

Package ‘ncvreg’

April 18, 2018

Title Regularization Paths for SCAD and MCP Penalized Regression Models

Version 3.10-0

Date 2018-04-17

Author Patrick Breheny [aut,cre]

Maintainer Patrick Breheny <patrick-breheny@uiowa.edu>

Suggests survival, parallel, knitr

VignetteBuilder knitr

Description Efficient algorithms for fitting regularization paths for linear or logistic regression models penalized by MCP or SCAD, with optional additional L2 penalty.

BugReports <http://github.com/pbreheny/ncvreg/issues>

License GPL-3

URL <http://myweb.uiowa.edu/pbreheny/publications/Breheny2011.pdf>

LazyData TRUE

NeedsCompilation yes

Repository CRAN

Date/Publication 2018-04-17 22:49:41 UTC

R topics documented:

| | |
|--------------------------|----|
| ncvreg-package | 2 |
| AUC.cv.ncvsurv | 3 |
| cv.ncvreg | 4 |
| cv.ncvsurv | 6 |
| fir | 8 |
| Heart | 9 |
| Lung | 10 |
| mfdR | 11 |
| ncvreg | 12 |

| | |
|-----------------------------|-----------|
| ncvsurv | 16 |
| perm.ncvreg | 19 |
| permres | 21 |
| plot.cv.ncvreg | 22 |
| plot.mfdr | 24 |
| plot.ncvreg | 25 |
| plot.ncvsurv.func | 26 |
| predict.ncvreg | 27 |
| predict.ncvsurv | 28 |
| Prostate | 30 |
| std | 31 |
| summary.cv.ncvreg | 32 |
| summary.ncvreg | 33 |
| Index | 36 |

| | |
|----------------|---|
| ncvreg-package | <i>Regularization paths for SCAD- and MCP-penalized regression models</i> |
|----------------|---|

Description

Efficient algorithms for fitting regularization paths for a variety of regression models (linear, logistic, Poisson, survival) penalized by MCP or SCAD, with optional additional L2 penalty.

Details

Accepts a design matrix X and vector of responses y , produces the regularization path over a grid of values for the tuning parameter λ . Also provides methods for plotting, cross-validation-based inference, and for determining locally convex regions of the coefficients paths.

See the "Quick start guide" for a brief overview of how the package works.

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Statist.*, 5: 232-253.

Examples

```
vignette("quick-start", package="ncvreg")
```

| | |
|----------------|--|
| AUC.cv.ncvsurv | <i>Calculates AUC for cv.ncvsurv objects</i> |
|----------------|--|

Description

Calculates the cross-validated AUC (concordance) from a "cv.ncvsurv" object.

Usage

```
## S3 method for class 'cv.ncvsurv'  
AUC(obj, ...)
```

Arguments

| | |
|-----|--|
| obj | A <code>cv.ncvsurv</code> object. You must run <code>cv.ncvsurv</code> with the option <code>returnY=TRUE</code> in order for AUC to work. |
| ... | For S3 method compatibility. |

Details

The area under the curve (AUC), or equivalently, the concordance statistic (C), is calculated according to the procedure outlined in the reference below. This calls the `survConcordance` function in the `survival` package, except the cross-validated linear predictors are used to guard against overfitting. Thus, the values returned by `AUC.cv.ncvsurv` will be lower than those you would obtain with `survConcordance` if you fit the full (unpenalized) model.

Author(s)

Patrick Breheny, Brandon Butcher, and Lawrence Hunsicker

References

van Houwelingen H, Putter H (2011). *Dynamic Prediction in Clinical Survival Analysis*. CRC Press.

See Also

[cv.ncvsurv](#), [survConcordance](#)

Examples

```
data(Lung)  
X <- Lung$X  
y <- Lung$y  
  
cvfit <- cv.ncvsurv(X, y, returnY=TRUE)  
head(AUC(cvfit))  
ll <- log(cvfit$fit$lambda)
```

```
plot(l1, AUC(cvfit), xlim=rev(range(l1)), lwd=3, type='l',
     xlab=expression(log(lambda)), ylab='AUC')
```

 cv.ncvreg

Cross-validation for ncvreg

Description

Performs k-fold cross validation for MCP- or SCAD-penalized regression models over a grid of values for the regularization parameter lambda.

Usage

```
cv.ncvreg(X, y, ..., cluster, nfolds=10, seed, fold, returnY=FALSE,
          trace=FALSE)
```

Arguments

| | |
|---------|--|
| X | The design matrix, without an intercept, as in ncvreg. |
| y | The response vector, as in ncvreg. |
| ... | Additional arguments to ncvreg. |
| cluster | cv.ncvreg can be run in parallel across a cluster using the parallel package. The cluster must be set up in advance using the makeCluster function from that package. The cluster must then be passed to cv.ncvreg (see example). |
| nfolds | The number of cross-validation folds. Default is 10. |
| fold | Which fold each observation belongs to. By default the observations are randomly assigned by cv.ncvreg. |
| seed | You may set the seed of the random number generator in order to obtain reproducible results. |
| returnY | Should cv.ncvreg return the fitted values from the cross-validation folds? Default is FALSE; if TRUE, this will return a matrix in which the element for row i, column j is the fitted value for observation i from the fold in which observation i was excluded from the fit, at the jth value of lambda. |
| trace | If set to TRUE, cv.ncvreg will inform the user of its progress by announcing the beginning of each CV fold. Default is FALSE. |

Details

The function calls ncvreg 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" or family="poisson".

