

EBglmnet Vignette

Anhui Huang Dianting Liu

Jan. 15, 2016

Contents

Introduction	2
Lasso and its Bayesian Interpretation	3
Empirical Bayesian Lasso (EBlasso)	4
Hierarchical Prior Distributions in EBglmnet	4
EBglmnet Implementation and Usage	5
Installation	6
Quick Start	6
GLM Family	7
Gaussian Model	7
Binomial Model	7
Prior and Hyperparameters	8
Example of $p > n$ Data	8
Gaussian Model with Main Effects	8
Binomial Model and Separation Problem	9
References	13
EBglmnet Algorithms:	13
Lasso and glmnet package	13

Introduction

Installation

Quick Start

GLM Family

Prior, Hyperparameters and Epistasis

Example of $p > n$ Data

Introduction

Acronyms and Concepts

EBglmnet is a package that implements the empirical Bayesian Lasso (EBlasso) and Elastic Net (EBEN) method for generalized linear models (GLMs). Additionally, in EBlasso, two different prior distributions are also developed: one with two-level hierarchical Normal + Exponential prior (denoted as NE), and the other one with three-level Normal + Exponential + Gamma prior (denoted as NEG). The major difference of the prior distributions are the probability density around zero and two tails:

- Probability density around zero: higher probability density around zero will enforce stronger shrinkage effect, resulting in more sparse model (i.e., the model set the coefficients of predictors with small or noise prediction power to zero);
- Probability density on tails: higher probability density around tails leave enough room for strong predictors to stay with non-zero coefficients.

A good prior distribution is a balance of the aforementioned properties as shown in the Figure 1. Generally, NE prior leads to more non-zero predictors with smaller absolute coefficient values, while NEG prior leads to less non-zero predictors with stronger signals in terms of both absolute coefficient values as well as significance level.

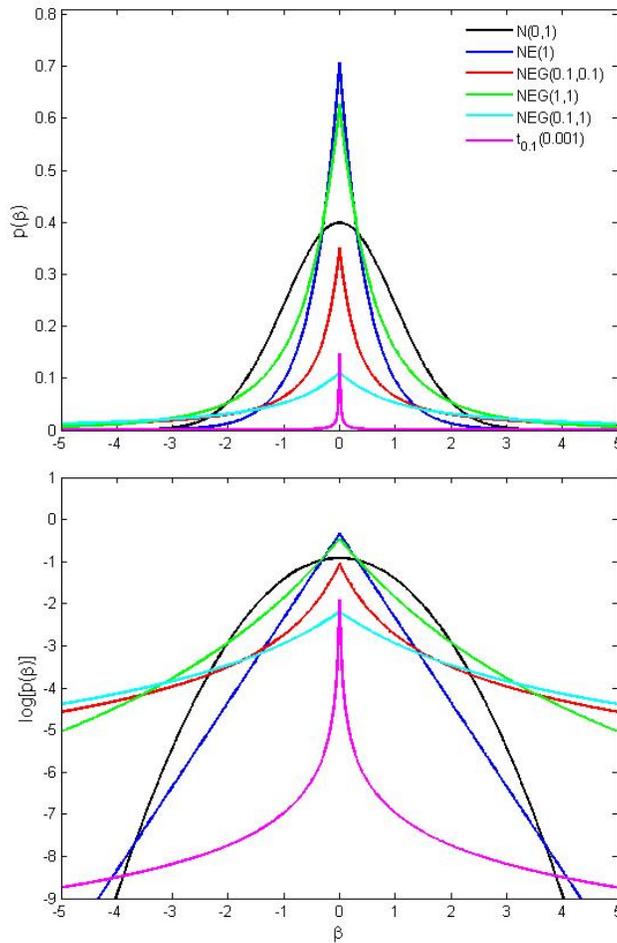


Figure 1: Prior distributions that penalise posterior distributions

The following names should not be confused with the `lasso` and `elastic net` method from the comparison package `glmnet`:

EBglmnet: package that implements **EBlasso** and **EBEN** methods.

EBlasso: empirical Bayesian method with `lasso` prior distribution, which includes two sets of prior distributions: **NE** and **NEG**.

EBEN: empirical Bayesian method with `elastic net` prior distribution.

lasso prior: the hierarchical prior distribution that is equivalent with `lasso` penalty term when the marginal probability distribution for the regression coefficients is considered.

elastic net prior: the hierarchical prior distribution that is equivalent with `elastic net` penalty term when the marginal probability distribution for the regression coefficients is considered.

EBlasso-NE: **EBlasso** method with **NE** prior.

EBlasso-NEG: **EBlasso** method with **NEG** prior.

