (4)$$

where f , g and h are densities supported on $(0, \infty)$. The λ_j s are usually referred to as local scale parameters while τ is a global scale parameter. Different choices of f and g lead to different classes of priors. For instance, a half-Cauchy distribution for f and g leads to the horseshoe prior of Carvalho et al. (2010). In the $p \gg n$ setting where most entries of β are assumed to be zero or close to zero, the choices of f and g play a key role in controlling the effective sparsity and concentration of the prior and posterior distributions (Polson & Scott, 2010; Pati et al., 2014).

Exploiting the scale-mixture representation (2), it is straightforward in principle to formulate a Gibbs sampler: the conditional posterior of β given $\lambda = (\lambda_1, \dots, \lambda_p)^T$, τ and σ is given by

$$\beta | y, \lambda, \tau, \sigma \sim N(A^{-1} X^T y, \sigma^2 A^{-1}), \quad A = (X^T X + \Lambda_*^{-1}), \quad \Lambda_* = \tau^2 \text{diag}(\lambda_1^2, \dots, \lambda_p^2). \quad (5)$$

Further, the p local scale parameters λ_j have conditionally independent posteriors and hence $\lambda = (\lambda_1, \dots, \lambda_p)^T$ can be updated in a block by slice sampling (Polson et al., 2014) if conditional posteriors are unavailable in closed form. However, unless care is exercised, sampling from (5) can be expensive for large values of p . Existing algorithms (Rue, 2001) to sample from (5) face a bottleneck for large p to perform a Cholesky decomposition of A at each iteration. One cannot resort to precomputing the Cholesky factors since the matrix Λ_* in (5) changes at each iteration. In this article, we present an exact sampling algorithm for Gaussian distributions (5) which relies on data augmentation. We show that the computational complexity of the algorithm scales linearly in p .

2. The algorithm

Suppose we aim to sample from $N_p(\mu, \Sigma)$, with

$$\Sigma = (\Phi^T \Phi + D^{-1})^{-1}, \quad \mu = \Sigma \Phi^T \alpha, \quad (6)$$

where $D \in \mathcal{R}^{p \times p}$ is symmetric positive definite, $\Phi \in \mathcal{R}^{n \times p}$, and $\alpha \in \mathcal{R}^{n \times 1}$; (5) is a special case of (6) with $\Phi = X/\sigma$, $D = \sigma^2 \Lambda^*$ and $\alpha = y/\sigma$. A similar sampling problem arises in all the applications mentioned in §1, and the proposed approach can be used in such settings. In the sequel, we do not require D to be diagonal, but we assume that D^{-1} is easy to calculate and that it is straightforward to sample from $N(0, D)$. This is the case, for example, if D corresponds to the covariance matrix of an AR(q) process or a Gaussian Markov random field.

Letting $Q = \Sigma^{-1} = (\Phi^T \Phi + D^{-1})$ denote the precision, or inverse covariance, matrix and $b = \Phi^T \alpha$, we can write $\mu = Q^{-1} b$. Rue (2001) proposed an efficient algorithm to sample from a $N(Q^{-1} b, Q^{-1})$ distribution that avoids explicitly calculating the inverse of Q , which is computationally expensive and numerically unstable. Instead, the algorithm in §3.1.2 of Rue (2001) performs a Cholesky decomposition of Q and uses the Cholesky factor to solve a series of linear systems to arrive at a sample from the desired Gaussian distribution. The original motivation was to efficiently sample from Gaussian Markov random fields where Q has a banded structure, so the Cholesky factor and the subsequent linear system solvers can be computed efficiently. Since $Q = (\Phi^T \Phi + D^{-1})$ does not have any special structure in the present setting, the Cholesky factorization has complexity $\mathcal{O}(p^3)$ (Golub & van Loan, 1996, Chapter 4.2.3), and becomes prohibitive for large p . We present an alternative exact mechanism to sample from a Gaussian distribution with parameters as in (6) below:

Algorithm 1

- i. Sample $u \sim N(0, D)$ and $\delta \sim N(0, I_n)$ independently.
- ii. Set $v = \Phi u + \delta$.
- iii. Solve $(\Phi D \Phi^T + I_n) w = (a - v)$ to obtain w .
- iv. Set $\theta = u + D \Phi^T w$.

Proposition 1—Suppose θ is obtained by following Algorithm 1. Then, $\theta \sim N(\mu, \Sigma)$, where μ and Σ are as in (6).