For family="binomial" models, the cross-validation fold assignments are balanced across the 0/1 outcomes, so that each fold has the same proportion of 0/1 outcomes (or as close to the same proportion as it is possible to achieve if cases do not divide evenly).

Value

An object with S3 class "cv.ncvreg" containing:

| | |
|------------|--|
| 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 ncvreg object for the whole data. |
| min | The index of lambda corresponding to lambda.min. |
| lambda.min | The value of lambda with the minimum cross-validation error. |
| null.dev | The deviance for the intercept-only model. |
| Bias | The estimated bias of the minimum cross-validation error, as in Tibshirani RJ and Tibshirani R (2009), "A Bias Correction for the Minimum Error Rate in Cross-Validation", Ann. Appl. Stat. 3:822-829. |
| pe | If family="binomial", the cross-validation prediction error for each value of lambda. |
| Y | If returnY=TRUE, the matrix of cross-validated fitted values (see above). |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>
Grant Brown helped with the parallelization support

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. Ann. Appl. Statist., 5: 232-253.

See Also

[ncvreg](#), [plot.cv.ncvreg](#), [summary.cv.ncvreg](#)

Examples

```
data(Prostate)

cvfit <- cv.ncvreg(Prostate$X, Prostate$y)
plot(cvfit)
summary(cvfit)

fit <- cvfit$fit
plot(fit)
beta <- fit$beta[,cvfit$min]

## requires loading the parallel package
## Not run:
library(parallel)
```

```
X <- Prostate$X
y <- Prostate$y
cl <- makeCluster(4)
cvfit <- cv.ncvreg(X, y, cluster=cl, nfolds=length(y))
## End(Not run)
```

cv.ncvsurv

Cross-validation for ncvsurv

Description

Performs k-fold cross validation for MCP- or SCAD-penalized survival models over a grid of values for the regularization parameter lambda.

Usage

```
cv.ncvsurv(X, y, ..., cluster, nfolds=10, seed, fold, se=c('quick',
'bootstrap'), returnY=FALSE, trace=FALSE)
```

Arguments

| | |
|---------|---|
| X | The design matrix, as in ncvsurv. |
| y | The response matrix, as in ncvsurv. |
| ... | Additional arguments to ncvsurv. |
| cluster | cv.ncvsurv can be run in parallel across a cluster using the parallel package. The cluster must be set up in advance using the makeCluster function from that package. The cluster must then be passed to cv.ncvsurv (see example). |
| nfolds | The number of cross-validation folds. Default is 10. |
| seed | You may set the seed of the random number generator in order to obtain reproducible results. |
| fold | Which fold each observation belongs to. By default the observations are randomly assigned by cv.ncvsurv in a manner that balances (as best as possible) the number of censored observations across the folds. |
| se | Method by which the cross-validation standard error (CVSE) is calculated. The 'quick' approach is based on a rough approximation, but can be calculated more or less instantly. The 'bootstrap' approach is more accurate, but requires additional computing time. |
| returnY | Should cv.ncvsurv return the linear predictors from the cross-validation folds? Default is FALSE; if TRUE, this will return a matrix in which the element for row i, column j is the fitted value for observation i from the fold in which observation i was excluded from the fit, at the jth value of lambda. NOTE: The rows of Y are ordered by time on study, and therefore do not correspond to the original order of observations passed to cv.ncvsurv. |
| trace | If set to TRUE, cv.ncvsurv will inform the user of its progress by announcing the beginning of each CV fold. Default is FALSE. |

Details

The function calls `ncvsurv` `nfolds` times, each time leaving out $1/nfolds$ of the data. Because of the semiparametric nature of Cox regression, cross-validation is not clearly defined. `cv.ncvsurv` uses the approach of calculating the full Cox partial likelihood using the cross-validated set of linear predictors. Unfortunately, using this approach there is no clear way (yet) of determining standard errors, so `cv.ncvsurv`, unlike `cv.ncvreg`, does not provide any.

Other approaches to cross-validation for the Cox regression model have been proposed; the strengths and weaknesses of the various methods for penalized regression in the Cox model are not well understood. Because of this, the approach used by `cv.ncvsurv` may change in the future as additional research is carried out.