Generalized Linear Models (GLMs) In a GLM

$$\boldsymbol{\eta} = \mu \mathbf{I} + \mathbf{X}\boldsymbol{\beta},$$

where \mathbf{X} is an $n \times p$ matrix containing p variables for n samples (p can be $\gg n$). $\boldsymbol{\eta}$ is an $n \times 1$ linear predictor and is related to the response variable \mathbf{y} through a link function g : $E(\mathbf{y}|\mathbf{X})=g^{-1}(\mu \mathbf{I} + \mathbf{X}\boldsymbol{\beta})$, and $\boldsymbol{\beta}$ is a $p \times 1$ vector of regression coefficients. Depending on certain assumption of the data distribution on \mathbf{y} , the GLM is generally inferred through finding the set of model parameters that maximize the model likelihood function $p(\mathbf{y}|\mu, \boldsymbol{\beta}, \varphi)$, where φ denotes the other model parameters of the data distribution. However, such Maximum Likelihood (ML) approach is no longer applicable when $p \gg n$. With Bayesian lasso and Bayesian elastic net (EN) prior distribution on $\boldsymbol{\beta}$, **EBglmnet** solves the problem by inferring a sparse posterior distribution for $\hat{\boldsymbol{\beta}}$, which includes exactly zero regression coefficients for irrelevant variables and both posterior mean and variance for non-zero ones. Comparing with the `glmnet` package, not only does **EBglmnet** provide features including both sparse outcome and hypothesis testing, simulation study and real data analysis in the reference papers also demonstrates the better performance in terms of Power of Detection (PD), False Discovery Rate (FDR), as well as Power Detecting Group Effects when applicable. While mathematical details of the **EBlasso** and **EBEN** methods can be found in the reference papers, the principle of the methods and differences on the prior distributions will be briefly introduced here.

Lasso and its Bayesian Interpretation

Lasso applies a penalty term on the log likelihood function and solves for $\hat{\boldsymbol{\beta}}$ by maximizing the following penalized likelihood:

$$\hat{\boldsymbol{\beta}} = \arg_{\boldsymbol{\beta}} \max [\log p(\mathbf{y}|\mu, \boldsymbol{\beta}, \varphi) - \lambda \|\boldsymbol{\beta}\|_1].$$

The l_1 penalty term can be regarded as a mixture of hierarchical prior distribution:

$$\beta_j \sim N(0, \sigma_j^2), \sigma_j^2 \sim \exp(\lambda), j = 1, \dots, p,$$

and maximizing the penalized likelihood function is equivalent to maximizing the marginal posterior distribution of $\boldsymbol{\beta}$:

$$\hat{\boldsymbol{\beta}} = \arg_{\boldsymbol{\beta}} \max \log p(\boldsymbol{\beta}|\mathbf{y}, \mathbf{X}, \mu, \lambda, \varphi) \approx \arg_{\boldsymbol{\beta}} \max \log \int \left[p(\mathbf{y}|\mu, \boldsymbol{\beta}, \varphi) \cdot (2\pi)^{-p/2} |\mathbf{A}|^{1/2} \exp\left\{-\frac{1}{2}\boldsymbol{\beta}^T \mathbf{A}\boldsymbol{\beta}\right\} \cdot \prod_{j=1}^p \lambda \exp\{-\lambda\sigma_j^2\} \right] d\boldsymbol{\sigma}^2,$$

where \mathbf{A} is a diagonal matrix with σ^{-2} on its diagonal. Of note, `lasso` integrates out the variance information σ^2 and estimates a posterior mode $\hat{\beta}$. The l_1 penalty ensures that a sparse solution can be achieved. In Bayesian lasso (Park and Casella, 2008), the prior probability distribution is also conditional on the residual variance so that it has a unique mode.

Empirical Bayesian Lasso (EBlasso)

`EBglmnet` keeps the variance information, while still enjoying the sparse property by taking a different and slightly complicated approach as showed below using `EBlasso-NE` as an example:

In contrary to the marginalization on β , the first step in `EBlasso-NE` is to obtain a marginal posterior distribution for σ^2 :

$$p(\sigma^2 | \mathbf{y}, \mathbf{X}, \mu, \lambda, \varphi) = \int c \left[p(\mathbf{y} | \mu, \beta, \varphi) \cdot (2\pi)^{-p/2} |\mathbf{A}|^{1/2} \exp\left\{-\frac{1}{2} \beta^T \mathbf{A} \beta\right\} \cdot \prod_{j=1}^p \lambda \exp\{-\lambda \sigma_j^2\} \right] d\beta,$$

where c is a constant. While the integral in `lasso` is achieved through the conjugated Normal + Exponential (NE) prior, the integral in `EBlasso-NE` is completed through mixture of two normal distributions: $p(\beta | \sigma^2)$ and $p(\mathbf{y} | \mu, \beta, \varphi)$, and the latter one is typically approximated to a normal distribution through Laplace approximation if itself is not a normal PDF. Then the estimate $\hat{\sigma}^2$ can be obtained by maximizing this marginal posterior distribution, which has the following form:

$$\hat{\sigma}^2 = \arg_{\sigma^2} \max \log p(\sigma^2 | \mathbf{y}, \mathbf{X}, \mu, \lambda, \varphi) = \arg_{\sigma^2} \max \left[\log p(\mathbf{y} | \mu, \sigma^2, \varphi, \lambda) - \sum_{j=1}^p \lambda \sigma_j^2 + c \right].$$

