HDS 8. Bayesian variable selection

Matti Pirinen, University of Helsinki

11.1.2024

In HDS6, we saw that LASSO has a feature of setting many of the coefficients to exactly zero at the optimum of the objective function and hence we said that LASSO does *variable selection*. A technical reason for this behavior was the LASSO's prior distribution for coefficients that had a sharp peak around zero. In HDS7, we embedded the LASSO model into a fully Bayesian model **blasso** that allowed us to estimate the full posterior distribution of the coefficients but that was computationally much more heavy to apply than the standard LASSO. Next, we will consider conceptually simpler Bayesian models for variable selection, where the prior distribution of a coefficient states that the coefficient is non-zero with probability π_1 and has a positive probability mass $1 - \pi_1$ of being exactly zero.

The spike and slab prior (SSP)

Under SSP, we model an individual regression coefficient β_i using a two-component **mixture distribution**

$$\beta_j \mid \pi_1, \tau^2 \sim (1 - \pi_1)\delta_0 + \pi_1 \mathcal{N}(0, \tau^2)$$

where δ_0 is the point mass at 0 (in mathematics called as *Dirac delta function*). Another way to describe this two-component mixture distribution is to say that

$$\begin{cases} \beta_j = 0, & \text{with probability } 1 - \pi_1, \\ \beta_j \sim \mathcal{N}(0, \tau^2), & \text{with probability } \pi_1. \end{cases}$$

Thus, if we wanted to generate example data from such a mixture distribution, we could proceed in two steps. First, we chose which component to sample from (component 1 with probability π_1 and component 0 with probability $1 - \pi_1$). Second, conditional on the chosen component, we would sample a value for β_j from the distribution specific to the chosen component.

Spike and slab with $\pi_1=0.5 \tau=1$



Figure visualizes the two-component spike and slab distribution. The spike component sits at exactly 0, and has a non-zero probability mass. The slab component is a Gaussian distribution centered at 0 and having SD > 0. The figure describes the probability mass (black) for the spike and the density function (blue) for the slab.

Suppose we have *n* samples and *p* predictors (columns of X) to model the outcome vector y. Let $\gamma = (\gamma_j)_{j=1}^p$ be a vector of binary indicators indicating which variables are non-zero, that is,

$$\gamma_j = \begin{cases} 1, & \text{if } \beta_j \neq 0, \\ 0, & \text{if } \beta_j = 0. \end{cases}$$

We can assign to each predictor, independently, an $SSP(\pi_1, \tau^2)$ prior where the values π_1 and τ^2 are shared between the predictors. Value of π_1 could be determined based on which proportion of predictors we expect to be non-zero, and τ^2 could be determined based on what magnitude of coefficient values we expect at the non-zero coefficients. Such a model is implemented in BoomSpikeSlab package. (Later we will discuss extensions that also estimate π_1 and τ^2 from data.)

Let's try it on the same data we analyzed using Bayesian LASSO in HDS7 that had n = 250, p = 30 and the first 3 variables had an effect size of 0.229 while the remaining 27 were exactly 0. We use as prior parameters $\pi_1 = 5/30 \approx 0.167$ (by setting expected.model.size = 5) and $\tau^2 = 1$ (by setting prior.beta.sd = rep(1,p)).

```
set.seed(122)
p = 30
n = 250
phi = 0.05 # variance explained by x_1, should be 0 < phi < 1.
b = rep(c(sqrt(phi / (1-phi)), 0), c(3, p-3)) # effects 1,2,3 are non-zero, see Lect. 0.1 for "phi"
X = scale(matrix(rnorm(n*p), nrow = n)) # cols 1,2,3 of X have effects, other cols are noise
eps = scale(rnorm(n, 0, 1)) # epsilon, error term
y = scale(X%*%b + eps, scale = FALSE) # makes y have mean = 0
library(BoomSpikeSlab)</pre>
```