Value

An object with S3 class "cv.ncvsurv" inheriting from "cv.ncvreg" and containing:

| | |
|-------------------------|---|
| <code>cve</code> | The error for each value of <code>lambda</code> , averaged across the cross-validation folds. |
| <code>cvse</code> | NULL; see Details. |
| <code>lambda</code> | The sequence of regularization parameter values along which the cross-validation error was calculated. |
| <code>fit</code> | The fitted <code>ncvsurv</code> object for the whole data. |
| <code>min</code> | The index of <code>lambda</code> corresponding to <code>lambda.min</code> . |
| <code>lambda.min</code> | The value of <code>lambda</code> with the minimum cross-validation error. |
| <code>null.dev</code> | The cross-validated deviance for the first model along the grid of <code>lambda</code> (i.e., the cross-validated deviance for <code>max(lambda)</code> , unless you have supplied your own <code>lambda</code> sequence, in which case this quantity is probably not meaningful). Although the actual null deviance can be calculated, it cannot be compared with the cross-validated deviance due to the manner in which deviance must be calculated for Cox regression models (see details). |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

- Breheny P and Huang J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, **5**: 232-253. myweb.uiowa.edu/pbreheny/publications/Breheny2011.pdf
- Verweij PJ and van Houwelingen HC. (1993) Cross-validation in survival analysis. *Statistics in Medicine*, **12**: 2305-2314.

See Also

[ncvsurv](#), [plot.cv.ncvreg](#), [summary.cv.ncvreg](#)

Examples

```
data(Lung)
X <- Lung$X
y <- Lung$y

cvfit <- cv.ncvsurv(X, y)
summary(cvfit)
plot(cvfit)
plot(cvfit, type="rsq")

## requires loading the parallel package
## Not run:
library(parallel)
cl <- makeCluster(4)
cvfit <- cv.ncvsurv(X, y, cluster=cl)
## End(Not run)
```

fir

Marginal false discovery rates

Description

Estimates the marginal false discovery rate (mFDR) of a penalized regression model.

Usage

```
fir(fit, ...)
```

Arguments

| | |
|-----|---|
| fit | An ncvreg or ncvsurv object. |
| ... | Arguments to pass to mfdr . |

Details

This function has been renamed and is currently deprecated. Use [mfdr](#) instead.

Heart

Risk factors associated with heart disease

Description

Data from a subset of the Coronary Risk-Factor Study baseline survey, carried out in rural South Africa.

- X: A design matrix with 462 observations (rows) and 9 predictor variables (columns). The columns are:
 - sbp: Systolic blood pressure
 - tobacco: Cumulative tobacco consumption, in kg
 - ldl: Low-density lipoprotein cholesterol
 - adiposity: Adipose tissue concentration
 - famhist: Family history of heart disease (1=Present, 0=Absent)
 - typea: Score on test designed to measure type-A behavior
 - obesity: Obesity
 - alcohol: Current consumption of alcohol
 - age: Age of subject
- y: Coronary heart disease at baseline; 1=Yes 0=No

Usage

```
data(Heart)
```

Source

<https://web.stanford.edu/~hastie/ElemStatLearn/>

References

- Hastie T, Tibshirani R, and Friedman J. (2001). *The Elements of Statistical Learning*. Springer.
- Rousseauw J, et al. (1983). Coronary risk factor screening in three rural communities. *South African Medical Journal*, **64**, 430-436.

Lung

VA lung cancer data set

Description

Data from a randomised trial of two treatment regimens for lung cancer. This is a standard survival analysis data set from the classic textbook by Kalbfleisch and Prentice.

- X: A design matrix with 137 observations (rows) and 9 predictor variables (columns). The columns are:
 - trt: Treatment indicator (1=control group, 2=treatment group)
 - karno: Karnofsky performance score (0=bad, 100=good)
 - diagtime: Time from diagnosis to randomization (months)
 - age: Age (years)
 - prior: Prior therapy (0=no, 1=yes)
 - squamous: Indicator for whether the cancer type is squamous cell carcinoma (0=no, 1=yes)
 - small: Indicator for whether the cancer type is small cell lung cancer (0=no, 1=yes)
 - adeno: Indicator for whether the cancer type is adenocarcinoma (0=no, 1=yes)
 - large: Indicator for whether the cancer type is large cell carcinoma (0=no, 1=yes)
- y: A two column matrix (Surv object) containing the follow-up time (in days) and an indicator variable for whether the patient died while on the study or not.

Usage

```
data(Lung)
```

Format

A list containing the design matrix X and response matrix y

Source

<https://cran.r-project.org/package=survival>

References

- Kalbfleisch D and Prentice RL (1980), *The Statistical Analysis of Failure Time Data*. Wiley, New York.

See Also

[ncvsurv](#)

mfd *Marginal false discovery rates*

Description

Estimates the marginal false discovery rate (mFDR) of a penalized regression model.

Usage

```
mfd(fit, X)
```

Arguments

| | |
|-----|--|
| fit | An ncvreg or ncvsurv object. |
| X | The model matrix corresponding to fit. This is not necessary for linear regression, but in logistic and Cox regression, the mFDR depends on X. It is not necessary to supply X if it is already contained in fit; i.e., if ncvreg/ncvsurv was run with returnX=TRUE. |

Details

The function estimates the marginal false discovery rate (mFDR) for a penalized regression model. The estimate tends to be accurate in most settings, but will be slightly conservative if predictors are highly correlated. For an alternative way of estimating the mFDR, typically more accurate in highly correlated cases, see [perm.ncvreg](#).