Given the constraint that $\sigma^2 > 0$, the above equation essentially maximizes the l_1 penalized marginal likelihood function of σ^2 , which images the l_1 penalty in `lasso` with the beauty of producing a sparse solution for $\hat{\sigma}^2$. Note that if $\hat{\sigma}_j^2 = 0$, $\hat{\beta}_j$ will also be zero and variable x_j will be excluded from the model. Finally, with the sparse estimate of $\hat{\sigma}^2$, the posterior estimate of $\hat{\beta}$ and other nuance parameters can then be obtained accordingly.

Hierarchical Prior Distributions in `EBglmnet`

Prior 1: `EBlasso-NE`

$$\beta_j \sim N(0, \sigma_j^2), \sigma_j^2 \sim \exp(\lambda), j = 1, \dots, p$$

As illustrated above, assuming a Normal + Exponential hierarchical prior distribution on β (`EBlasso-NE`) will yield exactly the lasso prior. `EBlasso-NE` accommodates the properties of sparse solution and hypothesis testing given both the estimated mean and variance information in $\hat{\beta}$ and $\hat{\sigma}^2$. The NE prior is “peak zero and flat tails”, which can select variables with relatively small effect size while shrinking most of irrelevant effects to exactly zero. `EBlasso-NE` can be applied to natural population analysis when effect sizes are relatively small.

Prior 2: `EBlasso-NEG` The prior in `EBlasso-NE` has a relatively large probability mass on the nonzero tails, resulting in a large number of non-zero small effects with large p -values in simulation and real data analysis. We further developed another well studied conjugated hierarchical prior distribution under the empirical Bayesian framework, the Normal + Exponential + Gamma (NEG) prior:

$$\beta_j \sim N(0, \sigma_j^2), \sigma_j^2 \sim \exp(\lambda), \lambda \sim \text{gamma}(a, b), j = 1, \dots, p$$

Comparing with EBlasso-NE, the NEG prior has both large probability density centered at 0 and two tails, and will only yield nonzero regression coefficients for effects having relatively large signal to noise ratio.

Prior 3: Elastic Net Prior for Grouping Effect Similar to lasso, EBlasso typically selects one variable out of a group of correlated variables. While `elastic net` was developed to encourage a grouping effect by incorporating an l_2 penalty term, EBglmnet implemented an innovative `elastic net` hierarchical prior:

$$\beta_j \sim N [0, (\lambda_1 + \tilde{\sigma}_j^{-2})^{-1}], \tilde{\sigma}_j^2 \sim \text{generalized gamma}(\lambda_1, \lambda_2), j = 1, \dots, p.$$

The generalized gamma distribution has PDF: $f(\tilde{\sigma}_j^2 | \lambda_1, \lambda_2) = c(\lambda_1 \tilde{\sigma}_j^2 + 1)^{-1/2} \exp\{-\lambda_2 \tilde{\sigma}_j^2\}$, $j = 1, \dots, p$, with c being a normalization constant. The property of this prior can be appreciated from the following aspects:

(1): $\lambda_1 = 0$ When $\lambda_1 = 0$, the generalized gamma distribution becomes an exponential distribution: $f(\tilde{\sigma}_j^2 | \lambda_2) = c \exp\{-\lambda_2 \tilde{\sigma}_j^2\}$, $j = 1, \dots, p$, with $c = \lambda_2$, and the elastic net prior is reduced to the two level EBlasso-NE prior.

(2): $\lambda_1 > 0$ When $\lambda_1 > 0$, the generalized gamma distribution can be written as a shift gamma distribution having the following PDF: $f(\tilde{\sigma}_j^2 | a, b, \gamma) = \frac{b^a}{\Gamma(a)} (\tilde{\sigma}_j^2 - \gamma)^{a-1} \exp\{-b(\tilde{\sigma}_j^2 - \gamma)\}$, where $a = 1/2$, $b = \lambda_2$, and $\gamma = -1/\lambda_1$. In (Huang A. 2015), it is proved that the marginal prior distribution of β_j can be obtained as $p(\beta_j) \propto \exp\{-\frac{\lambda_1}{2} \beta_j^2 - \sqrt{2\lambda_2} |\beta_j|\}$, which is equivalent to the `elastic net` method in `glmnet`.