summary(lm.ss)

```
##
  coefficients:
##
            mean
                       sd mean.inc sd.inc inc.prob
## X3
        3.25e-01 0.06140
                          0.32500 0.0614
                                              1.000
## X1
        2.65e-01 0.07120
                          0.26900 0.0642
                                              0.986
        2.39e-01 0.08060
                           0.25000 0.0640
                                              0.957
## X2
## X6
        9.79e-03 0.03860
                           0.12700 0.0667
                                              0.077
## X11
                                              0.045
        4.50e-03 0.02410
                          0.10000 0.0585
## X23 -4.01e-03 0.02320 -0.09550 0.0645
                                              0.042
## X27 -4.48e-03 0.02560 -0.11500 0.0650
                                              0.039
       -2.02e-03 0.01570 -0.06310 0.0631
                                              0.032
## X9
## X15 -1.91e-03 0.01500 -0.07080 0.0596
                                              0.027
        2.97e-04 0.00816
                                              0.020
## X22
                          0.01480 0.0572
        1.14e-03 0.01220
## X20
                          0.06020 0.0671
                                              0.019
## X10 -1.06e-03 0.01210 -0.05600 0.0695
                                              0.019
## X17 -6.31e-04 0.01000 -0.03510 0.0681
                                              0.018
## X29
        4.77e-04 0.00818
                          0.02800 0.0580
                                              0.017
## X24 -1.24e-03 0.01120 -0.07310 0.0477
                                              0.017
## X28
       1.51e-04 0.00559
                          0.00942 0.0446
                                              0.016
## X12 -1.06e-03 0.01230 -0.06650 0.0734
                                              0.016
## X19 -2.48e-04 0.00690 -0.01660 0.0557
                                              0.015
## X16
       1.87e-04 0.00974
                          0.01330 0.0842
                                              0.014
## X14 -3.41e-04 0.00873 -0.02440 0.0723
                                              0.014
## X5
       -1.90e-04 0.00786 -0.01360 0.0674
                                              0.014
## X13 -7.43e-05 0.00630 -0.00571 0.0572
                                              0.013
## X7
       -1.00e-03 0.00990 -0.07720 0.0424
                                              0.013
## X26
       6.95e-05 0.00637
                          0.00579 0.0604
                                              0.012
## X18
       4.45e-04 0.00681
                          0.03700 0.0523
                                              0.012
## X4
       -2.93e-04 0.01060 -0.02440 0.0981
                                              0.012
        6.06e-04 0.00900
                          0.05510 0.0692
                                              0.011
## X25
##
  X21 -1.26e-05 0.00669 -0.00126 0.0705
                                              0.010
       -3.50e-04 0.00573 -0.03500 0.0479
##
  Х8
                                              0.010
  X30 -3.42e-04 0.00631 -0.04270 0.0601
##
                                              0.008
##
##
   residual.sd =
##
      Min. 1st Qu.
                    Median
                               Mean 3rd Qu.
                                                Max.
            0.9638
                    0.9917
                             0.9946
##
    0.8597
                                    1.0245
                                              1.1778
##
##
  r-square
               =
##
                                                Max.
      Min. 1st Qu.
                    Median
                               Mean 3rd Qu.
## -0.1208
            0.1519
                    0.2054
                             0.1990
                                     0.2494
                                             0.4028
```

We see that the 3 variables with true non-zero effects have posterior inclusion probabilities (PIP) > 0.90 whereas all others have PIPs < 0.10. In the results, we have two sets of posterior means and standard deviations: first from the full posterior distribution and second conditional on that the variable has been included in the model. When the variable has a large PIP, then the two are similar, but for small PIPs, the two differ and the full unconditional mean is a shrunk version of the conditional mean where the shrinkage

factor is the PIP. For example, for variable X6, the posterior mean is 0.1270995 conditional on the variable being non-zero, and X6 is non-zero with a posterior probability of 0.077 Thus, the (unconditional) posterior mean is

 $0.077 \cdot 0.1270995 + 0.923 \cdot 0 = 0.0097867.$

Note that, in a similar way, one could extend the LASSO method to **relaxed LASSO**, where one first runs LASSO to choose the non-zero variables (corresponding to the unconditional analysis) and then fits an unpenalized model using only the chosen variables (corresponding to the conditional analysis). Statistical inference for such a stepwise procedure is, however, complicated, whereas, in the Bayesian approach, the posterior distribution is directly available and forms a natural basis for inference.

