

# GlmNet for Exponential Response Variables via RcppEigen

Xin Chen, Doug Martin, Aleksandr Aravkin, Dan Hanson, Anthony Christidis

September 1, 2019

## Abstract

This package implements the Generalized Linear Model with Elastic Net regularization for Exponential Response Variables (GLM-EN-EXP). This package is intended to fill the current gap in the R software ecosystem where an implementation that 1) supports exponential distribution, 2) supports Elastic Net model selection, 3) is easy to parallelize on multicore machine, does not exist, to the best knowledge of the authors. Significant speed improvement has been shown compared with both native R and H2O implementations in simple benchmark tests.

## 1 Theoretical Background

### 1.1 Generalized Linear Models For Exponential Distributions

The theory and methodology of generalized linear models (GLM's) is well established for independent observations  $Y_1, Y_2, \dots, Y_N$  whose mean values are  $\mu_1, \mu_2, \dots, \mu_N$  and whose distributions are members of the *exponential families*

$$f(y_k; \theta_k) = \exp \{ (y_k \theta_k - b(\theta_k)) / a(\phi_k) + c(y_k, \phi_k) \} \quad k = 1, 2, \dots, N \quad (1.1)$$

along with a link function  $g$  from the  $\mu_k$  to a linear model form  $\mathbf{x}'_k \boldsymbol{\beta}$ :

$$g(\mu_k) = \mathbf{x}'_k \boldsymbol{\beta} = \eta_k. \quad (1.2)$$

The form stated above with  $y_k$  in the first term in the exponent rather than some function  $a(y_k)$  is called the *canonical* form, and  $\theta_k$  is called the *natural parameter*. See for example McCullagh and Nelder (1989) and Fox (2016)<sup>1</sup>, where one finds simple derivations of the following expressions for the mean and variance of the observations:

$$\mu_k = E(Y_k) = b'(\theta_k) \quad (1.3)$$

$$\text{var}(Y_k) = b''(\theta_k) a(\phi_k). \quad (1.4)$$

### 1.2 Maximum Likelihood Model Estimation for GLM

The likelihood function for observed values  $y_1, y_2, \dots, y_N$  of the independent set  $Y_1, Y_2, \dots, Y_N$  is

$$L(\mathbf{y}; \boldsymbol{\theta}) = \prod_{k=1}^N f(y_k; \theta_k)$$

and the log-likelihood function is

$$l(\mathbf{y}; \boldsymbol{\theta}) = \log \prod_{k=1}^N f(y_k; \theta_k) = \sum_{k=1}^N \log f(y_k; \theta_k) = \sum_{k=1}^N l(y_k; \theta_k).$$

The classic GLM model is usually fitted using a method called iterative reweighted least squares (IRWLS) method.

---

<sup>1</sup>Other authors use slightly different but equivalent mathematical forms of an exponential family, e.g., see Dobson(2002)

### 1.3 GLM with Elastic Net Regularization

One problem with the standard GLM formulation is that it does not have regularization penalty for overfitting. More Regularization usually comes in the form of augmenting the original cost function (such as the log-likelihood function) with a penalty term that expresses the modeler's belief about the properties of the model. For example, in ridge regularization, the sum of least square errors is augmented by the L2 norm of the coefficients

$$\hat{\beta}_{Ridge} = \underset{\beta, \lambda}{\operatorname{argmin}} LL(\mathbf{y}; \boldsymbol{\theta}) + \lambda \|\boldsymbol{\beta}\|_2$$

This implies that the modeler believes there is unlikely to be particularly large coefficients in the model. In deed, large coefficients will be penalized heavily by the L2 regularization, resulting in smaller values for the coefficients.

A different kind of regularization, which became very popular in recent years, is called least absolute shrinkage and selection operator (LASSO). The idea is similar to that of ridge regression, except that instead of the believing the coefficients should not be too large, the modeler believes the model should be sparse with respect to the variables. This is achieved by augmenting the MSE with L1 norm of the coefficients.

$$\hat{\beta}_{LASSO} = \underset{\beta, \lambda}{\operatorname{argmin}} LL(\mathbf{y}; \boldsymbol{\theta}) + \lambda \|\boldsymbol{\beta}\|_1$$

The sparsity of  $\hat{\beta}_{LASSO}$  is controlled by adjusting the value of  $\lambda$ . One thing to note is that even though  $\hat{\beta}_{LASSO}$  tends to be sparse, it could very well have large coefficients, which may or may not be desirable.