(3): structure of σ^2 and interpretation of the elastic net prior Note that the prior variance for the regression coefficients has this form: $\sigma^2 = \tilde{\sigma}^2 / (\lambda_1 \tilde{\sigma}^2 + \mathbf{I})$. This structure seems counter-intuitive at first glance. However, if we look at it from precision point of view, i.e., $\alpha = \sigma^{-2}$, and $\tilde{\alpha} = \tilde{\sigma}^{-2}$, then we have:

$$\alpha = \lambda_1 \mathbf{I} + \tilde{\alpha}.$$

The above equation demonstrates that we actually decompose the precision of the normal prior into a fixed component λ_1 shared by all explanatory variables and a random component $\tilde{\alpha}$ that is unique for each explanatory variable. This design represents the mathematical balance between the inter-group independence and intra-group correlation among explanatory variables, and is aligned with the objective of sparseness while encouraging grouping effects.

The empirical Bayesian elastic net (EBEN) in EBglmnet is solved similarly as EBlasso using the aforementioned empirical Bayesian approach. Research studies presented in the reference papers demonstrated that EBEN has better performance comparing with `elastic net`, in terms of PD, FDR, and most importantly, Power of Detecting Groups.

EBglmnet Implementation and Usage

The EBglmnet algorithms use greedy coordinate descent, which successively optimizes the objective function over each parameter with others fixed, and cycles repeatedly until convergence. Key algorithms are implemented in C/C++ with matrix computation using the BLAS/LAPACK packages. Due to closed form solutions for $\tilde{\sigma}^2$ in all prior setups and other algorithmic and programming techniques, the algorithms can compute the solutions very fast.

We recommend using EBlasso-NEG when there are a large number of candidate effects (eg., $\geq 10^6$ number of effects such as whole-genome epistasis analysis and GWAS), and using EBEN when groups of highly correlated variables are of interest.

The authors of `EBglmnet` are Anhui Huang and Dianting Liu. This vignette describes the principle and usage of `EBglmnet` in R. Users are referred to the papers in the Reference section for details of the algorithms.

Installation

`EBglmnet` can be installed directly from CRAN using the following command in R console:

```
install.packages("EBglmnet", repos = "http://cran.us.r-project.org")
```

which will download and install the package to the default directory. Alternatively, users can download the pre-compiled binary file from CRAN and install it from local package.

[Back to Top](#)

Quick Start

We will give users a general idea of the package by using a simple example that demonstrates the basis package usage. Through running the main functions and examining the outputs, users may have a better idea on how the package works, what functions are available, which parameters to choose, as well as where to seek help. More details are given in later sections.

Let us first clear up the workspace and load the `EBglmnet` package:

```
rm(list = ls())  
library(EBglmnet)
```

We will use an R built-in dataset `state.x77` as an example, which includes a matrix with 50 rows and 8 columns giving the following measurements in the respective columns: Population, Income, Illiteracy, Life Expectancy, Murder Rate, High School Graduate Rate, Days Below Freezing Temperature, and Land Area. The default model used in the package is the Gaussian linear model, and we will demonstrate it using Life Expectancy as the response variable and the remaining as explanatory variables. We create the input data as showed below, and users can load their own data and prepare variable `y` and `x` following this example.

```
varNames = colnames(state.x77);  
varNames  
y = state.x77[, "Life Exp"]  
xNames = c("Population", "Income", "Illiteracy", "Murder", "HS Grad", "Frost", "Area")  
x = state.x77[, xNames]
```

We fit the model using the most basic call to `EBglmnet` with default prior

```
set.seed(1)  
output = EBglmnet(x, y, hyperparameters = c(0.1, 0.1))
```

“output” is a list containing all the relevant information of the fitted model. Users can examine the output by directly looking at each element in the list. Particularly, the sparse regression coefficients can be extracted as showed below:

```
glmfit = output$fit  
variables = xNames[glmfit[, 1, drop=FALSE]]  
cbind(variables, as.data.frame(round(glmfit[, 2:5, drop=FALSE], 4)))
```

The hyperparameters in each of the prior distributions control the number of non-zero effects to be selected, and cross-validation (CV) is perhaps the simplest and most widely used method in deciding their values. `cv.EBglmnet` is the main function to do cross-validation, which can be called using the following code.

```
cvfit = cv.EBglmnet(x, y)
```

`cv.EBglmnet` returns a `cv.EBglmnet` object, which is a list with all the ingredients of CV and the final fit results using CV selected hyperparameters. We can view the CV results, selected hyperparameters and the corresponding coefficients. For example, result using different hyperparameters and the corresponding prediction errors are shown below:

```
cvfit$CrossValidation
```

The selected parameters and the corresponding fitting results are:

```
cvfit$hyperparameters
cvfit$fit
```

[Back to Top](#)

GLM Family

Two families of models have been developed in `EBglmnet`, the **gaussian** family and the **binomial** family, which are essentially different probability distribution assumptions on the response variable y .