Proof: By the Sherman–Morrison–Woodbury identity (Hager, 1989) and some algebra, $\mu = D \Phi^T (\Phi D \Phi^T + I_n)^{-1} \alpha$. By construction, $v \sim N(0, \Phi D \Phi^T + I_n)$. Combining steps (iii) and (iv) of Algorithm 1, we have $\theta = u + D \Phi^T (\Phi D \Phi^T + I_n)^{-1} (a - v)$. Hence θ has a normal distribution with mean $D \Phi^T (\Phi D \Phi^T + I_n)^{-1} \alpha = \mu$. Since $\text{cov}(u, v) = D \Phi^T$, we obtain $\text{cov}(\theta) = D - D \Phi^T (\Phi D \Phi^T + I_n)^{-1} \Phi D = \Sigma$, again by the Sherman–Morrison–Woodbury identity. This completes the proof; a constructive proof is given in the Appendix.

While Algorithm 1 is valid for all n and p , the computational gains are biggest when $p \gg n$ and $N(0, D)$ is easily sampled. Indeed, the primary motivation is to use data augmentation to cheaply sample $\zeta = (v^T, u^T)^T \in \mathbb{R}^{n+p}$ and obtain a desired sample from (6) via linear transformations and marginalization. When D is diagonal, as in the case of global-local priors (2), the complexity of the proposed algorithm is $\mathcal{O}(n^2p)$; the proof uses standard results about complexity of matrix multiplications and linear system solutions; see §1.3 & 3.2 of Golub & van Loan (1996). For non-sparse D , calculating $D\Phi^T$ has a worst-case complexity of $\mathcal{O}(np^2)$, which is the dominating term in the complexity calculations. In comparison to the $\mathcal{O}(p^3)$ complexity of the competing algorithm in Rue (2001), Algorithm 1 therefore offers huge gains when $p \gg n$. For example, to run 6000 iterations of a Gibbs sampler for the horseshoe prior (Carvalho et al., 2010) with sample size $n = 100$ in MATLAB on a INTEL(E5-2690) 2.9 GHz machine with 64 GB DDR3 memory, Algorithm 1 takes roughly the same time as Rue (2001) when $p = 500$ but offers a speed-up factor of over 250 when $p = 5000$. MATLAB code for the above comparison and subsequent simulations is available at <https://github.com/antik015/Fast-Sampling-of-Gaussian-Posteriors.git>.

The first line of the proof implies that Algorithm 1 outputs μ if one sets $u = 0$, $\delta = 0$ in step (i). The proof also indicates that the log-density of (6) can be efficiently calculated at any $x \in \mathbb{R}^p$. Indeed, since Σ^{-1} is readily available, $x^T \Sigma^{-1} x$ and $x^T \Sigma^{-1} \mu$ are cheaply calculated and $\log |\Sigma^{-1}|$ can be calculated in $\mathcal{O}(n^3)$ steps using the identity $|\mathbf{I}_r + AB| = |\mathbf{I}_s + BA|$ for $A \in \mathbb{R}^{r \times s}$, $B \in \mathbb{R}^{s \times r}$. Finally, from the proof, $\mu^T \Sigma^{-1} \mu = a^T \Phi \Sigma^{-1} \Phi^T a = a^T \Phi \{D - D\Phi^T(\Phi D\Phi^T + \mathbf{I}_n)^{-1} \Phi D\} \Phi^T a = a^T \{(\Phi D\Phi^T + \mathbf{I}_n)^{-1} \Phi D\} \Phi^T a$, which can be calculated in $\mathcal{O}(n^3)$ operations.

3. Frequentist operating characteristics in high dimensions

The proposed algorithm provides an opportunity to compare the frequentist operating characteristics of shrinkage priors in high-dimensional regression problems. We compare various aspects of the horseshoe prior (Carvalho et al., 2010) to frequentist procedures and obtain highly promising results. We expect similar results for the Dirichlet–Laplace (Bhattacharya et al., 2015), normal-gamma (Griffin & Brown, 2010) and generalized double-Pareto (Armagan et al., 2013) priors, which we hope to report elsewhere.

