

Package ‘biglasso’

January 31, 2021

Version 1.4.1

Date 2021-01-29

Title Extending Lasso Model Fitting to Big Data

Author Yaohui Zeng [aut,cre], Chuyi Wang [aut,cre], Patrick Breheny [ctb]

Maintainer Chuyi Wang <wwaa0208@gmail.com>

Description Extend lasso and elastic-net model fitting for ultrahigh-dimensional, multi-gigabyte data sets that cannot be loaded into memory. It's much more memory- and computation-efficient as compared to existing lasso-fitting packages like 'glmnet' and 'ncvreg', thus allowing for very powerful big data analysis even with an ordinary laptop.

License GPL-3

URL <https://github.com/YaohuiZeng/biglasso>,
<https://arxiv.org/abs/1701.05936>

BugReports <https://github.com/YaohuiZeng/biglasso/issues>

Depends R (>= 3.2.0), bigmemory (>= 4.5.0), Matrix, ncvreg

Imports Rcpp (>= 0.12.1), methods

LinkingTo Rcpp, RcppArmadillo (>= 0.8.600), bigmemory, BH

VignetteBuilder knitr

Suggests parallel, testthat, R.rsp, glmnet, survival, knitr, rmarkdown

RoxygenNote 7.1.1

NeedsCompilation yes

Repository CRAN

Date/Publication 2021-01-31 15:30:06 UTC

R topics documented:

biglasso-package	2
biglasso	5
colon	9

cv.biglasso	10
plot.biglasso	12
plot.cv.biglasso	13
predict.biglasso	14
predict.cv.biglasso	16
setupX	17
summary.cv.biglasso	19

Index	21
--------------	-----------

biglasso-package	<i>Extending Lasso Model Fitting to Big Data</i>
------------------	--

Description

Extend lasso and elastic-net linear, logistic and cox regression models for ultrahigh-dimensional, multi-gigabyte data sets that cannot be loaded into available RAM. This package utilizes memory-mapped files to store the massive data on the disk and only read those into memory whenever necessary during model fitting. Moreover, some advanced feature screening rules are proposed and implemented to accelerate the model fitting. As a result, this package is much more memory- and computation-efficient and highly scalable as compared to existing lasso-fitting packages such as [glmnet](#) and [ncvreg](#), thus allowing for powerful big data analysis even with only an ordinary laptop.

Details

Package: biglasso
 Type: Package
 Version: 1.4-1
 Date: 2021-01-29
 License: GPL-3

Penalized regression models, in particular the lasso, have been extensively applied to analyzing high-dimensional data sets. However, due to the memory limit, existing R packages are not capable of fitting lasso models for ultrahigh-dimensional, multi-gigabyte data sets which have been increasingly seen in many areas such as genetics, biomedical imaging, genome sequencing and high-frequency finance.

This package aims to fill the gap by extending lasso model fitting to Big Data in R. Version $\geq 1.2-3$ represents a major redesign where the source code is converted into C++ (previously in C), and new feature screening rules, as well as OpenMP parallel computing, are implemented. Some key features of biglasso are summarized as below:

1. it utilizes memory-mapped files to store the massive data on the disk, only loading data into memory when necessary during model fitting. Consequently, it's able to seamlessly data-larger-than-RAM cases.
2. it is built upon pathwise coordinate descent algorithm with warm start, active set cycling, and feature screening strategies, which has been proven to be one of fastest lasso solvers.

3. in incorporates our newly developed hybrid and adaptive screening that outperform state-of-the-art screening rules such as the sequential strong rule (SSR) and the sequential EDPP rule (SEDPP) with additional 1.5x to 4x speedup.
4. the implementation is designed to be as memory-efficient as possible by eliminating extra copies of the data created by other R packages, making it at least 2x more memory-efficient than `glmnet`.
5. the underlying computation is implemented in C++, and parallel computing with OpenMP is also supported.