Gaussian Model

`EBglmnet` assumes a Gaussian distribution on y by default, i.e., $p(\mathbf{y}|\mu, \beta, \varphi) = N(\mu\mathbf{I} + \mathbf{X}\beta, \sigma_0^2\mathbf{I})$, where $\varphi = \sigma_0^2$ is the residual variance. In the above example, both $\hat{\mu}$ and $\hat{\sigma}_0^2$ are listed in the output:

```
output$Intercept
output$residual
```

Binomial Model

If there are two possible outcomes in the response variable, a binomial distribution assumption on y is available in `EBglmnet`, which has $p(\mathbf{y}|\mu, \beta, \varphi)$ following a binomial distribution and $\varphi \in \emptyset$. Same as the widely-used logistic regression model, the link function is $\eta_i = \text{logit}(p_i) = \log\left(\frac{\text{Pr}(y_i=1)}{1-\text{Pr}(y_i=1)}\right)$, $i = 1, \dots, n$. To run `EBglmnet` with binomial models, users need to specify the parameter `family` as **binomial**:

```
yy = y>mean(y);
output = EBglmnet(x,yy,family="binomial", hyperparameters = c(0.1, 0.1))
```

For illustration purpose, the above codes created a binary variable `yy` by set the cutoff at the mean Life Expectancy value.

[Back to Top](#)

Prior and Hyperparameters

The three sets of hierarchical prior distribution can be specified by `prior` option in `EBglmnet`. By default, `EBglmnet` assumes the `lassoNEG` prior. The following example changes the prior via `prior` parameter:

```
output = EBglmnet(x,yy,family="binomial", prior = "elastic net", hyperparameters = c(0.1, 0.1))
output$fit
```

Note that the hyperparameters setup is associated with a specific prior. In `lasso` prior, only one hyperparameter λ is required, while in `elastic net` and `lassoNEG`, two hyperparameters need to be specified. For `EBEN` having the `elastic net` prior distribution, the two hyperparameters λ_1 and λ_2 are defined in terms of other two parameters $\alpha \in [0, 1]$ and $\lambda > 0$ same as in `glmnet` package, such that $\lambda_1 = (1 - \alpha)\lambda$ and $\lambda_2 = \alpha\lambda$. Therefore, users are asked to specify `hyperparameters = c(α, λ)`.

Back to Top

Example of $p > n$ Data

Gaussian Model with Main Effects

We will demonstrate the application of `EBglmnet` in multiple QTL mapping using a simulated F2 population, which is available along with `EBglmnet` package. The genotype of the F2 population is simulated from cross of two inbred lines. A total of $p = 481$ genetic markers were simulated to be evenly spaced on a large chromosome of 2400 centi-Morgan (cM) with an interval of $d = 5$ cM. Theoretically, two adjacent markers have a correlation coefficient $R = e^{-2d} = 0.9048$ since the Haldane map function is assumed. The dummy variable for the three genotypes, AA, Aa and aa of individual i at marker j was defined as $x_{ij} = 1, 0, -1$, respectively.

```
data(BASIS)#this is the predictors matrix
N = nrow(BASIS)
p = ncol(BASIS)
j = sample((p-2),1)
cor(BASIS[,c(j,j+1,j+2)]) #Correlation structure among predictors
```

```
##           m464           m465           m466
## m464 1.0000000 0.9172078 0.8249249
## m465 0.9172078 1.0000000 0.8955560
## m466 0.8249249 0.8955560 1.0000000
```

Let us first simulate a quantitative phenotype with population mean 100 and residual variance that contribute to 10% of population variance. Ten QTLs were simulated from the 481 markers, and effect sizes are randomly generated from [2,3]. We assume that QTLs were coincided with markers. If QTLs were not on markers, they may still be detected given the above correlation structure, although a slightly larger sample size may be needed to give the same PD:

```
set.seed(1);
Mu = 100; #population mean;
nTrue = 10; # we assume 10 out of the 481 predictors are true causal ones
trueLoc = sort(sample(p,nTrue));
trueEff = runif(nTrue,2,3); #effect size from 2-3
xbeta = BASIS[,trueLoc]%%trueEff;
s2 = var(xbeta)*0.1/0.9 #residual variance with 10% noise
residual = rnorm(N,mean=0,sd=sqrt(s2))
y = Mu + xbeta + residual;
```

To demonstrate the performance of EBglmnet in analyzing dataset with $p > n$, we will analyze the simulated dataset using a smaller sample size $n = 300$:

```
n = 300;
index = sample(N,n);
CV = cv.EBglmnet(x=BASIS[index,],y=y[index],family="gaussian",prior="lassoNEG",nfold= 5)
```

With 5 fold CV, EBlasso-NEG identified the following effects:

```
CV$fit
trueLoc
```

Comparing with the true QTL locations, EBlasso-NEG successfully identified all QTLs. Of note, an identified marker that is <20cM from a true QTL is generally not considered as a false effect, given the small distance and high genotype correlation. Since there is limited prior information in this simulation, other methods in EBglmnet will yield similar results.