The latest development in regularization comes in the form of Elastic Net (EN). It combines Ridge Regression and LASSO regression by using weighted sum of L1 norm and L2 norm of the coefficients as the penalty function.

$$\hat{\beta}_{EN} = \underset{\beta, \lambda}{\operatorname{argmin}} LL(\mathbf{y}; \boldsymbol{\theta}) + \lambda(\alpha \|\boldsymbol{\beta}\|_1 + (1 - \alpha) \|\boldsymbol{\beta}\|_2)$$

The parameter  $\alpha \in [0, 1]$  expresses our belief about the relative importance of sparsity and coefficient value.

### 1.4 GLM-EN via Proximal Gradient Descent

The proximal gradient descent algorithm is designed to solve problems of the form

$$\min_x f(x) + g(x)$$

where  $f(x)$  and  $g(x)$  are closed proper convex and  $f(x)$  is differentiable. The proximal gradient method is

$$x^{k+1} = \operatorname{prox}_{t^k g}(x^k - \nabla f(x))$$

where

$$\operatorname{prox}_{tg}(v) = \min_x f(x) + \frac{1}{2t} \|x - v\|_2^2$$

For our problem,  $f(x)$  would be the sum negative log-likelihood of the exponential response variables and the L2 regularization and  $g(x)$  would be the L1 regularization. Therefore,

$$\begin{aligned} \operatorname{prox}_{tg}(v) &= \operatorname{prox}_{t\|\cdot\|_1}(v) = \operatorname{argmin}_x \left( \|x\|_1 + \frac{1}{2t} \|x - v\|_2^2 \right) \\ &= \begin{cases} v_i - t & v_i \geq t \\ v_i + t & v_i \leq -t \\ 0 & |v_i| \leq t \end{cases} \end{aligned}$$

## 2 Sample Code

### 2.1 Structure of the Sample Code for Timings and Results Comparison between RPEGLMEN and H2O

To compare our implementation with H2O, we simulate test data and then use both algorithms to try to find the true coefficients used to generate the simulated data. We assume that the data follows the model

$$b_i \sim \exp(\lambda_i)$$

where

$$\log \lambda_i = -A_i * x$$

The test data set is generated by the following steps:

1. Randomly generate the independent variable matrix  $A$  of size  $n \times p$ , where  $n$  is the number of observations,  $p$  is the number of independent variables.
2. Randomly generate the true coefficient of the model  $x.true$ , which is a vector of length  $p$ .
3. Compute the vector of parameters for the exponential distributions,  $\lambda = \exp(-A * x.true)$
4. For each  $\lambda_i$ , generate a sample from the exponential distribution. This is the response variable  $b_i$ .

Once the test data is generated, we apply `glmnet_exp()` and `h2o.glm()` to compute the estimated coefficients  $x$ , and compare it with the true coefficients  $x.true$ .