For more information:

- Benchmarking results: <https://github.com/YaohuiZeng/biglasso>.
- Tutorial: <http://yaohuizeng.github.io/biglasso/articles/biglasso.html>
- Technical paper: <https://arxiv.org/abs/1701.05936>

Note

The input design matrix X must be a `big.matrix` object. This can be created by the function `as.big.matrix` in the R package `bigmemory`. If the data (design matrix) is very large (e.g. 10 GB) and stored in an external file, which is often the case for big data, X can be created by calling the function `setupX`. **In this case, there are several restrictions about the data file:**

1. the data file must be a well-formatted ASCII-file, with each row corresponding to an observation and each column a variable;
2. the data file must contain only one single type. Current version only supports double type;
3. the data file must contain only numeric variables. If there are categorical variables, the user needs to create dummy variables for each categorical variable (by adding additional columns).

Future versions will try to address these restrictions.

Denote the number of observations and variables be, respectively, n and p . It's worth noting that the package is more suitable for wide data (ultrahigh-dimensional, $p \gg n$) as compared to long data ($n \gg p$). This is because the model fitting algorithm takes advantage of sparsity assumption of high-dimensional data. To just give the user some ideas, below are some benchmarking results of the total computing time (in seconds) for solving lasso-penalized linear regression along a sequence of 100 values of the tuning parameter. In all cases, assume 20 non-zero coefficients equal ± 2 in the true model. (Based on Version 1.2-3, screening rule "SSR-BEDPP" is used)

- For wide data case ($p > n$), $n = 1,000$:

p	1,000	10,000	100,000	1,000,000
Size of X	9.5 MB	95 MB	950 MB	9.5 GB
Elapsed time (s)	0.11	0.83	8.47	85.50

Author(s)

Yaohui Zeng, Chuyi Wang and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com> and Chuyi Wang <wwaa0208@gmail.com>

References

- Zeng, Y., and Breheny, P. (2017). The biglasso Package: A Memory- and Computation-Efficient Solver for Lasso Model Fitting with Big Data in R. <https://arxiv.org/abs/1701.05936>.
- 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. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **74**(2), 245-266.
- Wang, J., Zhou, J., Wonka, P., and Ye, J. (2013). Lasso screening rules via dual polytope projection. *In Advances in Neural Information Processing Systems*, pp. 1070-1078.
- Xiang, Z. J., and Ramadge, P. J. (2012). Fast lasso screening tests based on correlations. *In Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on* (pp. 2137-2140). IEEE.
- Wang, J., Zhou, J., Liu, J., Wonka, P., and Ye, J. (2014). A safe screening rule for sparse logistic regression. *In Advances in Neural Information Processing Systems*, pp. 1053-1061.

Examples

```
## Not run:
## Example of reading data from external big data file, fit lasso model,
## and run cross validation in parallel

# simulated design matrix, 1000 observations, 500,000 variables, ~ 5GB
# there are 10 true variables with non-zero coefficient 2.
xfname <- 'x_e3_5e5.txt'
yfname <- 'y_e3_5e5.txt' # response vector
time <- system.time(
  X <- setupX(xfname, sep = '\t') # create backing files (.bin, .desc)
)
print(time) # ~ 7 minutes; this is just one-time operation
dim(X)

# the big.matrix then can be retrieved by its descriptor file (.desc) in any new R session.
rm(X)
xdesc <- 'x_e3_5e5.desc'
X <- attach.big.matrix(xdesc)
dim(X)

y <- as.matrix(read.table(yfname, header = F))
time.fit <- system.time(
  fit <- biglasso(X, y, family = 'gaussian', screen = 'Hybrid')
)
print(time.fit) # ~ 44 seconds for fitting a lasso model along the entire solution path

# cross validation in parallel
seed <- 1234
time.cvfit <- system.time(
  cvfit <- cv.biglasso(X, y, family = 'gaussian', screen = 'Hybrid',
    seed = seed, ncores = 4, nfolds = 10)
)
print(time.cvfit) # ~ 3 minutes for 10-fold cross validation
```

```

plot(cvfit)
summary(cvfit)

## End(Not run)

```

biglasso

Fit lasso penalized regression path for big data

Description

Extend lasso model fitting to big data that cannot be loaded into memory. Fit solution paths for linear or logistic regression models penalized by lasso, ridge, or elastic-net over a grid of values for the regularization parameter lambda.