We first report comparisons with smoothly clipped absolute deviation (Fan & Li, 2001) and minimax concave penalty (Zhang, 2010) methods. We considered model (1) with $n = 200$, $p = 5000$ and $\sigma = 1.5$. Letting x_i denote the i th row of X , the x_i s were independently generated from $N_p(0, \Sigma)$, with (i) $\Sigma = \mathbf{I}_p$ and (ii) $\Sigma_{jj} = 1$, $\Sigma_{jj'} = 0.5$ ($j, j' = 1, \dots, p$), compound symmetry. The true β_0 had 5 non-zero entries in all cases, with the non-zero entries having magnitude (a) $\{1.5, 1.75, 2, 2.25, 2.5\}$ and (b) $\{0.75, 1, 1.25, 1.5, 1.75\}$, multiplied by a random sign. For each case, we considered 100 simulation replicates. The frequentist penalization approaches were implemented using the R package `ncvreg` via 10-fold cross-validation. For the horseshoe prior, we considered the posterior mean and the point-wise posterior median as point estimates. Figures 1 and 2 report boxplots for ℓ_1 , $\|\hat{\beta} - \beta_0\|_1$, ℓ_2 , $\|\hat{\beta} - \beta_0\|_2$, and prediction, $\|X\hat{\beta} - X\beta_0\|_2$, errors across the 100 replicates for the two signal strengths. The horseshoe prior is highly competitive across all simulation settings, in particular when the signal strength is weaker. An interesting observation is the somewhat

superior performance of the pointwise median even under an ℓ_2 loss; a similar fact has been observed about point mass mixture priors (Castillo & van der Vaart, 2012) in high dimensions. We repeated the simulation with $p = 2500$ with similar conclusions. Overall, out of the 24 settings, the horseshoe prior had the best average performance over the simulation replicates in 22 cases.

While there is now a huge literature on penalized point estimation, uncertainty characterization in $p > n$ settings has received attention only recently (Zhang & Zhang, 2014; van de Geer et al., 2014; Javanmard & Montanari, 2014). Although Bayesian procedures provide an automatic characterization of uncertainty, the resulting credible intervals may not possess the correct frequentist coverage in nonparametric/high-dimensional problems (Szabó et al., 2015). This led us to investigate the frequentist coverage of shrinkage priors in $p > n$ settings; it is trivial to obtain element-wise credible intervals for the β_j s from the posterior samples. We compared the horseshoe prior with van de Geer et al. (2014) and Javanmard & Montanari (2014), which can be used to obtain asymptotically optimal elementwise confidence intervals for the β_j s. We considered a similar simulation scenario as before. We let $p \in \{500, 1000\}$, and considered a Toeplitz structure, $\Sigma_{jj'} = 0.9^{|j-j'|}$, for the covariate design (van de Geer et al., 2014) in addition to the independent and compound symmetry cases stated already. The first two rows of Table 1 report the average coverage percentages and $100 \times$ lengths of confidence intervals over 100 simulation replicates, averaged over the 5 signal variables. The last two rows report the same averaged over the $(p - 5)$ noise variables.

Table 1 shows that the horseshoe prior gives superior performance. An attractive adaptive property of shrinkage priors emerges, where the lengths of the intervals automatically adapt between the signal and noise variables, maintaining the nominal coverage. The frequentist procedures seem to yield approximately equal sized intervals for the signals and noise variables. The default choice of the tuning parameter $\lambda \asymp (\log p/n)^{1/2}$ suggested in van de Geer et al. (2014) seemed to provide substantially poorer coverage for the signal variables at the cost of improved coverage for the noise, and substantial tuning was required to arrive at the coverage probabilities reported. The default approach of Javanmard & Montanari (2014) produced better coverages for the signals compared to van de Geer et al. (2014). The horseshoe and other shrinkage priors are free of tuning parameters. The same procedure used for estimation automatically provides valid frequentist uncertainty characterization.

4. Discussion

Our numerical results warrant additional numerical and theoretical investigations into properties of shrinkage priors in high dimensions. The proposed algorithm can be used for essentially all the shrinkage priors in the literature and should prove useful in an exhaustive comparison of existing priors. Its scope extends well beyond linear regression. For example, extensions to logistic and probit regression are immediate using standard data augmentation tricks (Albert & Chib, 1993; Holmes & Held, 2006). Multivariate regression problems where one has a matrix of regression coefficients can be handled by block updating of the vectorized coefficient matrix; even if $p < n$, the number of regression coefficients may be large if the dimension of the response is moderate. Shrinkage priors have been used as a

prior of factor loadings in Bhattacharya & Dunson (2011), who update the $p > n$ rows of the factor loadings independently, exploiting the assumption of independence in the idiosyncratic components, but their algorithm does not extend to approximate factor models, where the idiosyncratic errors are dependent. The proposed algorithm can be adapted to such situations by block updates of the vectorized loadings. Finally, we envision applications in high-dimensional additive models where each of a large number of functions is expanded in a basis, and the basis coefficients are updated in a block.