Let's visualize the PIPs and the posterior distributions of the coefficients.

```
par(mfrow = c(1, 2))
plot(lm.ss) # plots PIPs
plot(lm.ss, "coefficients") # plots coeffs
```



The posterior probability on the model size can also be plotted.

plot(lm.ss, "size", freq = F)



Histogram of size



Extension to estimating π_1 and τ^2

To make the model fully Bayesian, we would like to estimate also the parameters π_1 and τ^2 from the data. For that, we will consider τ^2 as being specified relative to the error variance σ^2 , i.e., we use a prior distribution

$$\beta_j \mid \pi_1, \tau^2, \sigma^2 \sim (1 - \pi_1)\delta_0 + \pi_1 \mathcal{N}(0, \sigma^2 \tau^2).$$

For example, the parameters π_1 and τ^2 can be given priors $\pi_1 \sim \text{Beta}(b_0, b_1)$ and $\tau^2 \sim \text{Inverse-Gamma}(1/2, s^2/2)$ where the hyper-parameters b_0, b_1 and s are given such fixed values that the priors cover all reasonable values of the parameters in a fairly uniform manner. Additionally, a flat prior on σ^2 could be an Inverse-Gamma(a, a) where a is a small positive value, such as 0.01, as such a distribution approximates a uniform prior on $\log(\sigma^2)$ scale (so called *Jeffrey's prior* for the variance parameter). By using such "flat priors" we hope to achieve a posterior distribution that has been guided by information from the data while the effect of the prior distribution on the posterior remains small.

In the following analysis, we will set the hyper-parameters to $b_0 = b_1 = 1$, s = 0.5 and a = 0.01 that lead to the following prior densities on π_1 and τ^2 .



Inv-Gamma(1/2,1/8)