Usage

```

biglasso(
  X,
  y,
  row.idx = 1:nrow(X),
  penalty = c("lasso", "ridge", "enet"),
  family = c("gaussian", "binomial", "cox"),
  alg.logistic = c("Newton", "MM"),
  screen = c("Adaptive", "SSR", "Hybrid", "None"),
  safe.thresh = 0,
  update.thresh = 1,
  ncores = 1,
  alpha = 1,
  lambda.min = ifelse(nrow(X) > ncol(X), 0.001, 0.05),
  nlambdas = 100,
  lambda.log.scale = TRUE,
  lambda,
  eps = 1e-07,
  max.iter = 1000,
  dfmax = ncol(X) + 1,
  penalty.factor = rep(1, ncol(X)),
  warn = TRUE,
  output.time = FALSE,
  return.time = TRUE,
  verbose = FALSE
)

```

Arguments

X The design matrix, without an intercept. It must be a `big.matrix` object. The function standardizes the data and includes an intercept internally by default during the model fitting.

y	The response vector for family="gaussian" or family="binomial". For family="cox", y should be a two-column matrix with columns 'time' and 'status'. The latter is a binary variable, with '1' indicating death, and '0' indicating right censored.
row.idx	The integer vector of row indices of X that used for fitting the model. 1:nrow(X) by default.
penalty	The penalty to be applied to the model. Either "lasso" (the default), "ridge", or "enet" (elastic net).
family	Either "gaussian", "binomial" or "cox", depending on the response.
alg.logistic	The algorithm used in logistic regression. If "Newton" then the exact hessian is used (default); if "MM" then a majorization-minimization algorithm is used to set an upper-bound on the hessian matrix. This can be faster, particularly in data-larger-than-RAM case.
screen	The feature screening rule used at each lambda that discards features to speed up computation: "SSR" (default if penalty="ridge" or penalty="enet") is the sequential strong rule; "Hybrid" is our newly proposed hybrid screening rules which combine the strong rule with a safe rule. "Adaptive" (default for penalty="lasso" without penalty.factor) is our newly proposed adaptive rules which reuse screening reference for multiple lambda values. Note that: (1) for linear regression with elastic net penalty, both "SSR" and "Hybrid" are applicable since version 1.3-0; (2) only "SSR" is applicable to elastic-net-penalized logistic regression or cox regression; (3) active set cycling strategy is incorporated with these screening rules.
safe.thresh	the threshold value between 0 and 1 that controls when to stop safe test. For example, 0.01 means to stop safe test at next lambda iteration if the number of features rejected by safe test at current lambda iteration is not larger than 1% of the total number of features. So 1 means to always turn off safe test, whereas 0 (default) means to turn off safe test if the number of features rejected by safe test is 0 at current lambda.
update.thresh	the non negative threshold value that controls how often to update the reference of safe rules for "Adaptive" methods. Smaller value means updating more often.
ncores	The number of OpenMP threads used for parallel computing.
alpha	The elastic-net mixing parameter that controls the relative contribution from the lasso (l1) and the ridge (l2) penalty. The penalty is defined as $\alpha\ \beta\ _1 + (1 - \alpha)/2\ \beta\ _2^2.$ alpha=1 is the lasso penalty, alpha=0 the ridge penalty, alpha in between 0 and 1 is the elastic-net ("enet") penalty.
lambda.min	The smallest value for lambda, as a fraction of lambda.max. Default is .001 if the number of observations is larger than the number of covariates and .05 otherwise.
nlambda	The number of lambda values. Default is 100.
lambda.log.scale	Whether compute the grid values of lambda on log scale (default) or linear scale.

<code>lambda</code>	A user-specified sequence of lambda values. By default, a sequence of values of length <code>nlambda</code> is computed, equally spaced on the log scale.
<code>eps</code>	Convergence threshold for inner coordinate descent. The algorithm iterates until the maximum change in the objective after any coefficient update is less than <code>eps</code> times the null deviance. Default value is $1e-7$.
<code>max.iter</code>	Maximum number of iterations. Default is 1000.
<code>dfmax</code>	Upper bound for the number of nonzero coefficients. Default is no upper bound. However, for large data sets, computational burden may be heavy for models with a large number of nonzero coefficients.
<code>penalty.factor</code>	A multiplicative factor for the penalty applied to each coefficient. If supplied, <code>penalty.factor</code> must be a numeric vector of length equal to the number of columns of X . The purpose of <code>penalty.factor</code> is to apply differential penalization if some coefficients are thought to be more likely than others to be in the model. Current package doesn't allow unpenalized coefficients. That is <code>penalty.factor</code> cannot be 0. <code>penalty.factor</code> is only supported for "SSR" screen.
<code>warn</code>	Return warning messages for failures to converge and model saturation? Default is TRUE.
<code>output.time</code>	Whether to print out the start and end time of the model fitting. Default is FALSE.
<code>return.time</code>	Whether to return the computing time of the model fitting. Default is TRUE.
<code>verbose</code>	Whether to output the timing of each lambda iteration. Default is FALSE.