Value

An object with S3 class mfd inheriting from data.frame and containing:

| | |
|------|---|
| EF | The number of variables selected at each value of lambda, averaged over the permutation fits. |
| S | The actual number of selected variables for the non-permuted data. |
| mFDR | The estimated marginal false discovery rate (EF/S). |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>, with contributions from Ryan Miller <ryan-e-miller@uiowa.edu>

See Also

[ncvreg](#), [ncvsurv](#), [plot.mfd](#), [perm.ncvreg](#)

Examples

```
# Linear regression -----
data(Prostate)
fit <- ncvreg(Prostate$X, Prostate$y)

obj <- mfdr(fit)
obj[1:10,]

# Comparison with perm.ncvreg
par(mfrow=c(2,2))
plot(obj)
plot(obj, type="EF")
pmfit <- perm.ncvreg(Prostate$X, Prostate$y)
plot(pmfit)
plot(pmfit, type="EF") # Note that mfdr() is more conservative

# Logistic regression -----
data(Heart)
fit <- ncvreg(Heart$X, Heart$y, family="binomial")
obj <- mfdr(fit)
head(obj)
plot(obj)
plot(obj, type="EF")

# Cox regression -----
data(Lung)
fit <- ncvsvr(Lung$X, Lung$y)
obj <- mfdr(fit)
head(obj)
plot(obj)
plot(obj, type="EF")
```

ncvreg

Fit an MCP- or SCAD-penalized regression path

Description

Fit coefficients paths for MCP- or SCAD-penalized regression models over a grid of values for the regularization parameter lambda. Fits linear and logistic regression models, with option for an additional L2 penalty.

Usage

```
ncvreg(X, y, family=c("gaussian", "binomial", "poisson"),
penalty=c("MCP", "SCAD", "lasso"), gamma=switch(penalty, SCAD=3.7, 3),
alpha=1, lambda.min=ifelse(n>p,.001,.05), nlambda=100, lambda, eps=1e-4,
max.iter=10000, convex=TRUE, dfmax=p+1, penalty.factor=rep(1, ncol(X)),
warn=TRUE, returnX, ...)
```

Arguments

| | |
|-----------------------------|--|
| <code>X</code> | The design matrix, without an intercept. <code>ncvreg</code> standardizes the data and includes an intercept by default. |
| <code>y</code> | The response vector. |
| <code>family</code> | Either "gaussian", "binomial", or "poisson", depending on the response. |
| <code>penalty</code> | The penalty to be applied to the model. Either "MCP" (the default), "SCAD", or "lasso". |
| <code>gamma</code> | The tuning parameter of the MCP/SCAD penalty (see details). Default is 3 for MCP and 3.7 for SCAD. |
| <code>alpha</code> | Tuning parameter for the Mnet estimator which controls the relative contributions from the MCP/SCAD penalty and the ridge, or L2 penalty. <code>alpha=1</code> is equivalent to MCP/SCAD penalty, while <code>alpha=0</code> would be equivalent to ridge regression. However, <code>alpha=0</code> is not supported; <code>alpha</code> may be arbitrarily small, but not exactly 0. |
| <code>lambda.min</code> | The smallest value for <code>lambda</code> , as a fraction of <code>lambda.max</code> . Default is .001 if the number of observations is larger than the number of covariates and .05 otherwise. |
| <code>nlambda</code> | The number of <code>lambda</code> values. Default is 100. |
| <code>lambda</code> | A user-specified sequence of <code>lambda</code> 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. The algorithm iterates until the RMSD for the change in linear predictors for each coefficient is less than <code>eps</code> . Default is $1e-4$. |
| <code>max.iter</code> | Maximum number of iterations (total across entire path). Default is 10000. |
| <code>convex</code> | Calculate index for which objective function ceases to be locally convex? Default is TRUE. |
| <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 <code>X</code> . 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. In particular, <code>penalty.factor</code> can be 0, in which case the coefficient is always in the model without shrinkage. |
| <code>warn</code> | Return warning messages for failures to converge and model saturation? Default is TRUE. |
| <code>returnX</code> | Return the standardized design matrix along with the fit? By default, this option is turned on if <code>X</code> is under 10 MB, but turned off for larger matrices to preserve memory. Note that certain methods, such as summary.ncvreg require access to the design matrix and may not be able to run if <code>returnX=FALSE</code> . |
| <code>...</code> | Not used. |

Details

The sequence of models indexed by the regularization parameter `lambda` is fit using a coordinate descent algorithm. For logistic regression models, some care is taken to avoid model saturation; the algorithm may exit early in this setting. The objective function is defined to be

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

for "gaussian" and

$$-\frac{1}{n} \ell + \text{penalty}$$

for "binomial" or "poisson", where the likelihood is from a traditional generalized linear model assuming the canonical link (logit for "binomial"; log for "poisson").

This algorithm is stable, very efficient, and generally converges quite rapidly to the solution. For GLMs, adaptive rescaling (see reference) is used.

The convexity diagnostics rely on a fine covering of `(lambda.min, lambda.max)`; choosing a low value of `nlambda` may produce unreliable results.