The following Gibbs sampler code to implement this model is adapted from Fabian Dablander https://fabiandablander.com/r/Spike-and-Slab.html

```
#' Spike-and-Slab Regression using Gibbs Sampling for p > 1 predictors
#'
#' Oparam y: vector of responses
#' Oparam X: matrix of predictor values
#' Oparam nr_samples: indicates number of samples drawn
#' Oparam a1: parameter a1 of Gamma prior on variance sigma2e
#' Oparam a2: parameter a2 of Gamma prior on variance sigma2e
#' Oparam pil: parameter of prior over mixture weight
#' Oparam a: 1st parameter of Beta prior for pi1
#' Oparam b: 2nd parameter of Beta prior for pi1
#' Oparam burnin: number of samples we discard ('burnin samples')
#'
#' Creturns matrix of posterior samples from parameters gamma, beta, tau2, sigma2e, pi1
ss_regress <- function(</pre>
 y, X, a1 = .01, a2 = .01, pi1 = .5,
 a = 1, b = 1, s = 1/2, nr_{samples} = 6000, nr_{burnin} = round(nr_{samples} / 4, 2)
 ) {
 p <- ncol(X)</pre>
  n \leftarrow nrow(X)
  # res is where we store the posterior samples
  res <- matrix(NA, nrow = nr_samples, ncol = 2*p + 1 + 1 + 1)</pre>
  colnames(res) <- c(</pre>
    paste0('gamma', seq(p)),
    paste0('beta', seq(p)),
    'sigma2', 'tau2', 'pi1'
  )
  # take the MLE estimate as the values for the first sample
  m < -lm(y ~ X - 1)
  res[1, ] <- c(rep(0, p), coef(m), var(predict(m) - y), 1, .5)</pre>
```

```
# compute only once
XtX <- t(X) %*% X
Xty <- t(X) %*% y
# we start running the Gibbs sampler
for (i in seq(2, nr_samples)) {
  # first, get all the values of the previous time point
  gamma_prev <- res[i-1, seq(p)]</pre>
 beta_prev <- res[i-1, seq(p + 1, 2*p)]</pre>
 sigma2_prev <- res[i-1, ncol(res) - 2]</pre>
 tau2_prev <- res[i-1, ncol(res) - 1]</pre>
 pi1_prev <- res[i-1, ncol(res)]</pre>
  ## Start sampling from the conditional posterior distributions
  # sample pi1 from a Beta
 pi1_new <- rbeta(1, a + sum(gamma_prev), b + sum(1 - gamma_prev))</pre>
  # sample sigma2e from an Inverse-Gamma
 err <- y - X %*% beta_prev
  sigma2_new <- 1 / rgamma(1, a1 + n/2, a2 + t(err) %*% err / 2)</pre>
  # sample tau2 from an Inverse Gamma
 tau2_new <- 1 / rgamma(</pre>
   1, 1/2 + 1/2 * sum(gamma_prev),
   s^2/2 + t(beta_prev) %*% beta_prev / (2*sigma2_new)
 )
  # sample beta from multivariate Gaussian
 beta_cov <- qr.solve((1/sigma2_new) * XtX + diag(1/(tau2_new*sigma2_new), p))</pre>
 beta_mean <- beta_cov %*% Xty * (1/sigma2_new)</pre>
 beta_new <- mvtnorm::rmvnorm(1, beta_mean, beta_cov)</pre>
  # sample each gamma_j in random order
 for (j in sample(seq(p))) {
    # get the betas for which beta_j is zero
    gamma0 <- gamma_prev
    gamma0[j] <- 0
    bp0 <- t(beta_new * gamma0)</pre>
    # compute the z variables and the conditional variance
   xj <- X[, j]
    z <- y - X %*% bp0
    cond_var <- sum(xj^2) + 1/tau2_new</pre>
    # compute chance parameter of the conditional posterior of qamma_j (Bernoulli)
   10 < -\log(1 - pi1_new)
   11 <- (
      log(pi1_new) - .5 * log(tau2_new*sigma2_new) +
      sum(xj*z)^2 / (2*sigma2_new*cond_var) + .5 * log(sigma2_new / cond_var)
```

```
/
    # sample gamma_j from a Bernoulli
    gamma_prev[j] <- rbinom(1, 1, exp(l1) / (exp(l1) + exp(l0)))
}
gamma_new <- gamma_prev
# add new samples
res[i, ] <- c(gamma_new, beta_new*gamma_new, sigma2_new, tau2_new, pi1_new)
}
# remove the first nr_burnin number of samples
res[-seq(nr_burnin), ]
}</pre>
```

Let's fit the model using this Gibbs sampler and visualize the posteriors of π_1 and τ^2 .

```
set.seed(18)
ss = ss_regress(
    y, X, a1 = .01, a2 = .01, pi1 = .5,
    a = 1, b = 1, s = 1/2, nr_samples = 5000
    )
par(mfrow = c(1,2))
plot(density(ss[,"pi1"]), xlab = expression(pi[1]), lwd = 2,
    main = "Posterior of pi1", col = "red")
plot(density(ss[,"tau2"]), xlab = expression(tau^2), lwd = 2,
    main = "Posterior of tau2", col = "orange", xlim = c(0,1))
```







summary(ss[,"pi1"])

Min. 1st Qu. Median Mean 3rd Qu. Max. ## 0.008434 0.126940 0.184983 0.199423 0.258069 0.688194 summary(ss[,"tau2"])

Min. 1st Qu. Median Mean 3rd Qu. Max. ## 0.01785 0.05738 0.09008 0.14707 0.15038 14.88438

The posterior of π_1 is quite wide and suggest that π_1 is somewhere between 0.05 and 0.4. For τ^2 , the bulk of the posterior is below 0.20.

We can also see the PIPs and the posterior distributions of the coefficients.



Unfortunately, these implementations of Bayesian variable selection using the Spike and Slab prior through Gibbs sampler do not scale to large data sets. The first problem is that, at each iteration of the Gibbs sampler, the matrix decomposition on $p \times p$ matrix related to $\mathbf{X}^T \mathbf{X}$ is done. The second problem is that, at

each iteration, every predictor is updated even when the model might be very sparse, i.e., even when only a few coefficients are non-zero. These issues become too demanding when p becomes truly large.