Binomial Model and Separation Problem

In many genetics and population analysis such as binary QTL mapping or GWAS, both genetic effects (eg., \mathbf{X} takes values of 1, 0, -1 denoting three genotype values AA, Aa, aa) and response variable (eg., \mathbf{y} takes values of 0 and 1 denoting the two phenotype classes) are discrete values. Separation problem (complete separation and semi-complete separation) occurs often in such data, especially when epistasis is considered due to the larger number of variables $p' = p(p+1)/2$. Of note, separation is a problem in logistic regression where there exist some coefficients β such that $y_i = 1$ whenever $\mathbf{x}_i^T \beta > 0$, and $y_i = 0$ whenever $\mathbf{x}_i^T \beta < 0, i = 1, \dots, n$. Unless the phenotype is a Mendelian trait, finding a genetic factor/set of factors that perfectly predict the phenotype outcome is too good to be true. While separation problem has been well documented in many statistical textbook (eg., Ch9 of Altman M, Gill J, McDonald M P. (2005)), it is less studied in high dimensional sparse modeling methods. We next examine the problem using lasso as an example.

Benchmarking using EBglmnet In the simulated F2 population, if we consider both main and epistatic effects, there will be a total number $p' = 115,921$ candidate effects. We will randomly select 10 main and 10 interaction effects as true effects. we will use sample size $n = 300$, and also randomly generate effect sizes from unif[2,3]. Note that an epistatic effect is generated by dot product of two interacting main effects:

```
n = 300;
set.seed(1)
index = sample(nrow(BASIS),n)
p = ncol(BASIS);
m = p*(p+1)/2;
#1. simulate true causal effects
nMain = 10;
nEpis = 10;
mainLoc = sample(p,nMain);
episLoc = sample(seq((p+1),m,1),nEpis);
trueLoc = sort(c(mainLoc,episLoc)); #a vector in [1,m]
nTrue = length(trueLoc);
trueLocs = ijIndex(trueLoc, p); #two columns denoting the pair (i,j)
#2. obtain true predictors
basis = matrix(0,n,nTrue);
for(i in 1:nTrue)
{
```

```

if(trueLocs[i,1]==trueLocs[i,2])
{
  basis[,i] = BASIS[index,trueLocs[i,1]]
}else
{
  basis[,i] = BASIS[index,trueLocs[i,1]]*BASIS[index,trueLocs[i,2]]
}
}
#3. simulate true effect size
trueEff = runif(nTrue,2,3);
#4. simulate response variables
xbeta = basis%%trueEff;
vary = var(xbeta);
Pr = 1/(1+ exp( -xbeta));
y = rbinom(n,1,Pr);

```

Now we have the phenotype simulated. Let us demonstrate the data analysis using EBlasso-NE as an example. Given the significant larger number of candidate effects (200 times more effects), this will take a longer time for CV to finish ($n_{folds} \times n_{hyperparameters} + 1$) times of calling EBglmnet algorithm. In fact, the computational time is mostly determined by the number of nonzero effects selected by the model. This can be seen if a larger sample size is used (note: if you set $n = 1000$, this will take several more hours to finish the $(n_{folds} \times n_{hyperparameters} + 1)$ times computation (Although a higher PD will be obtained).

```
CV = cv.EBglmnet(x=BASIS[index,],y=y,family="binomial",prior="lasso",nfold=5)
```

```
## Empirical Bayes LASSO Logistic Model (Normal + Exponential prior) 5 fold cross-validation
```

```
ind = which(CV$fit[,5]<=0.1)#p-value cutoff
CV$fit[ind,]
```

##	predictor	beta	posterior variance	t-value	p-value	
##	[1,]	54	0.3574937	0.04064460	1.773238	0.077207261
##	[2,]	56	0.6570230	0.05495714	2.802648	0.005399570
##	[3,]	72	0.5370808	0.03918611	2.713148	0.007051139
##	[4,]	91	0.5854443	0.04522653	2.752890	0.006268404
##	[5,]	210	0.5296459	0.03667812	2.765554	0.006036039
##	[6,]	225	0.5406408	0.05317435	2.344540	0.019704733
##	[7,]	227	0.4852761	0.04927319	2.186168	0.029578514
##	[8,]	239	0.4313054	0.04712116	1.986904	0.047845428
##	[9,]	240	0.3445439	0.04333749	1.655056	0.098962334
##	[10,]	270	0.5462582	0.04604777	2.545621	0.011410310
##	[11,]	293	-0.3105339	0.02784121	1.861080	0.063714242
##	[12,]	336	0.3962398	0.03394442	2.150671	0.032303559

As discussed earlier, lasso prior assigns a large probability mass in the two tails, resulting in a large number of small effects with large p -values.