Value

An object with S3 class "ncvreg" containing:

| | |
|-----------------------------|---|
| <code>beta</code> | The fitted matrix of coefficients. The number of rows is equal to the number of coefficients, and 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>gamma</code> | Same as above. |
| <code>alpha</code> | Same as above. |
| <code>convex.min</code> | The last index for which the objective function is locally convex. The smallest value of <code>lambda</code> for which the objective function is convex is therefore <code>lambda[convex.min]</code> , with corresponding coefficients <code>beta[,convex.min]</code> . |
| <code>loss</code> | A vector containing either the residual sum of squares ("gaussian") or negative log-likelihood ("binomial" and "poisson") of the fitted model at each value of <code>lambda</code> . |
| <code>penalty.factor</code> | Same as above. |
| <code>n</code> | Sample size. |

Additionally, if `returnX=TRUE`, the object will also contain

| | |
|----------------|--|
| <code>X</code> | The standardized design matrix. |
| <code>y</code> | The response, centered if <code>family='gaussian'</code> . |

Author(s)

Patrick Breheny <patrick-breheny@uiow.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5: 232-253.

See Also

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

Examples

```
# Linear regression -----
data(Prostate)

par(mfrow=c(2,2))
fit <- ncvreg(Prostate$X, Prostate$y)
plot(fit, main=expression(paste(gamma,"=",3)))
fit <- ncvreg(Prostate$X, Prostate$y, gamma=10)
plot(fit, main=expression(paste(gamma,"=",10)))
fit <- ncvreg(Prostate$X, Prostate$y, gamma=1.5)
plot(fit, main=expression(paste(gamma,"=",1.5)))
fit <- ncvreg(Prostate$X, Prostate$y, penalty="SCAD")
plot(fit, main=expression(paste("SCAD, ",gamma,"=",3)))

par(mfrow=c(2,2))
fit <- ncvreg(Prostate$X, Prostate$y)
plot(fit, main=expression(paste(alpha,"=",1)))
fit <- ncvreg(Prostate$X, Prostate$y, alpha=0.9)
plot(fit, main=expression(paste(alpha,"=",0.9)))
fit <- ncvreg(Prostate$X, Prostate$y, alpha=0.5)
plot(fit, main=expression(paste(alpha,"=",0.5)))
fit <- ncvreg(Prostate$X, Prostate$y, alpha=0.1)
plot(fit, main=expression(paste(alpha,"=",0.1)))

par(mfrow=c(2,2))
fit <- ncvreg(Prostate$X, Prostate$y)
plot(mfdr(fit))          ## Independence approximation
plot(mfdr(fit), type="EF") ## Independence approximation
perm.fit <- perm.ncvreg(Prostate$X, Prostate$y)
plot(perm.fit)
plot(perm.fit, type="EF")

# Logistic regression -----
data(Heart)

par(mfrow=c(2,2))
fit <- ncvreg(Heart$X, Heart$y, family="binomial")
plot(fit, main=expression(paste(gamma,"=",3)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", gamma=10)
```

```

plot(fit, main=expression(paste(gamma,"=",10)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", gamma=1.5)
plot(fit, main=expression(paste(gamma,"=",1.5)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", penalty="SCAD")
plot(fit, main=expression(paste("SCAD, ",gamma,"=",3)))

par(mfrow=c(2,2))
fit <- ncvreg(Heart$X, Heart$y, family="binomial")
plot(fit, main=expression(paste(alpha,"=",1)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", alpha=0.9)
plot(fit, main=expression(paste(alpha,"=",0.9)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", alpha=0.5)
plot(fit, main=expression(paste(alpha,"=",0.5)))
fit <- ncvreg(Heart$X, Heart$y, family="binomial", alpha=0.1)
plot(fit, main=expression(paste(alpha,"=",0.1)))

```

ncvsurv

*Fit an MCP- or SCAD-penalized survival model***Description**

Fit coefficients paths for MCP- or SCAD-penalized Cox regression models over a grid of values for the regularization parameter lambda, with option for an additional L2 penalty.

Usage

```

ncvsurv(X, y, penalty=c("MCP", "SCAD", "lasso"),
gamma=switch(penalty, SCAD=3.7, 3), alpha=1,
lambda.min=ifelse(n>p,.001,.05), nlambdas=100, lambda, eps=1e-4,
max.iter=10000, convex=TRUE, dfmax=p, penalty.factor=rep(1, ncol(X)),
warn=TRUE, returnX, ...)

```

Arguments

| | |
|---------|---|
| X | The design matrix of predictor values. <code>ncvsurv</code> standardizes the data prior to fitting. |
| y | The time-to-event outcome, as a two-column matrix or <code>Surv</code> object. The first column should be time on study (follow up time); the second column should be a binary variable with 1 indicating that the event has occurred and 0 indicating (right) censoring. |
| penalty | The penalty to be applied to the model. Either "MCP" (the default), "SCAD", or "lasso". |
| gamma | The tuning parameter of the MCP/SCAD penalty (see details). Default is 3 for MCP and 3.7 for SCAD. |
| alpha | Tuning parameter for the Mnet estimator which controls the relative contributions from the MCP/SCAD penalty and the ridge, or L2 penalty. $\alpha=1$ is equivalent to MCP/SCAD penalty, while $\alpha=0$ would be equivalent to ridge regression. However, $\alpha=0$ is not supported; α may be arbitrarily small, but not exactly 0. |