### 2.2 Sample Code for Timings and Results Comparison between RPEGLMEN and H2O

```
rm(list = ls())

##### load package
library(RPEGLMEN)
library(h2o)

##### helper function
compute_mean_vector = function(data_list){
  tmp_mat = matrix(unlist(data_list), nrow = length(data_list[[1]]))
  apply(tmp_mat, 1, mean)
}

##### set parameters
set.seed(20170827)
nobs = 50
nvars = 7
ntests = 10 # takes a while to run h2o
alpha = 0.5
lasso.lambda = 0.1
x.true = rnorm(nvars)
x.true[sample(2:nvars, floor(nvars / 2) - 1)] = 0

##### generate test data

params_list = list()
for(j in 1:ntests){
  ## random normal matrix A
  A = matrix(rnorm(nvars * nobs), ncol = nvars)
  exp.lambdas = exp(-A %*% x.true)
  b = sapply(exp.lambdas, function(x) rexp(1, x))
```

```

params_list[[j]] = list(
  b = b,
  # The response variables
  A = A,
  # The matrix
  alpha = alpha,
  # 1 means lasso
  lambda = lasso.lambda
) # the regularization coefficient
}

#### try doing things parallel using foreach
#### note that h2o.glm does not work with foreach
library(doParallel)
library(foreach)

h2o.init()
h2o_time_single = system.time({h2o_single_res_list = lapply(params_list,
  function(dat){

    b = dat[["b"]]
    A = dat[["A"]]
    alpha = dat[["alpha"]]
    x.h2o.df = as.h2o(data.frame(b, A))
    predictors = colnames(x.h2o.df)[-1]
    response = colnames(x.h2o.df)[1]
    my.glm.lasso = h2o.glm(
      x = predictors,
      y = response,
      family = 'gamma',
      intercept = FALSE,
      training_frame = x.h2o.df,
      ignore_const_cols = TRUE,
      link = "log",
      # lambda = enet_lambda,
      lambda_search = TRUE,
      alpha = alpha,
      standardize = FALSE
    )
    return(my.glm.lasso@model$coefficients[-1])
  }
)
}
)
h2o.shutdown(FALSE)

cl <- makeCluster(3)
registerDoParallel(cl)

time_single = system.time({cpp_single_res_list = foreach(i = 1:length(params_list),
  .packages = 'RPEGLMEN') %do% {
  x = params_list[[i]]
  glmnet_exp(x[["A"]], x[["b"]], x[["alpha"]])
}
})

time_multi = system.time({cpp_multi_res_list = foreach(i = 1:length(params_list),
  .packages = 'RPEGLMEN') %dopar% {
  x = params_list[[i]]
  glmnet_exp(x[["A"]], x[["b"]], x[["alpha"]])
}
})

```

```

})
stopCluster(c1)
res_cpp_single = compute_mean_vector(cpp_single_res_list)
res_cpp_multi = compute_mean_vector(cpp_multi_res_list)
res_h2o_single = compute_mean_vector(h2o_single_res_list)

```

```

### timings
data.frame(cpp_time_single = time_single[3],
           cpp_time_multi = time_multi[3],
           h2o_time_single = h2o_time_single[3])

##           cpp_time_single cpp_time_multi h2o_time_single
## elapsed                0.03             0.26             18.13

```

```

### fitted coefficients
data.frame(x.true = x.true,
           x.single = res_cpp_single,
           x.multi = res_cpp_multi,
           x.h2o = res_h2o_single)

##           x.true      x.single      x.multi      x.h2o
## 1  1.23913750  1.194611941  1.191183236  0.96848325
## 2  0.00000000  0.043368134  0.034279220  0.01700905
## 3  0.44916967  0.464882100  0.472978607  0.33655192
## 4  0.02784615 -0.037922908 -0.040167674 -0.01955223
## 5 -1.58000686 -1.543750634 -1.550456911 -1.35342157
## 6  0.00000000 -0.000354347  0.005115291  0.01172726
## 7  0.54055576  0.466366411  0.475149283  0.32159844

```

### 2.3 Testing RPEGLMEN for estimating the standard errors of sample mean

The following code chunk tests using RPEGLMEN to estimate the standard error of sample for IID normal data.

```

rm(list=ls())
library(RPEGLMEN)

#' Compute the periodogram as defined in H&W 1981
#'
#' @param data Vector of data
#' @param max.freq Maximum frequency to be computed
#'
#' @return list of frequencies and corresponding periodograms
#' @export
#'
#' @examples
#' myperiodogram(rnorm(10))
myperiodogram=function(data,max.freq=0.5){
  ## data.fft=myfft(data) This is very slow
  data.fft=fft(data)
  N=length(data)
  # tmp=1:N
  # inset=tmp[(1:N)<floor(N/2)]
  tmp = Mod(data.fft[1:floor(N/2)])^2/N
  tmp = sapply(tmp, function(x) max(0.01,x))
  return(list(spec=tmp,freq=((0:(floor(N/2)-1))/N)))
}

