

Package ‘ADMMnet’

December 12, 2015

Type Package

Title Regularized Model with Selecting the Number of Non-Zeros

Version 0.1

Date 2015-12-10

Author Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang

Maintainer Xiang Li <x12473@cumc.columbia.edu>

Description Fit linear and cox models regularized with net (L1 and Laplacian), elastic-net (L1 and L2) or lasso (L1) penalty, and their adaptive forms, such as adaptive lasso and net adjusting for signs of linked coefficients. In addition, it treats the number of non-zero coefficients as another tuning parameter and simultaneously selects with the regularization parameter. The package uses one-step coordinate descent algorithm and runs extremely fast by taking into account the sparsity structure of coefficients.

License GPL (>= 2)

Imports Rcpp (>= 0.12.2)

LinkingTo Rcpp, RcppEigen

Depends Matrix (>= 1.2-3)

NeedsCompilation yes

Repository CRAN

Date/Publication 2015-12-12 09:41:54

R topics documented:

| | |
|---------------------------|---|
| ADMMnet-package | 2 |
| ADMMnet | 3 |
| print.ADMMnet | 7 |

| | |
|--------------|----------|
| Index | 8 |
|--------------|----------|

Description

This package fits linear and cox models regularized with net (L1 and Laplacian), elastic-net (L1 and L2) or lasso (L1) penalty, and their adaptive forms, such as adaptive lasso and net adjusting for signs of linked coefficients. In addition, it treats the number of non-zero coefficients as another tuning parameter and simultaneously selects with the regularization parameter λ . This is motivated by formulating L0 variable selection in ADMM form. By selecting the regularization parameter and the number of non-zeros, it shows significant improvement over the commonly used regularized methods, which depend on the regularization parameter only.

The package uses one-step coordinate descent algorithm and runs extremely fast by taking into account the sparsity structure of coefficients.

Details

Package: ADMMnet
Type: Package
Version: 0.1
Date: 2015-12-10
License: GPL (>= 2)

Functions: `ADMMnet`, `print.ADMMnet`

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
Maintainer: Xiang Li <xl2473@cumc.columbia.edu>

References

- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. *Foundations and Trends in Machine Learning*, 3(1), 1-122.
<http://dl.acm.org/citation.cfm?id=2185816>
- Friedman, J., Hastie, T. and Tibshirani, R. (2010). *Regularization paths for generalized linear models via coordinate descent*, *Journal of Statistical Software*, Vol. 33(1), 1.
<http://www.jstatsoft.org/v33/i01/>
- Li, C., and Li, H. (2010). *Variable selection and regression analysis for graph-structured covariates with an application to genomics*. *The annals of applied statistics*, 4(3), 1498.
<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3423227/>
- Sun, H., Lin, W., Feng, R., and Li, H. (2014) *Network-regularized high-dimensional cox regression for analysis of genomic data*, *Statistica Sinica*.

<http://www3.stat.sinica.edu.tw/statistica/j24n3/j24n319/j24n319.html>

Examples

```
### Linear model ###
set.seed(1213)
N=100;p=30;p1=5
x=matrix(rnorm(N*p),N,p)
beta=rnorm(p1)
xb=x[,1:p1]
y=rnorm(N,xb)

fiti=ADMMnet(x,y,penalty="Lasso",nlambda=10,nfolds=10) # Lasso
# attributes(fiti)

### Cox model ###
set.seed(1213)
N=100;p=30;p1=5
x=matrix(rnorm(N*p),N,p)
beta=rnorm(p1)
xb=x[,1:p1]
ty=rexp(N,exp(xb))
tcens=rbinom(n=N,prob=.3,size=1) # censoring indicator
y=cbind(time=ty,status=1-tcens)

fiti=ADMMnet(x,y,family="cox",penalty="Lasso",nlambda=10,nfolds=10) # Lasso
# attributes(fiti)
```

ADMMnet

Fit a Model with Various Regularization Forms

Description

Fit a linear or cox model regularized with net (L1 and Laplacian), elastic-net (L1 and L2) or lasso (L1) penalty, and their adaptive forms, such as adaptive lasso and net adjusting for signs of linked coefficients. In addition, it treats the number of non-zero coefficients as another tuning parameter and simultaneously selects with the regularization parameter lambda. The package uses one-step coordinate descent algorithm and runs extremely fast by taking into account the sparsity structure of coefficients.