| | |
|-----------------------------|---|
| <code>lambda.min</code> | 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. |
| <code>nlambda</code> | The number of lambda values. Default is 100. |
| <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. The algorithm iterates until the RMSD for the change in linear predictors for any coefficient is less than <code>eps</code> . Default is $1e-4$. |
| <code>max.iter</code> | Maximum number of iterations (total across entire path). Default is 1000. |
| <code>convex</code> | Calculate index for which objective function ceases to be locally convex? Default is TRUE. |
| <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. In particular, <code>penalty.factor</code> can be 0, in which case the coefficient is always in the model without any penalization/shrinkage. |
| <code>warn</code> | Return warning messages for failures to converge and model saturation? Default is TRUE. |
| <code>returnX</code> | Return the standardized design matrix along with the fit? By default, this option is turned on if X is under 10 MB, but turned off for larger matrices to preserve memory. Note that certain methods, such as <code>summary.ncvsurv</code> require access to the design matrix and may not be able to run if <code>returnX=FALSE</code> . |
| <code>...</code> | Not used. |

Details

The sequence of models indexed by the regularization parameter `lambda` is fit using a coordinate descent algorithm. In order to accomplish this, the second derivative (Hessian) of the Cox partial log-likelihood is diagonalized (see references for details). The objective function is defined to be

$$-\frac{1}{n}L(\beta|X, y) + \text{penalty},$$

where L is the partial log-likelihood from the Cox regression model.

Presently, ties are not handled by `ncvsurv` in a particularly sophisticated manner. This will be improved upon in a future release of `ncvreg`.

Adaptive rescaling (see references) is used for MCP and SCAD models. The convexity diagnostics rely on a fine covering of $(\text{lambda.min}, \text{lambda.max})$; choosing a low value of `nlambda` may produce unreliable results.

Value

An object with S3 class "ncvsurv" containing:

| | |
|----------------|--|
| beta | The fitted matrix of coefficients. The number of rows is equal to the number of coefficients, and the number of columns is equal to nlambda. |
| iter | A vector of length nlambda containing the number of iterations until convergence at each value of lambda. |
| lambda | The sequence of regularization parameter values in the path. |
| penalty | Same as above. |
| model | Same as above. |
| gamma | Same as above. |
| alpha | Same as above. |
| convex.min | The last index for which the objective function is locally convex. The smallest value of lambda for which the objective function is convex is therefore lambda[convex.min], with corresponding coefficients beta[,convex.min]. |
| loss | The negative partial log-likelihood of the fitted model at each value of lambda. |
| penalty.factor | Same as above. |
| n | The number of observations. |

For Cox models, the following objects are also returned (and are necessary to estimate baseline survival conditional on the estimated regression coefficients), all of which are ordered by time on study. I.e., the *i*th row of *W* does not correspond to the *i*th row of *X*):

| | |
|------|---|
| W | Matrix of exp(beta) values for each subject over all lambda values. |
| time | Times on study. |
| fail | Failure event indicator. |

Additionally, if returnX=TRUE, the object will also contain

| | |
|---|---------------------------------|
| X | The standardized design matrix. |
|---|---------------------------------|

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

- Breheny P and Huang J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Annals of Applied Statistics*, **5**: 232-253. myweb.uiowa.edu/pbreheny/publications/Breheny2011.pdf
- Simon N, Friedman JH, Hastie T, and Tibshirani R. (2011) Regularization Paths for Cox's Proportional Hazards Model via Coordinate Descent. *Journal of Statistical Software*, **39**: 1-13. <http://www.jstatsoft.org/v39/i05>

See Also

[plot.ncvreg](#), [cv.ncvsurv](#)

Examples

```

data(Lung)
X <- Lung$X
y <- Lung$y

par(mfrow=c(2,2))
fit <- ncvsurv(X, y)
plot(fit, main=expression(paste(gamma,"=",3)))
fit <- ncvsurv(X, y, gamma=10)
plot(fit, main=expression(paste(gamma,"=",10)))
fit <- ncvsurv(X, y, gamma=1.5)
plot(fit, main=expression(paste(gamma,"=",1.5)))
fit <- ncvsurv(X, y, penalty="SCAD")
plot(fit, main=expression(paste("SCAD, ",gamma,"=",3)))

fit <- ncvsurv(X,y)
ll <- log(fit$lambda)
par(mfrow=c(2,1))
plot(ll, BIC(fit), type="l", xlim=rev(range(ll)))
lam <- fit$lambda[which.min(BIC(fit))]
b <- coef(fit, lambda=lam)
b[b!=0]
plot(fit)
abline(v=lam)

S <- predict(fit, X, type='survival', lambda=lam)
par(mfrow=c(1,1))
plot(S, xlim=c(0,200))

```

perm.ncvreg

Permutation fitting for ncvreg

Description

Fits multiple penalized regression models in which the outcome is randomly permuted, thereby allowing estimation of the false inclusion rate.