Separation Problem Compared with EBglmnet, glmnet does not have a reasonable result in analyzing this dataset, partially because of the separation problem. Let us first show the result using lasso approach.

Since glmnet has no built-in facility for epistasis analysis, we will manually create a genotype matrix \mathbf{X} containing all main and epistasis effects.

```

X = matrix(0,n,m);
X[,1:p] = BASIS[index,];
kk = p + 1;
for(i in 1:(p-1))
{
  for(j in (i+1):p)
  {
    X[,kk] = BASIS[index,i] * BASIS[index,j];
    kk = kk + 1;
  }
}

```

Let us analyze the dataset using lasso, examine the lasso selection path:

```
library(glmnet);
```

```
## Loading required package: Matrix
```

```
## Loaded glmnet 4.1-6
```

```

alpha = 1
lambdaRatio = 1e-4; #same as in EBlasso
cv = cv.glmnet(X, y, alpha = alpha,family="binomial",nfolds = 5,lambda.min.ratio=lambdaRatio)
nLambda = length(cv$lambda)
nLambda

```

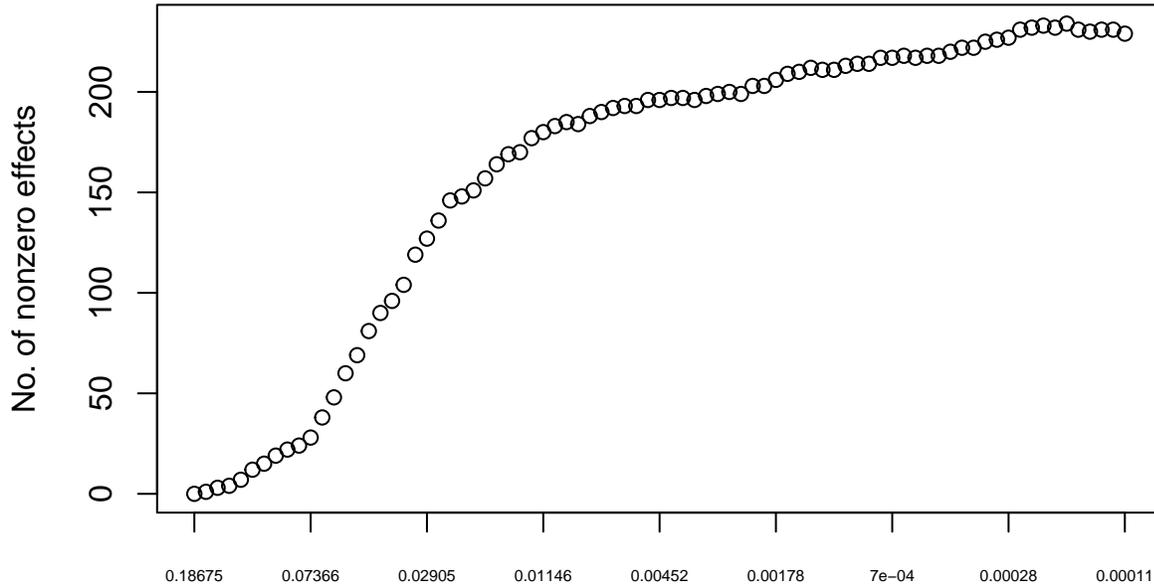
```
## [1] 81
```

```

nbeta = rep(0,nLambda);
fit0 = cv$glmnet.fit;
for(i in 1:nLambda)
{
  nbeta[i] = length(which(fit0$beta[,i]!=0))
}
plot(nbeta,xlab=expression(paste(lambda, " in lasso selection path(n=300,p=115,921)")),
     ylab="No. of nonzero effects",xaxt="n")#
ticks = seq(1,nLambda,10)
axis(side=1, at= ticks,labels=round(cv$lambda[ticks],5), las=1,cex.axis = 0.5)
title("Number of nonzero effects in lasso selection path")

```

Number of nonzero effects in lasso selection path



λ in lasso selection path($n=300, p=115,921$)

The result above demonstrated that lasso was not able to complete the selection path, and exited at the 81 out of 100 candidate λ s. From the scatterplot, it is shown that the number of nonzero effects has stabilized after the 25th λ due to semi-complete separation that fewer candidate variables can be selected, even lambda is decreasing exponentially (will be explained in lasso discarding rule in the ensuing section). See the `glmnet` user manual for the built-in exit mechanism.