Details

The objective function for linear regression (`family = "gaussian"`) is

$$\frac{1}{2n} \text{RSS} + \lambda * \text{penalty},$$

for logistic regression (`family = "binomial"`) it is

$$-\frac{1}{n} \text{loglike} + \lambda * \text{penalty},$$

, for cox regression, breslow approximation for ties is applied.

Several advanced feature screening rules are implemented. For lasso-penalized linear regression, all the options of `screen` are applicable. Our proposal adaptive rule - "Adaptive" - achieves highest speedup so it's the recommended one, especially for ultrahigh-dimensional large-scale data sets. For cox regression and/or the elastic net penalty, only "SSR" is applicable for now. More efficient rules are under development.

Value

An object with S3 class "biglasso" with following variables.

<code>beta</code>	The fitted matrix of coefficients, store in sparse matrix representation. The number of rows is equal to the number of coefficients, whereas the number of columns is equal to <code>nlambda</code> .
-------------------	---

<code>iter</code>	A vector of length <code>nlambda</code> containing the number of iterations until convergence at each value of <code>lambda</code> .
<code>lambda</code>	The sequence of regularization parameter values in the path.
<code>penalty</code>	Same as above.
<code>family</code>	Same as above.
<code>alpha</code>	Same as above.
<code>loss</code>	A vector containing either the residual sum of squares (for "gaussian") or negative log-likelihood (for "binomial") of the fitted model at each value of <code>lambda</code> .
<code>penalty.factor</code>	Same as above.
<code>n</code>	The number of observations used in the model fitting. It's equal to <code>length(row.idx)</code> .
<code>center</code>	The sample mean vector of the variables, i.e., column mean of the sub-matrix of <code>X</code> used for model fitting.
<code>scale</code>	The sample standard deviation of the variables, i.e., column standard deviation of the sub-matrix of <code>X</code> used for model fitting.
<code>y</code>	The response vector used in the model fitting. Depending on <code>row.idx</code> , it could be a subset of the raw input of the response vector <code>y</code> .
<code>screen</code>	Same as above.
<code>col.idx</code>	The indices of features that have 'scale' value greater than $1e-6$. Features with 'scale' less than $1e-6$ are removed from model fitting.
<code>rejections</code>	The number of features rejected at each value of <code>lambda</code> .
<code>safe_rejections</code>	The number of features rejected by safe rules at each value of <code>lambda</code> .

Author(s)

Yaohui Zeng, Chuyi Wang and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com> and Chuyi Wang <wwaa0208@gmail.com>

See Also

[biglasso-package](#), [setupX](#), [cv.biglasso](#), [plot.biglasso](#), [ncvreg](#)

Examples

```
## Linear regression
data(colon)
X <- colon$X
y <- colon$y
X.bm <- as.big.matrix(X)
# lasso, default
par(mfrow=c(1,2))
fit.lasso <- biglasso(X.bm, y, family = 'gaussian')
plot(fit.lasso, log.l = TRUE, main = 'lasso')
```

```

# elastic net
fit.enet <- biglasso(X.bm, y, penalty = 'enet', alpha = 0.5, family = 'gaussian')
plot(fit.enet, log.l = TRUE, main = 'elastic net, alpha = 0.5')

## Logistic regression
data(colon)
X <- colon$X
y <- colon$y
X.bm <- as.big.matrix(X)
# lasso, default
par(mfrow = c(1, 2))
fit.bin.lasso <- biglasso(X.bm, y, penalty = 'lasso', family = "binomial")
plot(fit.bin.lasso, log.l = TRUE, main = 'lasso')
# elastic net
fit.bin.enet <- biglasso(X.bm, y, penalty = 'enet', alpha = 0.5, family = "binomial")
plot(fit.bin.enet, log.l = TRUE, main = 'elastic net, alpha = 0.5')

## Cox regression
set.seed(10101)
N <- 1000; p <- 30; nzc <- p/3
X <- matrix(rnorm(N * p), N, p)
beta <- rnorm(nzc)
fx <- X[, seq(nzc)] %*% beta/3
hx <- exp(fx)
ty <- rexp(N, hx)
tcens <- rbinom(n = N, prob = 0.3, size = 1) # censoring indicator
y <- cbind(time = ty, status = 1 - tcens) # y <- Surv(ty, 1 - tcens) with library(survival)
X.bm <- as.big.matrix(X)
fit <- biglasso(X.bm, y, family = "cox")
plot(fit)