Usage

```
perm.ncvreg(X, y, ..., permute=c("outcome", "residuals"), N=10, seed,
trace=FALSE)
```

Arguments

| | |
|-----|--|
| X | The design matrix, without an intercept, as in ncvreg. |
| y | The response vector, as in ncvreg. |
| ... | Additional arguments to ncvreg. |

| | |
|---------|--|
| permute | What to permute. If 'outcome', the response vector, y, is permuted. If 'residuals', the residuals are permuted. This is only available for linear regression (i.e., for family='gaussian'). Note that permuting the residuals may take a long time, as the residuals differ for each value of lambda, so separate permutations are required at every value of lambda. See also permres . |
| N | The number of permutation replications. Default is 10. |
| seed | You may set the seed of the random number generator in order to obtain reproducible results. |
| trace | If set to TRUE, perm.ncvreg will inform the user of its progress by announcing the beginning of each permutation fit. Default is FALSE. |

Details

The function fits a penalized regression model to the actual data, then repeats the process N times with a permuted version of the response vector. This allows estimation of the expected number of variables included by chance for each value of lambda. The ratio of this expected quantity to the number of selected variables using the actual (non-permuted) response is called the marginal false discovery rate (mFDR).

Value

An object with S3 class "perm.ncvreg" containing:

| | |
|------|--|
| EF | The number of variables selected at each value of lambda, averaged over the permutation fits. |
| S | The actual number of selected variables for the non-permuted data. |
| mFDR | The estimated marginal false discovery rate (EF/S). |
| fit | The fitted ncvreg object for the original (non-permuted) data. |
| loss | The loss/deviance for each value of lambda, averaged over the permutation fits. This is an estimate of the explanatory power of the model under null conditions, and can be used to adjust the loss of the fitted model in a manner akin to the idea of an adjusted R-squared in classical regression. |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

See Also

[ncvreg](#), [plot.mfdr](#), [mfdr](#)

Examples

```
# Linear regression -----
data(Prostate)
pmfit <- perm.ncvreg(Prostate$X, Prostate$y)

par(mfcol=c(2,2))
```

```

plot(pmfit)
plot(pmfit, type="EF")
plot(pmfit$fit)
lam <- pmfit$fit$lambda

pmfit.r <- perm.ncvreg(Prostate$X, Prostate$y, permute='residuals')
plot(pmfit.r, col="red") # Permuting residuals is
lines(lam, pmfit$mFDR, col="gray60") # less conservative

# Logistic regression -----
data(Heart)
pmfit <- perm.ncvreg(Heart$X, Heart$y, family="binomial")

par(mfrow=c(2,2))
plot(pmfit)
plot(pmfit, type="EF")
plot(pmfit$fit)

```

permres

*Permute residuals for a fitted ncvreg model***Description**

Fits multiple penalized regression models in which the residuals are randomly permuted, thereby allowing estimation of the false inclusion rate.

Usage

```

permres(fit, ...)
## S3 method for class 'ncvreg'
permres(fit, lambda, N=10, seed, trace=FALSE, ...)

```

Arguments

| | |
|--------|---|
| fit | A fitted ncvreg model, as produced by <code>ncvreg()</code> . To use with <code>permres</code> , the model must be fit using the <code>returnX=TRUE</code> option. |
| lambda | The regularization parameter to use for estimating residuals. Unlike <code>perm.ncvreg</code> , <code>permres</code> calculates EF and FIR for a specific lambda value, not an entire path. As a result, it runs much faster. |
| N | The number of permutation replications. Default is 10. |
| seed | You may set the seed of the random number generator in order to obtain reproducible results. |
| trace | If set to TRUE, <code>perm.ncvreg</code> will inform the user of its progress by announcing the beginning of each permutation fit. Default is FALSE. |
| ... | Not used. |

Details

The function fits a penalized regression model to the actual data, then repeats the process N times with a permuted version of the response vector. This allows estimation of the expected number of variables included by chance for each value of λ . The ratio of this expected quantity to the number of selected variables using the actual (non-permuted) response is called the false inclusion rate (FIR).

Value

A list with the following components:

| | |
|------|---|
| EF | The number of variables selected at each value of λ , averaged over the permutation fits. |
| S | The actual number of selected variables for the non-permuted data. |
| FIR | The estimated false inclusion rate (EF/S). |
| loss | The loss/deviance, averaged over the permutation fits. This is an estimate of the explanatory power of the model under null conditions, and can be used to adjust the loss of the fitted model in a manner akin to the idea of an adjusted R-squared in classical regression. |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

See Also

[ncvreg](#), [fir](#), [perm.ncvreg](#)

Examples

```
data(Prostate)
fit <- ncvreg(Prostate$X, Prostate$y, N=50)
permres(fit, lambda=0.15)
```

plot.cv.ncvreg

Plots the cross-validation curve from a "cv.ncvreg" object

Description

Plots the cross-validation curve from a "cv.ncvreg" object, along with standard error bars.

Usage

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

Arguments

| | |
|---------------|--|
| x | A "cv.ncvreg" 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 zero, and will hopefully be improved upon in later versions of ncvreg.