Acknowledgments

The authors would like to thank the Editor, an Associate Editor, and three reviewers for their suggestions to improve the presentation and content of the paper. Dr. Bhattacharya acknowledges support from the Office of Naval Research (ONR BAA 14-0001) and the National Science Foundation (NSF-DMS PD 08-1269). Dr. Mallick's research was supported by national cancer Institute of the National Institutes of Health under award number R01CA194391.

References

- Albert JH, Chib S. Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*. 1993; 88:669–679.
- Armagan A, Dunson D, Lee J. Generalized double Pareto shrinkage. *Statistica Sinica*. 2013; 23:119. [PubMed: 24478567]
- Bhattacharya A, Dunson D. Sparse Bayesian infinite factor models. *Biometrika*. 2011; 98:291–306. [PubMed: 23049129]
- Bhattacharya A, Pati D, Pillai NS, Dunson DB. Dirichlet–Laplace priors for optimal shrinkage. *Journal of the American Statistical Association*. 2015; 110:1479–1490. [PubMed: 27019543]
- Caron, F., Doucet, A. Sparse Bayesian nonparametric regression. *Proceedings of the 25th International Conference on Machine learning*; ACM; 2008.
- Carvalho C, Polson N, Scott J. The horseshoe estimator for sparse signals. *Biometrika*. 2010; 97:465–480.
- Castillo I, van der Vaart A. Needles and straw in a haystack: Posterior concentration for possibly sparse sequences. *The Annals of Statistics*. 2012; 40:2069–2101.
- Durante D, Scarpa B, Dunson DB. Locally adaptive factor processes for multivariate time series. *The Journal of Machine Learning Research*. 2014; 15:1493–1522.
- Fan J, Li R. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American statistical Association*. 2001; 96:1348–1360.
- George EI, McCulloch RE. Approaches for Bayesian variable selection. *Statistica sinica*. 1997; 7:339–373.
- Golub, GH., van Loan, CF. *Matrix Computations*. 3. John Hopkins University Press; 1996.
- Griffin J, Brown P. Inference with normal-gamma prior distributions in regression problems. *Bayesian Analysis*. 2010; 5:171–188.
- Hager WW. Updating the inverse of a matrix. *SIAM Review*. 1989; 31:221–239.
- Hahn PR, Carvalho CM. Decoupling shrinkage and selection in Bayesian linear models: a posterior summary perspective. *Journal of the American Statistical Association*. 2015; 110:435–448.
- Holmes CC, Held L. Bayesian auxiliary variable models for binary and multinomial regression. *Bayesian Analysis*. 2006; 1:145–168.
- Javanmard A, Montanari A. Confidence intervals and hypothesis testing for high-dimensional regression. *The Journal of Machine Learning Research*. 2014; 15:2869–2909.
- Khare K, Hobert JP. Geometric ergodicity of the Bayesian lasso. *Electronic Journal of Statistics*. 2013; 7:2150–2163.
- Pal S, Khare K. Geometric ergodicity for Bayesian shrinkage models. *Electronic Journal of Statistics*. 2014; 8:604–645.

- Pati D, Bhattacharya A, Pillai N, Dunson D. Posterior contraction in sparse Bayesian factor models for massive covariance matrices. *The Annals of Statistics*. 2014; 42:1102–1130.
- Polson NG, Scott JG. Shrink globally, act locally: sparse Bayesian regularization and prediction. *Bayesian Statistics*. 2010; 9:501–538.
- Polson NG, Scott JG, Windle J. The Bayesian bridge. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2014; 76:713–733.
- Rue H. Fast sampling of Gaussian Markov random fields. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2001; 63:325–338.
- Scott JG, Berger JO. Bayes and empirical-bayes multiplicity adjustment in the variable-selection problem. *The Annals of Statistics*. 2010; 38:2587–2619.
- Szabó B, van der Vaart A, van Zanten J. Frequentist coverage of adaptive nonparametric Bayesian credible sets. *The Annals of Statistics*. 2015; 43:1391–1428.
- van de Geer S, Bühlmann P, Ritov Y, Dezeure R. On asymptotically optimal confidence regions and tests for high-dimensional models. *The Annals of Statistics*. 2014; 42:1166–1202.
- van der Pas S, Kleijn B, van der Vaart A. The horseshoe estimator: Posterior concentration around nearly black vectors. *Electronic Journal of Statistics*. 2014; 8:2585–2618.
- Zhang CH. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*. 2010; 38:894–942.
- Zhang CH, Zhang SS. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*. 2014; 76:217–242.