```

colon

Gene expression data from colon-cancer patients

Description

The data file contains gene expression data of 62 samples (40 tumor samples, 22 normal samples) from colon-cancer patients analyzed with an Affymetrix oligonucleotide Hum6000 array.

Format

A list of 2 variables included in colon:

- X: a 62-by-2000 matrix that records the gene expression data. Used as design matrix.
- y: a binary vector of length 62 recording the sample status: 1 = tumor; 0 = normal. Used as response vector.

Source

The raw data can be found on Bioconductor: <https://bioconductor.org/packages/release/data/experiment/html/colonCA.html>.

References

- U. Alon et al. (1999): Broad patterns of gene expression revealed by clustering analysis of tumor and normal colon tissue probed by oligonucleotide arrays. *Proc. Natl. Acad. Sci. USA* **96**, 6745-6750. <https://www.pnas.org/content/96/12/6745.short>.

Examples

```
data(colon)
X <- colon$X
y <- colon$y
str(X)
dim(X)
X.bm <- as.big.matrix(X, backingfile = "") # convert to big.matrix object
str(X.bm)
dim(X.bm)
```

cv.biglasso

Cross-validation for biglasso

Description

Perform k-fold cross validation for penalized regression models over a grid of values for the regularization parameter lambda.

Usage

```
cv.biglasso(
  X,
  y,
  row.idx = 1:nrow(X),
  eval.metric = c("default", "MAPE"),
  ncores = parallel::detectCores(),
  ...,
  nfolds = 5,
  seed,
  cv.ind,
  trace = FALSE
)
```

Arguments

X	The design matrix, without an intercept, as in biglasso .
y	The response vector, as in biglasso .
row.idx	The integer vector of row indices of X that used for fitting the model. as in biglasso .
eval.metric	The evaluation metric for the cross-validated error and for choosing optimal lambda. "default" for linear regression is MSE (mean squared error), for logistic regression is misclassification error. "MAPE", for linear regression only, is the Mean Absolute Percentage Error.
ncores	The number of cores to use for parallel execution across a cluster created by the parallel package. (This is different from ncores in biglasso , which is the number of OpenMP threads.)
...	Additional arguments to biglasso .
nfolds	The number of cross-validation folds. Default is 5.
seed	The seed of the random number generator in order to obtain reproducible results.
cv.ind	Which fold each observation belongs to. By default the observations are randomly assigned by cv.biglasso .
trace	If set to TRUE, cv.biglasso will inform the user of its progress by announcing the beginning of each CV fold. Default is FALSE.

Details

The function calls [biglasso](#) `nfolds` times, each time leaving out $1/nfolds$ of the data. The cross-validation error is based on the residual sum of squares when `family="gaussian"` and the binomial deviance when `family="binomial"`.

The S3 class object `cv.biglasso` inherits class `cv.ncvreg`. So S3 functions such as `"summary"`, `"plot"` can be directly applied to the `cv.biglasso` object.

Value

An object with S3 class `"cv.biglasso"` which inherits from class `"cv.ncvreg"`. The following variables are contained in the class (adopted from [cv.ncvreg](#)).

cve	The error for each value of lambda, averaged across the cross-validation folds.
cvse	The estimated standard error associated with each value of for cve.
lambda	The sequence of regularization parameter values along which the cross-validation error was calculated.
fit	The fitted biglasso object for the whole data.
min	The index of lambda corresponding to <code>lambda.min</code> .
lambda.min	The value of lambda with the minimum cross-validation error.
null.dev	The deviance for the intercept-only model.
pe	If <code>family="binomial"</code> , the cross-validation prediction error for each value of lambda.
cv.ind	Same as above.