Usage

```
ADMMnet(x, y, family = c("gaussian", "cox"), penalty = c("Lasso", "Enet", "Net"),
  Omega = NULL, alpha = 1.0, lambda = NULL, nlambda = 50, rlambda = NULL,
  nfolds = 1, foldid = NULL, inzero = TRUE, adaptive = c(FALSE, TRUE), aini = NULL,
  isd = FALSE, keep.beta = FALSE, ifast = TRUE, thresh = 1e-07, maxit = 1e+05)
```

Arguments

| | |
|----------------------|---|
| <code>x</code> | input matrix. Each row is an observation vector. |
| <code>y</code> | response variable. For <code>family = "gaussian"</code> , <code>y</code> is a continuous vector. For <code>family = "cox"</code> , <code>y</code> is a two-column matrix with columns named 'time' and 'status'. 'status' is a binary variable, with '1' indicating event, and '0' indicating right censored. |
| <code>family</code> | type of outcome. Can be "gaussian" or "cox". |
| <code>penalty</code> | penalty type. Can choose "Net", "Enet" (elastic net) and "Lasso". For "Net", need to specify <code>Omega</code> ; otherwise, "Enet" is performed. For <code>penalty = "Net"</code> , the penalty is defined as |

$$\lambda * \alpha * \|\beta\|_1 + (1 - \alpha)/2 * (\beta^T L \beta),$$

where L is a Laplacian matrix calculated from `Omega`.

| | |
|-----------------------|---|
| <code>Omega</code> | correlation/adjacency matrix with zero diagonal, used for <code>penalty = "Net"</code> to calculate Laplacian matrix. |
| <code>alpha</code> | ratio between L1 and Laplacian for "Net", or between L1 and L2 for "Enet". Default is <code>alpha = 1.0</code> , i.e. lasso. |
| <code>lambda</code> | a user supplied decreasing sequence. If <code>lambda = NULL</code> , a sequence of <code>lambda</code> is generated based on <code>nlambda</code> and <code>rlambda</code> . Supplying a value of <code>lambda</code> overrides this. |
| <code>nlambda</code> | number of <code>lambda</code> values. Default is 50. |
| <code>rlambda</code> | fraction of <code>lambda.max</code> to determine the smallest value for <code>lambda</code> . The default is <code>rlambda = 0.0001</code> when the number of observations is larger than or equal to the number of variables; otherwise, <code>rlambda = 0.01</code> . |
| <code>nfolds</code> | number of folds. With <code>nfolds = 1</code> and <code>foldid = NULL</code> by default, cross-validation is not performed. For cross-validation, smallest value allowable is <code>nfolds = 3</code> . Specifying <code>foldid</code> overrides <code>nfolds</code> . |
| <code>foldid</code> | an optional vector of values between 1 and <code>nfolds</code> specifying which fold each observation is in. |
| <code>inzero</code> | logical flag for simultaneously selecting the number of non-zero coefficients with <code>lambda</code> . Default is <code>inzero = TRUE</code> . |
| <code>adaptive</code> | logical flags for adaptive version. Default is <code>adaptive = c(FALSE, TRUE)</code> . The first element is for adaptive on β in L1 and the second for adjusting for signs of linked coefficients in Laplacian matrix. |
| <code>aini</code> | a user supplied initial estimate of β . It is a list including <code>wbeta</code> for adaptive L1 and <code>sgn</code> for adjusting Laplacian matrix. <code>wbeta</code> is the absolute value of inverse initial estimates. If <code>aini = NULL</code> but <code>adaptive</code> is required, <code>aini</code> is generated from regularized model with <code>penalty = "Enet"</code> and <code>alpha = 0.0</code> , i.e. a ridge regression. |
| <code>isd</code> | logical flag for outputting standardized coefficients. <code>x</code> is always standardized prior to fitting the model. Default is <code>isd = FALSE</code> , returning β on the original scale. |

| | |
|-----------|--|
| keep.beta | logical flag for returning estimates for all lambda values. For keep.beta = FALSE, only return the estimate with the minimum cross-validation value. |
| ifast | logical flag for efficient calculation of risk set updates for family = "cox". Default is ifast = TRUE. |
| thresh | convergence threshold for coordinate descent. Default value is 1E-7. |
| maxit | Maximum number of iterations for coordinate descent. Default is 10^5. |

Details

One-step coordinate descent algorithm is applied for each lambda. For family = "cox", ifast = TRUE adopts an efficient way to update risk set and sometimes the algorithm ends before all nlambda values of lambda have been evaluated. To evaluate small values of lambda, use ifast = FALSE. The two methods only affect the efficiency of algorithm, not the estimates.

x is always standardized prior to fitting the model and the estimate is returned on the original scale. For family = "gaussian", y is centered by removing its mean, so there is no intercept output.