We can also take a closer look by re-fitting the lasso selected effects in an ordinary logistic regression model, which will explicitly print out the warning message of separation:

```
lambda= cv$lambda.min
coefs = fit0$beta
ind = which(cv$lambda==cv$lambda.min)
beta = coefs[,ind]
betaij = which(beta!=0)
Xdata = X[,betaij];
colnames(Xdata) = betaij;
refit = glm(y ~ Xdata, family = binomial(link = "logit"))#separation occurs
```

```
## Warning: glm.fit: algorithm did not converge
```

```
## Warning: glm.fit: fitted probabilities numerically 0 or 1 occurred
```

The warning message describes that separation occurs. Separation problem is detrimental to lasso/elastic net due to the discarding rules (Tibshirani et al., 2012). Of note, the lasso discarding rule for logistic regression states that variable x_j can be discarded if:

$$|\mathbf{x}_j^T(\mathbf{y} - \mathbf{p}(\hat{\boldsymbol{\beta}}_{\lambda_{k-1}}))| < 2\lambda - \lambda_{k-1}, \forall j$$

where $\hat{\boldsymbol{\beta}}_{\lambda_{k-1}}$ is the nonzero coefficients found by lasso using λ_{k-1} at the $k-1$ step. Suppose with λ_{k-1} , lasso selected a set of variables $\tilde{\mathbf{X}}$ that perfectly separate \mathbf{y} , which lead to $\mathbf{y} - \mathbf{p}(\hat{\boldsymbol{\beta}}_{\lambda_{k-1}}) \approx 0$. Then, the

above discarding rule will have most of the remaining variables discarded. Note that $\mathbf{y} - \mathbf{p}(\hat{\boldsymbol{\beta}}_{\lambda_{k-1}})$ will not be exactly 0 due to numerical estimations, and as λ decreases exponentially, a few variables can still pass the discarding rule (shown from 25th - 81st λ in the above example). However, finding a genetic factor/set of factors that perfectly predict the phenotype outcome is unlikely, and it is the case given the simulation setup. With the perfect separation provided by $\tilde{\mathbf{X}}$, other true effects cannot be selected into the non-zero set, resulting in limited PD.

EBglmnet doesn't implement such type of discarding rule and is more numerically stable. In the above simulation, EBglmnet can still identify several true effects with reasonable FDR. More simulation results using EBglmnet are available in (Huang et al., 2013) and the EBglmnet Application Note.

References

EBglmnet Algorithms:

Anhui Huang, Shizhong Xu, and Xiaodong Cai. (2015). Empirical Bayesian elastic net for multiple quantitative trait locus mapping. *Heredity*, Vol. 114(1), 107-115.

Anhui Huang, Shizhong Xu, and Xiaodong Cai. (2014a). Whole-genome quantitative trait locus mapping reveals major role of epistasis on yield of rice. *PLoS ONE*, Vol. 9(1) e87330.

Anhui Huang, Eden Martin, Jeffery Vance, and Xiaodong Cai (2014b). Detecting genetic interactions in pathway-based genome-wide association studies. *Genetic Epidemiology*, 38(4), 300-309.

Anhui Huang, Shizhong Xu, and Xiaodong Cai. (2013). Empirical Bayesian LASSO-logistic regression for multiple binary trait locus mapping. *BMC Genetics*, 14(1),5.

Xiaodong Cai, Anhui Huang, and Shizhong Xu (2011). Fast empirical Bayesian LASSO for multiple quantitative trait locus mapping. *BMC Bioinformatics*, 12(1),211.

Lasso and glmnet package

Tibshirani, R., Bien, J., Friedman, J., Hastie, T., Simon, N., Taylor, J., and Tibshirani, R.J., 2012. Strong rules for discarding predictors in lasso-type problems. *J. R. Stat. Soc. Series B. Stat. Methodol.* 74, 245-266.

Friedman, J., T. Hastie, et al. (2010). Regularization paths for generalized linear models via coordinate descent. *J. Stat. Softw.* 33(1): 1-22.

Park, T. and G. Casella (2008). The Bayesian lasso. *J. Am. Stat. Assoc.* 103(482): 681-686.

Yi, N., and S. Xu (2008). Bayesian LASSO for quantitative trait loci mapping. *Genetics* 179(2): 1045-1055.

Zou, H., and T. Hastie (2005). Regularization and variable selection via the elastic net. *J. Roy. Stat. Soc. B. Met.* 67(2): 301-320.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *J. Roy. Stat. Soc. B. Met.* 58(1): 267-288.