Let's next discuss some recent work that has made it possible to apply these kinds of models to very large p.

Efficient computation of Bayes factors under SSP

Let's continue with the assumption that each predictor has an independent $SSP(\pi_1, \tau^2)$ prior. It has been shown by Benner et al. that the Bayes factor (BF) comparing a model where the non-zero coefficients are specified by a configuration γ and the null model $\gamma_0 = (0, \ldots, 0)$ where all coefficients are zero can be computed using data on only the non-zero variables (Details in Slides). This allows computing BF($\gamma : \gamma_0$) between γ and γ_0 in $\mathcal{O}(k^3)$ operations, where $k = \sum_{j=1}^p \gamma_j$ is the number of non-zero coefficients in γ , compared to $\mathcal{O}(p^3)$ operations that are required when the computation is done in the standard way. When we consider sparse models where $k \sim 10$ and $p \sim 10^5$ this can make all the difference.

Note that under the $SSP(\pi_1, \tau^2)$ prior, the prior probability of a configuration γ is $\pi_1^{k_{\gamma}}(1-\pi_1)^{p-k_{\gamma}}$ and only depends on k_{γ} . Thus, by knowing the BFs for all possible configurations, we can compute an (unnormalized) posterior probabilities for the configurations as

$$\Pr(\boldsymbol{\gamma} \mid \text{Data}) \propto \text{BF}(\boldsymbol{\gamma} : \boldsymbol{\gamma}_0) \pi_1^{k_{\gamma}} (1 - \pi_1)^{p - k_{\gamma}}.$$

By normalizing these quantities with respect to their sum over γ , we have computed the posteriors for every configuration.

However, if we allowed, say, 10 non-zero variables out of 10^5 candidates, we would still have $\binom{10^5}{10} \approx 10^{43}$ configurations to evaluate, and that is far too many even when we can do every one of them quickly. For this reason, stochastic search algorithms have been introduced to approximate the posterior by finding the most relevant configurations from the whole space of configurations and then normalizing their un-normalized posteriors with respect to the relevant subset of configurations. (See slides for an example.)

Sum of single effect model (SUSIE)

Another recent approach to analyze Bayesian variable selection model is using the 'single-effect regression' (SER) model as a building block (Wang et al.). SER model is a multiple-regression model in which exactly one of the p predictors has a non-zero regression coefficient. The idea is to combine some number K of these SERs together to get a model where there are K non-zero predictors, or K effects. The key is that fitting each SER model conditional on the other SER models is very quick.

SUSIE model with K effects Let's assume that the number of non-zero effects K, the error variance σ^2 and the prior variance on the non-zero effect τ^2 are fixed. For each non-zero effect k = 1, ..., K, define $\gamma^{(k)} \in \{0, 1\}^p$ as the indicator vector that shows which is the kth non-zero coefficient, that is, $\sum_{j=1}^p \gamma_j^{(k)} = 1$. The model is

$$\begin{aligned} \boldsymbol{y} &\sim \mathcal{N}(\boldsymbol{X}\boldsymbol{\beta},\sigma^{2}\boldsymbol{I}) \\ \boldsymbol{\beta} &= \sum_{k=1}^{K} \boldsymbol{\beta}^{(k)} \\ \boldsymbol{\beta}^{(k)} &= \boldsymbol{\beta}^{(k)}\boldsymbol{\gamma}^{(k)}, \text{ for each } k = 1,\ldots,K, \\ \boldsymbol{\beta}^{(k)} &\sim \mathcal{N}(0,\tau^{2}), \text{ for each } k = 1,\ldots,K, \\ \boldsymbol{\gamma}^{(k)} &\sim \text{Multinomial}(1,(1/p,\ldots,1/p)), \text{ for each } k = 1,\ldots,K. \end{aligned}$$

The last line means that exactly one of the elements of $\gamma^{(k)}$ is non-zero and, *a priori*, it is equally likely to be any one of the *p* elements.