Cross-validation is used for tuning parameters. For inzero = TRUE, we further select the number of non-zero coefficients obtained from regularized model at each lambda. This is motivated by formulating L0 variable selection in ADMM form, which shows significant improvement over the commonly used regularized methods without this technique.

Value

An object with S3 class "ADMMnet".

| | |
|------------|--|
| Beta | a sparse Matrix of coefficients, stored in class "dgCMatrix". |
| Beta0 | coefficients after additionally tuning the number of non-zeros, for inzero = TRUE. |
| fit | a data.frame containing lambda and the number of non-zero coefficients nzero. With cross-validation, additional results are reported, such as average cross-validation partial likelihood cvm and its standard error cvse, and index with '*' indicating the minimum cvm. For family = "gaussian", rsq is also reported. |
| fit0 | a data.frame containing lambda, cvm and nzero based on inzero = TRUE. |
| lambda.min | value of lambda that gives minimum cvm. |
| lambda.opt | value of lambda based on inzero = TRUE. |
| penalty | penalty type. |
| adaptive | logical flags for adaptive version (see above). |
| flag | convergence flag (for internal debugging). flag = 0 means converged. |

Warning

It may terminate and return NULL.

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
 Maintainer: Xiang Li <xl2473@cumc.columbia.edu>, Shanghong Xie <sx2168@cumc.columbia.edu>

References

Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. *Foundations and Trends in Machine Learning*, 3(1), 1-122.

<http://dl.acm.org/citation.cfm?id=2185816>

Friedman, J., Hastie, T. and Tibshirani, R. (2010). *Regularization paths for generalized linear models via coordinate descent*, *Journal of Statistical Software*, Vol. 33(1), 1.

<http://www.jstatsoft.org/v33/i01/>

Li, C., and Li, H. (2010). *Variable selection and regression analysis for graph-structured covariates with an application to genomics*. *The annals of applied statistics*, 4(3), 1498.

<http://www.ncbi.nlm.nih.gov/pmc/articles/PMC3423227/>

Sun, H., Lin, W., Feng, R., and Li, H. (2014) *Network-regularized high-dimensional cox regression for analysis of genomic data*, *Statistica Sinica*.

<http://www3.stat.sinica.edu.tw/statistica/j24n3/j24n319/j24n319.html>

See Also

[print.ADMMnet](#)

Examples

```
### Linear model ###
set.seed(1213)
N=100;p=30;p1=5
x=matrix(rnorm(N*p),N,p)
beta=rnorm(p1)
xb=x[,1:p1]
y=rnorm(N,xb)

fiti=ADMMnet(x,y,penalty="Lasso",nlambda=10,nfolds=10) # Lasso
# attributes(fiti)

### Cox model ###
set.seed(1213)
N=100;p=30;p1=5
x=matrix(rnorm(N*p),N,p)
beta=rnorm(p1)
xb=x[,1:p1]
ty=rexp(N,exp(xb))
tcens=rbinom(n=N,prob=.3,size=1) # censoring indicator
y=cbind(time=ty,status=1-tcens)

fiti=ADMMnet(x,y,family="cox",penalty="Lasso",nlambda=10,nfolds=10) # Lasso
# attributes(fiti)
```

print.ADMMnet *Print a ADMMnet Object*

Description

Print a summary of results along the path of lambda.

Usage

```
## S3 method for class 'ADMMnet'  
print(x, digits = 4, ...)
```

Arguments

| | |
|--------|--------------------------------|
| x | fitted ADMMnet object |
| digits | significant digits in printout |
| ... | additional print arguments |

Details

The performed model is printed, followed by `fit` and `fit0` (if any) from a fitted ADMMnet object.

Value

The data frame above is silently returned

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
Maintainer: Xiang Li <xl2473@cumc.columbia.edu>

See Also

[ADMMnet](#)

Examples

```
### Linear model ###  
set.seed(1213)  
N=100;p=30;p1=5  
x=matrix(rnorm(N*p),N,p)  
beta=rnorm(p1)  
xb=x[,1:p1]  
y=rnorm(N,xb)  
  
fiti=ADMMnet(x,y,penalty="Lasso",nlambda=10,nfolds=10) # Lasso  
fiti
```

Index

*Topic **Number of non-zeros**

ADMMnet, [3](#)

ADMMnet-package, [2](#)

*Topic **Package**

ADMMnet-package, [2](#)

*Topic **Print**

print.ADMMnet, [7](#)

*Topic **Regularization**

ADMMnet, [3](#)

ADMMnet-package, [2](#)

ADMMnet, [2](#), [3](#), [7](#)

ADMMnet-package, [2](#)

print.ADMMnet, [2](#), [6](#), [7](#)