Author(s)

Yaohui Zeng and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also[biglasso](#), [plot.cv.biglasso](#), [summary.cv.biglasso](#), [setupX](#)**Examples**

```
## Not run:
## cv.biglasso
data(colon)
X <- colon$X
y <- colon$y
X.bm <- as.big.matrix(X)

## logistic regression
cvfit <- cv.biglasso(X.bm, y, family = 'binomial', seed = 1234, ncores = 2)
par(mfrow = c(2, 2))
plot(cvfit, type = 'all')
summary(cvfit)

## End(Not run)
```

plot.biglasso

Plot coefficients from a "biglasso" object

Description

Produce a plot of the coefficient paths for a fitted [biglasso](#) object.

Usage

```
## S3 method for class 'biglasso'
plot(x, alpha = 1, log.l = TRUE, ...)
```

Arguments

x	Fitted "biglasso" model.
alpha	Controls alpha-blending, helpful when the number of covariates is large. Default is alpha=1.
log.l	Should horizontal axis be on the log scale? Default is TRUE.
...	Other graphical parameters to plot

Author(s)

Yaohui Zeng and Patrick Breheny
 Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also

[biglasso](#), [cv.biglasso](#)

Examples

```
## See examples in "biglasso"
```

plot.cv.biglasso	<i>Plots the cross-validation curve from a "cv.biglasso" object</i>
------------------	---

Description

Plot the cross-validation curve from a [cv.biglasso](#) object, along with standard error bars.

Usage

```
## S3 method for class 'cv.biglasso'
plot(
  x,
  log.l = TRUE,
  type = c("cve", "rsq", "scale", "snr", "pred", "all"),
  selected = TRUE,
  vertical.line = TRUE,
  col = "red",
  ...
)
```

Arguments

x	A "cv.biglasso" object.
log.l	Should horizontal axis be on the log scale? Default is TRUE.
type	What to plot on the vertical axis. cve plots the cross-validation error (deviance); rsq plots an estimate of the fraction of the deviance explained by the model (R-squared); snr plots an estimate of the signal-to-noise ratio; scale plots, for family="gaussian", an estimate of the scale parameter (standard deviation); pred plots, for family="binomial", the estimated prediction error; all produces all of the above.
selected	If TRUE (the default), places an axis on top of the plot denoting the number of variables in the model (i.e., that have a nonzero regression coefficient) at that value of lambda.

vertical.line If TRUE (the default), draws a vertical line at the value where cross-validation error is minimized.

col Controls the color of the dots (CV estimates).

... Other graphical parameters to plot

Details

Error bars representing approximate 68% confidence intervals are plotted along with the estimates at value of lambda. For rsq and snr, these confidence intervals are quite crude, especially near.

Author(s)

Yaohui Zeng and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also

[biglasso](#), [cv.biglasso](#)

Examples

```
## See examples in "cv.biglasso"
```

predict.biglasso *Model predictions based on a fitted biglasso object*

Description

Extract predictions (fitted response, coefficients, etc.) from a fitted [biglasso](#) object.

Usage

```
## S3 method for class 'biglasso'
predict(
  object,
  X,
  row.idx = 1:nrow(X),
  type = c("link", "response", "class", "coefficients", "vars", "nvars"),
  lambda,
  which = 1:length(object$lambda),
  ...
)

## S3 method for class 'biglasso'
coef(object, lambda, which = 1:length(object$lambda), drop = TRUE, ...)
```

Arguments

object	A fitted "biglasso" model object.
X	Matrix of values at which predictions are to be made. It must be a big.matrix object. Not used for type="coefficients".
row.idx	Similar to that in biglasso , it's a vector of the row indices of X that used for the prediction. 1:nrow(X) by default.
type	Type of prediction: "link" returns the linear predictors; "response" gives the fitted values; "class" returns the binomial outcome with the highest probability; "coefficients" returns the coefficients; "vars" returns a list containing the indices and names of the nonzero variables at each value of lambda; "nvars" returns the number of nonzero coefficients at each value of lambda.
lambda	Values of the regularization parameter lambda at which predictions are requested. Linear interpolation is used for values of lambda not in the sequence of lambda values in the fitted models.
which	Indices of the penalty parameter lambda at which predictions are required. By default, all indices are returned. If lambda is specified, this will override which.
...	Not used.
drop	If coefficients for a single value of lambda are to be returned, reduce dimensions to a vector? Setting drop=FALSE returns a 1-column matrix.

