

R functions for the Bayesian Lasso Variable Selection (BLVS)

Anastasia Lykou & Ioannis Ntzoufras

The function `blvs`

The function `blvs` performs the Bayesian version of the Lasso by imposing the Double-Exponential prior distribution on the linear regression coefficients.

Usage:

```
blvs(X, Y, lambda=NULL, r=NULL, delta=NULL, alpha2=0.01, gamma2=0.01,  
      nburn=1000, ndraw=5000, stand = TRUE)
```

Input arguments:

`X`: numeric matrix ($n \times p$) containing the \mathbf{X} variables.

`Y`: numeric vector containing the responses.

`lambda`: shrinkage parameter. A default of 'NULL' requires numeric values for the hyperparameters `r` and `delta`.

`r`, `delta`: hyperparameter for the Gamma hyperprior imposed on `lambda`. Shape is defined by `r` and scale by `delta`. If both `r` and `delta` are not available, then a value for `lambda` is required.

`alpha2`, `gamma2`: rate and scale parameters for the Gamma prior imposed on the precision `tau`.

`nburn`: number of updates discarded as the burn-in period.

`ndraw`: number of updates generated after the burn-in period.

`stand`: if true, variables are standardised to have zero mean and unit variance.

Output components:

`draws`: a numeric matrix with $(2 * p + 2)$ columns:

$1, \dots, p$ columns contain the β updates

$p + 1, \dots, 2 * p$ columns contain the indicator parameters updates

$2 * p + 1$ column contains the precision updates

$2 * p + 2$ column contains the shrinkage parameter updates when a hyperprior is imposed.

Required R packages: `msm`

Examples We perform the Bayesian Lasso Variable Selection method on the diabetes data set that can be found in the `lars` R package. All the variables

are standardized to have zero mean, unit variance and the data matrix is transformed appropriately to ensure that all the covariates are positively correlated with the response.

The Bayesian Lasso is performed for a fixed shrinkage value $\lambda = 1.459$ as proposed in Lykou and Ntzoufras (2011)

```
> bLasso = blvs(X, Y, lambda=1.459, alpha2=0.0001, gamma2=10000,
  nburn=1000, ndraw=20000)
```

The posterior updates for the parameters of interest can be derived as follows

```
> beta.post = bLasso[,1:p]
> gamma.post = bLasso[(p+1):(2*p)]
> tau.post = bLasso[,2*p+1]
```

The effect of each covariate X_j is evaluated through the model averaged versions of the product of the Lasso regression parameters β and the indicator γ .

```
> apply(beta.post*gamma.post, 2, median)
[1] 0.000 -0.135 0.329 0.198 -0.073 0.000 0.126 0.000 0.331 0.000
```

The corresponding posterior inclusion probabilities are extracted below.

```
> apply(gamma.post, 2, mean)
[1] 0.107 0.985 1.000 1.000 0.627 0.420 0.700 0.378 1.000 0.183
```

The variables X_2, X_3, X_4 and X_9 seem to be the most important predictors, since their posterior inclusion probabilities are almost equal to 1. The variables X_5 and X_7 have lower posterior probabilities (0.63, 0.70 respectively) indicating that are of moderate significance.

The example is repeated by considering the hierarchical model and imposing a Gamma hyperprior on the shrinkage parameter. The hyperparameters are specified according to Lykou and Ntzoufras (2011) corresponding to prior mean 1.459 and prior variance 0.337.

```
> bLasso = blvs(X, Y, r=6.322, delta=0.231, alpha2=0.0001, gamma2=10000,
  nburn=1000, ndraw=20000)
```

```
> beta.post = bLasso[,1:p]
> gamma.post = bLasso[(p+1):(2*p)]
> tau.post = bLasso[,2*p+1]
> lambda.post = bLasso[,2*p+2]
```

The posterior summaries are found likewise.

```
> apply(beta.post*gamma.post, 2, median)
[1] 0.000 -0.132 0.330 0.197 -0.081 0.000 0.119 0.000 0.329 0.000
> apply(gamma.post, 2, mean)
[1] 0.134 0.985 1.000 1.000 0.648 0.395 0.707 0.448 1.000 0.215
> mean(lambda.post)
[1] 1.830
```

The function `bf.rho`

The function `bf.rho` returns the benchmark correlation for a given sample size and a Bayes factor (BF) value.

Usage:

```
bf.rho(n, crit.value)
```

Input arguments:

`n`: sample size

`crit.value`: the critical value of BF for which the benchmark correlation will be returned.

Examples

```
> bf.rho(50,3)
[1] 0.316
```

The Bayes factors of all the covariates that are correlated with the response with $\rho \leq 0.316$ are upper bounded by the value 3 for all the shrinkage parameters when $n = 50$.

The function `bf.lambda`

For given values of the sample size and the correlation, the function `bf.lambda` returns the value of the shrinkage parameter that makes the Bayes factor to be equal to a specified value.

Usage:

```
bf.lambda(rho, n, bf.value)
```

Input arguments:

`rho`: correlation value

`n`: sample size

`bf.value`: the given value for the BF.

Examples

```
> bf.lambda(rho=0.316, n=50, bf.value=1)
[1] 0.446
```

The value of shrinkage parameter 0.446 corresponds to $BF = 1$ when the candidate covariate is correlated with the response with $\rho = 0.316$ and the sample size is 50.

Bibliography

Lykou, A. and Ntzoufras, I. (2011). On Bayesian lasso variable selection and the specification of the shrinkage parameter. Technical report, Department of Statistics, Athens University of Economics and Business; available at <http://stat-athens.aueb.gr/~jbn/papers/paper25.htm>.