Appendix

We provide a constructive proof of Proposition 1 as suggested by a reviewer. Without loss of generality let $D = I_p$ since if $z \sim N(\Sigma\Phi^T\alpha, \Sigma)$ then $D^{1/2}z \sim N(\tilde{\mu}, \tilde{\Sigma})$ where $\tilde{\Sigma} = (\tilde{\Phi}^T\tilde{\Phi} + I_p)^{-1}$ and $\tilde{\mu} = \tilde{\Sigma}\tilde{\Phi}^T\alpha$ with $\tilde{\Phi} = \Phi D^{-1/2}$. Now using the Sherman–Morrison–Woodbury matrix identity we can write

$$(\Phi^T\Phi + I_p)^{-1} = I_p - \Phi^T(\Phi\Phi^T + I_n)^{-1}\Phi = \{I_p - \Phi^T(\Phi\Phi^T + I_n)^{-1}\Phi\}^2 + \Phi^T(\Phi\Phi^T + I_n)^{-2}\Phi.$$

Then, if $u \sim N(0, I_p)$ and $\delta \sim N(0, I_n)$ are independent and we set

$$\theta = \{I_p - \Phi^T(\Phi\Phi^T + I_n)^{-1}\Phi\}u + \Phi^T(\Phi\Phi^T + I_n)^{-1}\delta + \Phi^T\alpha - \Phi^T(\Phi\Phi^T + I_n)^{-1}\Phi\Phi^T\alpha,$$

then θ has the desired distribution (6). Algorithm 1 follows from some straightforward algebra.

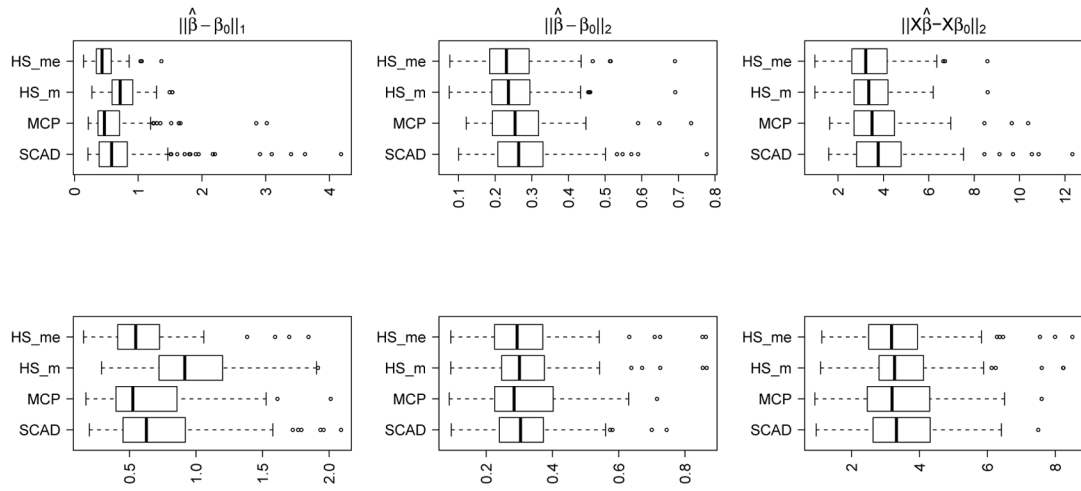


Fig. 1. Boxplots of l_1 , l_2 and prediction error across 100 simulation replicates. HSme and HSm respectively denote posterior point wise median and mean for the horeshoe prior. True β_0 is 5-sparse with nonzero entries $\pm\{1.5, 1.75, 2, 2.25, 2.5\}$. Top row: $\Sigma = I_p$ (independent). Bottom row: $\Sigma_{jj} = 1, \Sigma_{jj'} = 0.5, j \neq j'$ (compound symmetry).