Value

The object returned depends on type.

Author(s)

Yaohui Zeng and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also

[biglasso](#), [cv.biglasso](#)

Examples

```
## Logistic regression
data(colon)
X <- colon$X
y <- colon$y
X.bm <- as.big.matrix(X, backingfile = "")
fit <- biglasso(X.bm, y, penalty = 'lasso', family = "binomial")
coef <- coef(fit, lambda=0.05, drop = TRUE)
coef[which(coef != 0)]
predict(fit, X.bm, type="link", lambda=0.05)
predict(fit, X.bm, type="response", lambda=0.05)
predict(fit, X.bm, type="class", lambda=0.1)
predict(fit, type="vars", lambda=c(0.05, 0.1))
predict(fit, type="nvars", lambda=c(0.05, 0.1))
```

predict.cv.biglasso *Model predictions based on a fitted cv.biglasso object*

Description

Extract predictions from a fitted `cv.biglasso` object.

Usage

```
## S3 method for class 'cv.biglasso'
predict(
  object,
  X,
  row.idx = 1:nrow(X),
  type = c("link", "response", "class", "coefficients", "vars", "nvars"),
  lambda = object$lambda.min,
  which = object$min,
  ...
)

## S3 method for class 'cv.biglasso'
coef(object, lambda = object$lambda.min, which = object$min, ...)
```

Arguments

<code>object</code>	A fitted "cv.biglasso" model object.
<code>X</code>	Matrix of values at which predictions are to be made. It must be a <code>big.matrix</code> object. Not used for <code>type="coefficients"</code> .
<code>row.idx</code>	Similar to that in <code>biglasso</code> , it's a vector of the row indices of <code>X</code> that used for the prediction. <code>1:nrow(X)</code> by default.
<code>type</code>	Type of prediction: "link" returns the linear predictors; "response" gives the fitted values; "class" returns the binomial outcome with the highest probability; "coefficients" returns the coefficients; "vars" returns a list containing the indices and names of the nonzero variables at each value of <code>lambda</code> ; "nvars" returns the number of nonzero coefficients at each value of <code>lambda</code> .
<code>lambda</code>	Values of the regularization parameter <code>lambda</code> at which predictions are requested. The default value is the one corresponding to the minimum cross-validation error.
<code>which</code>	Indices of the penalty parameter <code>lambda</code> at which predictions are requested. The default value is the index of <code>lambda</code> corresponding to <code>lambda.min</code> . Note: this is overridden if <code>lambda</code> is specified.
<code>...</code>	Not used.

Value

The object returned depends on `type`.

Author(s)

Yaohui Zeng and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also[biglasso](#), [cv.biglasso](#)**Examples**

```
## Not run:
## predict.cv.biglasso
data(colon)
X <- colon$X
y <- colon$y
X.bm <- as.big.matrix(X, backingfile = "")
fit <- biglasso(X.bm, y, penalty = 'lasso', family = "binomial")
cvfit <- cv.biglasso(X.bm, y, penalty = 'lasso', family = "binomial", seed = 1234, ncores = 2)
coef <- coef(cvfit)
coef[which(coef != 0)]
predict(cvfit, X.bm, type = "response")
predict(cvfit, X.bm, type = "link")
predict(cvfit, X.bm, type = "class")

## End(Not run)
```

 setupX

Set up design matrix X by reading data from big data file

Description

Set up the design matrix X as a `big.matrix` object based on external massive data file stored on disk that cannot be fully loaded into memory. The data file must be a well-formatted ASCII-file, and contains only one single type. Current version only supports double type. Other restrictions about the data file are described in [biglasso-package](#). This function reads the massive data, and creates a `big.matrix` object. By default, the resulting `big.matrix` is file-backed, and can be shared across processors or nodes of a cluster.