```

```

#' Influence Function of Mean
#'
#' @param data Vector of the data
#' @param ... other parameters
#'
#' @return IF of Mean
#' @export
#' @author Xin Chen, \email{chenx26@uw.edu}
#'
#' @examples
#' mu.IF(rnorm(10))
mu.IF=function(data,...){
  mu.hat=mean(data)
  return(data-mu.hat)
}

#' Compute the standard error of the mean of the data using glmnet_exp
#'
#' @param data vector of data
#' @param d degree of polynomial
#' @param alpha weight for Elastic Net regularizer
#' @param keep what portion of sample spectral density to use for model fitting
#'
#' @return variance of the mean of the data
#' @export
#'
SE.glmnet_exp=function(data, d=7, alpha=0.5, keep=1){

  N=length(data)
  # Step 1: compute the periodograms
  my.periodogram=myperiodogram(data)
  my.freq=my.periodogram$freq
  my.periodogram=my.periodogram$spec

  # remove values of frequency 0 as it does not contain information about the variance
  my.freq=my.freq[-1]
  my.periodogram=my.periodogram[-1]

  # implement cut-off
  nfreq=length(my.freq)
  my.freq=my.freq[1:floor(nfreq*keep)]
  my.periodogram=my.periodogram[1:floor(nfreq*keep)]

  # Step 2: use GLM with BFGS optimization

  # create 1, x, x^2, ..., x^d
  x.mat=rep(1,length(my.freq))
  for(col.iter in 1:d){
    x.mat=cbind(x.mat,my.freq^col.iter)
  }

  # b0 = rnorm(d + 1)

  # fit the glmnet_exp model
  res = glmnet_exp(x.mat, my.periodogram, alpha = alpha)

  # Step 3: return the estimated variance
  return(exp(-res[1])/N)
}

```

```

# set random seed
# set.seed(20180828)

# number of observations for each simulation run
nobs = 50

# number of simulation runs
n.mcsim = 2000

# mean of the data
mu = 0

# sd of the data
StdDev = 1

# generate data matrix
test.mat = matrix(rnorm(nobs * n.mcsim, mean = mu, sd = StdDev), nrow = nobs)

# compute se using glmnet_exp
res.glmnet_exp = apply(test.mat, 2,
  function(x) {
    data.IF = mu.IF(x)
    sqrt(SE.glmnet_exp(data.IF))
  }
)

```

```

# comparison of results
data.frame(theoretical.se =StdDev/sqrt(nobs), # theoretical SE of mean is StdDev / sqrt(nobs)
  mc.se = sd(apply(test.mat, 2, mean)), # SE from Monte Carlo assuming true params are known, which is
  glmnet_exp.se = mean(res.glmnet_exp) # SE estimated by glmnet_exp
)

## theoretical.se    mc.se glmnet_exp.se
## 1      0.1414214 0.145348    0.1461552

```

## 2.4 Estimating the standard error of sample sharpe ratio using RPEGLMEN

```

rm(list=ls())
library(RPEGLMEN)
library(nse)

#' Compute the periodogram as defined in H&W 1981
#'
#' @param data Vector of data
#' @param max.freq Maximum frequency to be computed
#'
#' @return list of frequencies and corresponding periodograms
#' @export
#'
#' @examples
#' myperiodogram(rnorm(10))
myperiodogram=function(data,max.freq=0.5){
  ## data.fft=myfft(data) This is very slow
  data.fft=fft(data)
  N=length(data)
  # tmp=1:N
  # inset=tmp[(1:N)<floor(N/2)]
  tmp = Mod(data.fft[1:floor(N/2)])^2/N
}