Note that the model does not rule out the possibility that $\boldsymbol{\gamma}^{(k)} = \boldsymbol{\gamma}^{(\ell)}$ for some $k \neq \ell$ even though conceptually it would be natural to avoid such overlaps. However, by leaving a possibility for an overlap, we can treat each effect independently and we get much simpler computations. Additionally, in practice, we won't observe overlapping effects in the posterior distribution estimated by the following algorithm.

Iterative Bayesian stepwise selection (IBSS) A key motivation for the SUSIE model is that, given $(\boldsymbol{\beta}^{(\ell)})_{\ell \neq k}$, fitting $\boldsymbol{\beta}^{(k)}$ involves only fitting an SER model whose posterior distribution on the indicator $\boldsymbol{\gamma}^{(k)}$ and the parameter $\beta^{(k)}$ is simple to compute. This suggests an iterative algorithm (called IBSS) to fit the model. Below we apply it to our existing data set by allowing K = 5 effects and setting the prior variance $\tau^2 = 1$.

```
p = ncol(X)
n = nrow(X)
K = 5
tau2 = 1
xx = colSums(X^2) \# = diagonal of t(X) \% X. Will be needed later.
b = matrix(0, nrow = K, ncol = p) # initialize each effect beta_k = 0
res.b = res.pip = res.v = b # initialize results to 0
diff.in.post = rep(0,p) # initialize to 0
converged = FALSE
tol = 1e-6 # stop when difference between iterations < tol
iter = 0
while(!converged){
  iter = iter + 1
  for (k in 1:K){
                     # Fitting SER model for effect k
    r = as.vector(y - rowSums(X%*%t(b[-k,]))) # residuals without effect k
    # Use Exercise 1.5 results here to get MLE of linear model
    b.k = t(X) %*% r / xx # univariate OLS estimate
    B = matrix(b.k, byrow = T, nrow = n, ncol = p)
    M = (X*B - r)^{2}
    v.k = colSums(M) / xx / (n-1) \# MLE for variance of b.k
    # Combine MLEs with prior distributions to get posteriors
    post.v.k = 1/(1/v.k + 1/tau2)
    post.b.k = post.v.k/v.k * b.k
    # Conditional posterior for beta_k is N(post.b.k, post.v.k)
    # conditional on the effect being at that position
    # PIPs for the effect being at position 1,...,p:
    \log.bf = 0.5 * \log(v.k/(v.k + tau2)) + 0.5 * b.k^2/v.k * tau2/(tau2 + v.k)
    tmp = exp(log.bf - max(log.bf))
    post.ip = tmp / sum(tmp)
    b[k,] = post.b.k*post.ip # save new effect k estimate for next iteration
    # compare to previous iteration to observe convergence:
    diff.in.post[k] = max(abs(b[k,] - res.b[k,]*res.pip[k,]))
    # save results
    res.b[k,] = post.b.k
    res.pip[k,] = post.ip
    res.v[k,] = post.v.k
  }
```

```
converged = (all(diff.in.post < tol))
}
print(paste("Ended in", iter, "iterations"))</pre>
```

```
## [1] "Ended in 8 iterations"
```

Let's examine the output of this algorithm. We allowed 5 non-zero effects. Let's see the distributions of PIPs for each effect k = 1, ..., 5.



We see that for the first 3 effects, there is a clear difference between predictors in their PIPs. For the last two, none of the coefficients has a large PIP and therefore we may want to ignore those effects. In other words, likely there are only 3 clear non-zero effects there and the remaining effects, if included in the SUSIE model, wouldn't provide any information about which of the predictors would have a non-zero effect. Note, however, that despite this reasonable ad hoc criterion to determine the number of effects, the SUSIE model cannot make any probabilistic assessment of the number of non-zero effects whereas the SSP model directly gave us the posterior distribution on the number of non-zero coefficients.