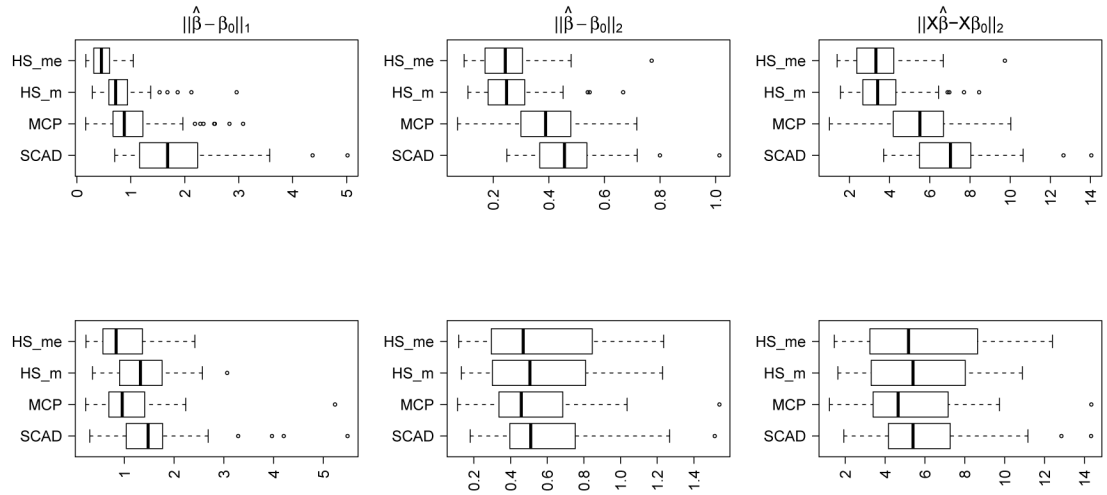


Fig. 2. Same setting as in Fig 1. True β_0 is 5-sparse with non-zero entries $\pm\{0.75, 1, 1.25, 1.5, 1.75\}$.

Table 1

Frequentist coverages (%) and 100×lengths of pointwise 95% intervals. Average coverages and lengths are reported after averaging across all signal variables (rows 1 and 2) and noise variables (rows 3 and 4). Subscripts denote standard errors (%) for coverages. LASSO and SS respectively stand for the methods in van de Geer et al. (2014) and Javanmard & Montanari (2014). The intervals for the horseshoe (HS) are the symmetric posterior credible intervals.

| Design | 500 | | | | | | 1000 | | | | | | | | | | | |
|-----------------|--------------------|--------------------|-------------------|--------------------|-------------------|-------------------|-------------------|-------------------|-------------------|-------------------|--------------------|-------------------|--------------------|-------------------|-------------------|--------------------|-------------------|-------------------|
| | Independent | | Comp Symm | | Toeplitz | | Independent | | Comp Symm | | Toeplitz | | | | | | | |
| | HS | LASSO | SS | HS | LASSO | SS | HS | LASSO | SS | HS | LASSO | SS | | | | | | |
| Signal Coverage | 93 _{1,0} | 75 _{12,0} | 82 _{5,7} | 95 _{0,9} | 73 _{4,0} | 80 _{4,0} | 94 _{4,0} | 80 _{7,0} | 79 _{5,6} | 94 _{2,0} | 78 _{12,0} | 85 _{5,1} | 94 _{1,0} | 77 _{2,0} | 82 _{7,4} | 95 _{1,0} | 76 _{3,0} | 80 _{8,3} |
| Signal Length | 42 | 46 | 41 | 85 | 71 | 75 | 86 | 79 | 74 | 39 | 41 | 42 | 82 | 76 | 77 | 105 | 96 | 95 |
| Noise Coverage | 100 _{0,0} | 99 _{0,8} | 99 _{1,0} | 100 _{0,0} | 98 _{1,0} | 99 _{0,8} | 98 ₁ | 98 _{1,0} | 99 _{0,6} | 99 _{1,0} | 99 _{1,0} | 98 _{0,9} | 100 _{0,0} | 99 _{1,0} | 99 _{0,1} | 100 _{0,0} | 99 _{1,0} | 99 _{0,2} |
| Noise Length | 2 | 43 | 40 | 4 | 69 | 73 | 5 | 78 | 73 | 0.6 | 0.42 | 41 | 0.7 | 76 | 77 | 0.3 | 98 | 93 |