```

```

tmp = sapply(tmp, function(x) max(0.01,x))
return(list(spec=tmp,freq=((0:(floor(N/2)-1))/N)))
}

#' Influence Function of Mean
#'
#' @param data Vector of the data
#' @param ... other parameters
#'
#' @return IF of Mean
#' @export
#' @author Xin Chen, \email{chenx26@uw.edu}
#'
#' @examples
#' mu.IF(rnorm(10))
mu.IF=function(data,...){
  mu.hat=mean(data)
  return(data-mu.hat)
}

#' Compute sample Sharpe Ratio from data
#'
#' @param data Vector of data
#' @param rf Risk-free rate
#'
#' @return Value of Sample Sharpe Ratio
#' @export
#' @author Xin Chen, \email{chenx26@uw.edu}
#'
#' @examples
#' SR(rnorm(10))
SR=function(data,rf=0){
  mu.hat=mean(data)
  sigma.hat=sd(data)
  return((mu.hat-rf)/sigma.hat)
}

#' Compute the Influence Function value of Sharpe Ratio
#'
#' @param data vector of data
#' @param rf risk free interest rate
#' @param ... other parameters
#'
#' @return vector of influence function
#' @export
#' @author Xin Chen, \email{chenx26@uw.edu}
#'
#' @examples
#' SR.IF(rnorm(10))
SR.IF=function(data, ..., rf=0){
  mu.hat=mean(data)
  sd.hat=sd(data)
  IF=1/sd.hat*(data-mu.hat)-1/2*mu.hat/sd.hat^3*((data-mu.hat)^2-sd.hat^2)
  return(IF)
}

#' Compute the standard error of the mean of the data using glmnet_exp
#'
#' @param data vector of data
#' @param d degree of polynomial
#' @param alpha weight for Elastic Net regularizer

```

```

#' @param keep what portion of sample spectral density to use for model fitting
#'
#' @return variance of the mean of the data
#' @export
#'

SE.glmnet_exp=function(data, d=7, alpha=0.5, keep=1){

  N=length(data)
  # Step 1: compute the periodograms
  my.periodogram=myperiodogram(data)
  my.freq=my.periodogram$freq
  my.periodogram=my.periodogram$spec

  # remove values of frequency 0 as it does not contain information about the variance
  my.freq=my.freq[-1]
  my.periodogram=my.periodogram[-1]

  # implement cut-off
  nfreq=length(my.freq)
  my.freq=my.freq[1:floor(nfreq*keep)]
  my.periodogram=my.periodogram[1:floor(nfreq*keep)]

  # Step 2: use GLM with BFGS optimization

  # create 1, x, x^2, ..., x^d
  x.mat=rep(1,length(my.freq))
  for(col.iter in 1:d){
    x.mat=cbind(x.mat,my.freq^col.iter)
  }

  # b0 = rnorm(d + 1)

  # fit the glmnet_exp model
  res = glmnet_exp(x.mat, my.periodogram, alpha = alpha)

  # Step 3: return the estimated variance
  return(exp(-res[1])/N)
}

# set random seed
set.seed(20180828)

# number of observations for each simulation run
nobs = 50

# number of simulation runs
n.mcsim = 2000

# mean of the data
mu = 0.05

# sd of the data
StdDev = 1

# generate data matrix
test.mat = matrix(rnorm(nobs * n.mcsim, mean = mu, sd = StdDev), nrow = nobs)

# compute se using glmnet_exp
res.glmnet_exp = apply(test.mat, 2,
  function(x) {

```

```

        data.IF = SR.IF(x)
        sqrt(SE.glmnet_exp(data.IF))
    }
)

# compute se using nse
res.nse = apply(test.mat, 2,
    function(x) {
        data.IF = SR.IF(x)
        nse.nw(data.IF)
    }
)

# comparison of results
data.frame(theoretical.se = sqrt((1 + 0.5 * (mu / StdDev)^2)/nobs), # theoretical SE of mean is StdDev / sqrt(n)
    mc.se = sd(apply(test.mat, 2, mean)), # SE from Monte Carlo assuming true params are known, which is
    glmnet_exp.se = mean(res.glmnet_exp), # SE estimated by glmnet_exp
    nse.se = mean(res.nse)# SE estimated by nse
)

##   theoretical.se    mc.se glmnet_exp.se   nse.se
## 1      0.1415097 0.1403275    0.1442201 0.129672

```