We can next estimate the coefficients by summing over the K SER models. Let's denote the posterior mean and variance of the SER model k conditional on the effect being at predictor j by $\mu_j^{(k)}$ and $\eta_j^{(k)}$, respectively. Then we can compute the mean and variance for each coefficient (marginalized over all K SER models) as follows.

$$E(\beta_{j}) = E\left(\sum_{k=1}^{K} \gamma_{j}^{(k)} \beta_{j}^{(k)}\right) = \sum_{k=1}^{K} \pi_{j}^{(k)} \mu_{j}^{(k)}$$

$$Var(\beta_{j}) = Var\left(\sum_{k=1}^{K} \gamma_{j}^{(k)} \beta_{j}^{(k)}\right) = \sum_{k=1}^{K} Var\left(\gamma_{j}^{(k)} \beta_{j}^{(k)}\right)$$

$$= \sum_{k=1}^{K} \left(E\left(\gamma_{j}^{(k)^{2}} \beta_{j}^{(k)^{2}}\right) - E\left(\gamma_{j}^{(k)} \beta_{j}^{(k)}\right)^{2}\right)$$

$$= \sum_{k=1}^{K} \left(\pi_{j}^{(k)} \left(\eta_{j}^{(k)^{2}} + \mu_{j}^{(k)^{2}}\right) - \pi_{j}^{(k)^{2}} \mu_{j}^{(k)^{2}}\right)$$

$$= \sum_{k=1}^{K} \left(\pi_{j}^{(k)} \left(\eta_{j}^{(k)^{2}}\right) + \pi_{j}^{(k)} (1 - \pi_{j}^{(k)}) \mu_{j}^{(k)^{2}}\right)$$

Let's find the posterior means and variances and then show the 95% credible intervals assuming that the shape of the posterior distributions are Gaussian.

```
b.mean = colSums(res.pip * res.b)
b.var = colSums(res.pip * res.v + res.pip * (1 - res.pip) * res.b^2)
plot(NULL, xlim = c(0, p + 1), ylim = c(-0.1, 0.5),
    ylab = expression(beta), xlab = "")
points(1:p, b.mean, pch = 19)
arrows(1:p, b.mean - 1.96 * sqrt(b.var), 1:p, b.mean + 1.96 * sqrt(b.var),
        code = 3, length = 0)
abline(h = 0)
```



Credible sets of non-zero variables A nice property of the SUSIE model is that it allows a straightforward computation of **credible set** of the non-zero variables.

Definition. A level ρ credible set (CS) is a subset of variables that has a probability $\geq \rho$ of containing at least one effect variable, i.e., a variable with non-zero regression coefficient.

For each effect in the SUSIE model, we can compute a credible set by sorting the PIPs in decreasing order and by forming the CS by including in the set the smallest possible number of variables whose PIPs sum to a value $\geq \rho$.

In our current example, the 95% CSs are below.

```
rho = 0.95
cs = list()
for(ii in 1:K){
  ord = order(res.pip[ii,], decreasing = T)
  cs[[ii]] = ord[1:min(which(cumsum(res.pip[ii,ord]) >= rho))]
}
cs
## [[1]]
## [1] 3
##
## [[2]]
## [1] 1
##
## [[3]]
## [1] 2
##
## [[4]]
         6 23 11 27 15 9 12 20 24 10 30 7 5 18 17 25 29 16 14 26 22 13 19 8 1
##
   [1]
         4 28 3
##
   [26]
##
## [[5]]
         6 23 11 27 15 9 12 20 24 10 30 7 5 18 17 25 29 16 14 26 22 13 19 8 1
##
   [1]
## [26]
         4 28 3
```

For the first three effects, the CSs only contain one variable each, while the CSs for the last two effects contain almost all variables. This again shows that the first three effects are clearly localized each to one variable whereas the two additional effects are not informative and we might want to just ignore them.

Note that the posterior distribution provided by the SSP model does not directly contain information by which we could build credible sets for the variables. Credible sets are useful when one needs to list the possible candidate variables that could replace each other in the model.

Question. Assume that there are three highly correlated variables of which one truly affects the outcome but no other variables associated with the outcome. What kind of PIPs SUSIE would give and how would the credible sets look like?