Usage

```
setupX(
  filename,
  dir = getwd(),
  sep = ",",
  backingfile = paste0(unlist(strsplit(filename, split = "\\."))[1], ".bin"),
  descriptorfile = paste0(unlist(strsplit(filename, split = "\\."))[1], ".desc"),
  type = "double",
  ...
)
```

Arguments

filename	The name of the data file. For example, "dat.txt".
dir	The directory used to store the binary and descriptor files associated with the <code>big.matrix</code> . The default is current working directory.
sep	The field separator character. For example, "," for comma-delimited files (the default); "\t" for tab-delimited files.
backingfile	The binary file associated with the file-backed <code>big.matrix</code> . By default, its name is the same as <code>filename</code> with the extension replaced by ".bin".
descriptorfile	The descriptor file used for the description of the file-backed <code>big.matrix</code> . By default, its name is the same as <code>filename</code> with the extension replaced by ".desc".
type	The data type. Only "double" is supported for now.
...	Additional arguments that can be passed into function <code>read.big.matrix</code> .

Details

For a data set, this function needs to be called only one time to set up the `big.matrix` object with two backing files (.bin, .desc) created in current working directory. Once set up, the data can be "loaded" into any (new) R session by calling `attach.big.matrix(descriptorfile)`.

This function is a simple wrapper of `read.big.matrix`. See `read.big.matrix` and the package `bigmemory` for more details.

Value

A `big.matrix` object corresponding to a file-backed `big.matrix`. It's ready to be used as the design matrix X in `biglasso` and `cv.biglasso`.

Author(s)

Yaohui Zeng and Patrick Breheny

Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also

`biglasso`, `cv.ncvreg`

Examples

```
## see the example in "biglasso-package"
```

summary.cv.biglasso *Summarizing inferences based on cross-validation*

Description

Summary method for cv.biglasso objects.

Usage

```
## S3 method for class 'cv.biglasso'
summary(object, ...)

## S3 method for class 'summary.cv.biglasso'
print(x, digits, ...)
```

Arguments

object	A cv.biglasso object.
...	Further arguments passed to or from other methods.
x	A "summary.cv.biglasso" object.
digits	Number of digits past the decimal point to print out. Can be a vector specifying different display digits for each of the five non-integer printed values.

Value

summary.cv.biglasso produces an object with S3 class "summary.cv.biglasso" which inherits class "summary.cv.ncvreg". The class has its own print method and contains the following list elements:

penalty	The penalty used by biglasso.
model	Either "linear" or "logistic", depending on the family option in biglasso.
n	Number of observations
p	Number of regression coefficients (not including the intercept).
min	The index of lambda with the smallest cross-validation error.
lambda	The sequence of lambda values used by cv.biglasso.
cve	Cross-validation error (deviance).
r.squared	Proportion of variance explained by the model, as estimated by cross-validation.
snr	Signal to noise ratio, as estimated by cross-validation.
sigma	For linear regression models, the scale parameter estimate.
pe	For logistic regression models, the prediction error (misclassification error).

Author(s)

Yaohui Zeng and Patrick Breheny
 Maintainer: Yaohui Zeng <yaohui.zeng@gmail.com>

See Also

[biglasso](#), [cv.biglasso](#), [plot.cv.biglasso](#)

Examples

```
## See examples in "cv.biglasso" and "biglasso-package"
```

Index

*** datasets**

colon, [9](#)

*** models**

plot.biglasso, [12](#)

plot.cv.biglasso, [13](#)

predict.biglasso, [14](#)

predict.cv.biglasso, [16](#)

summary.cv.biglasso, [19](#)

*** package**

biglasso-package, [2](#)

*** regression**

plot.biglasso, [12](#)

plot.cv.biglasso, [13](#)

predict.biglasso, [14](#)

predict.cv.biglasso, [16](#)

summary.cv.biglasso, [19](#)

big.matrix, [3](#), [5](#), [15](#), [16](#)

biglasso, [5](#), [11–18](#), [20](#)

biglasso-package, [2](#)

coef.biglasso (predict.biglasso), [14](#)

coef.cv.biglasso (predict.cv.biglasso),
[16](#)

colon, [9](#)

cv.biglasso, [8](#), [10](#), [13–18](#), [20](#)

cv.ncvreg, [11](#), [18](#)

ncvreg, [8](#)

plot.biglasso, [8](#), [12](#)

plot.cv.biglasso, [12](#), [13](#), [20](#)

predict.biglasso, [14](#)

predict.cv.biglasso, [16](#)

print.summary.cv.biglasso
(summary.cv.biglasso), [19](#)

read.big.matrix, [18](#)

setupX, [3](#), [8](#), [12](#), [17](#)

summary.cv.biglasso, [12](#), [19](#)