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Statist.*, 5: 232-253.

See Also

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

Examples

```
# Linear regression -----
data(Prostate)
cvfit <- cv.ncvreg(Prostate$X, Prostate$y)
plot(cvfit)
par(mfrow=c(2,2))
plot(cvfit, type="all")

# Logistic regression -----
data(Heart)
cvfit <- cv.ncvreg(Heart$X, Heart$y, family="binomial")
```

```
plot(cvfit)
par(mfrow=c(2,2))
plot(cvfit, type="all")
```

plot.mfdr

Plot marginal false discovery rate curves

Description

Plot marginal false discovery rate curves from an "mfdr" or "perm.ncvreg" object.

Usage

```
## S3 method for class 'mfdr'
plot(x, type=c("mFDR", "EF"), log.l=FALSE, selected=TRUE,
      legend=TRUE, ...)
```

Arguments

| | |
|----------|---|
| x | A "perm.ncvreg" or "mfdr" object. |
| type | What to plot on the vertical axis. mFDR plots the marginal false discovery rate; EF plots the expected number of false discoveries along with the actual number of variables included in the model. |
| log.l | Should horizontal axis be on the log scale? Default is FALSE. |
| 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. |
| legend | For type="EF" plots, draw a legend to indicate which line is for the actual selections and which line is for the expected number of false discoveries? Default is TRUE. |
| ... | Other graphical parameters to pass to plot |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

See Also

[mfdr](#), [perm.ncvreg](#)

Examples

```

data(Prostate)
fit <- ncvreg(Prostate$X, Prostate$y)

obj <- mfdr(fit)
obj[1:10,]

# Some plotting options
plot(obj)
plot(obj, type="EF")
plot(obj, log=TRUE)

# Comparison with perm.ncvreg
par(mfrow=c(2,2))
plot(obj)
plot(obj, type="EF")
pmfit <- perm.ncvreg(Prostate$X, Prostate$y)
plot(pmfit)
plot(pmfit, type="EF") # Note that mfdr() is more conservative

```

plot.ncvreg

Plot coefficients from a "ncvreg" object

Description

Produces a plot of the coefficient paths for a fitted "ncvreg" object.

Usage

```

## S3 method for class 'ncvreg'
plot(x, alpha=1, log.l=FALSE, shade=TRUE, ...)

```

Arguments

| | |
|-------|--|
| x | Fitted "ncvreg" 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 FALSE. |
| shade | Should nonconvex region be shaded? Default is TRUE. |
| ... | Other graphical parameters to plot |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5: 232-253.

See Also

[ncvreg](#)

Examples

```
data(Prostate)

fit <- ncvreg(Prostate$X, Prostate$y)
plot(fit)
plot(fit, col="black")
plot(fit, log=TRUE)
```

plot.ncvsurv.func *Plot survival curve for ncvsurv model*

Description

Plot survival curve for a model that has been fit using `ncvsurv` followed by a prediction of the survival function using `predict.ncvsurv`

Usage

```
## S3 method for class 'ncvsurv.func'
plot(x, alpha=1, ...)
```

Arguments

| | |
|-------|--|
| x | A 'ncvsurv.func' object, which is returned by <code>predict.ncvsurv</code> if <code>type='survival'</code> is specified. See examples. |
| alpha | Controls alpha-blending (i.e., transparency). Useful if many overlapping lines are present. |
| ... | Other graphical parameters to pass to <code>plot</code> |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

See Also

[ncvsurv](#), [predict.ncvsurv](#)

Examples

```

data(Lung)
X <- Lung$X
y <- Lung$y

fit <- ncvsurv(X, y)

# A single survival curve
S <- predict(fit, X[1,], type='survival', lambda=.15)
plot(S, xlim=c(0,200))

# Lots of survival curves
S <- predict(fit, X, type='survival', lambda=.15)
plot(S, xlim=c(0,200))

```

predict.ncvreg *Model predictions based on a fitted "ncvreg" object.*

Description

Similar to other predict methods, this function returns predictions from a fitted "ncvreg" object.

Usage

```

## S3 method for class 'ncvreg'
predict(object, X, type=c("link", "response", "class",
"coefficients", "vars", "nvars"), lambda, which=1:length(object$lambda),
...)
## S3 method for class 'ncvreg'
coef(object, lambda, which=1:length(object$lambda),
drop=TRUE, ...)

```

Arguments

| | |
|--------|--|
| object | Fitted "ncvreg" model object. |
| X | Matrix of values at which predictions are to be made. Not used for type="coefficients" or for some of the type settings in predict. |
| lambda | Values of the regularization parameter lambda at which predictions are requested. For values of lambda not in the sequence of fitted models, linear interpolation is used. |
| 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. |
| 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. |

| | |
|------|---|
| 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. |
| ... | Not used. |

Value

The object returned depends on type.

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5: 232-253.

See Also

[ncvreg](#)

Examples

```
data(Heart)

fit <- ncvreg(Heart$X, Heart$y, family="binomial")
coef(fit, lambda=0.05)
head(predict(fit, Heart$X, type="link", lambda=0.05))
head(predict(fit, Heart$X, type="response", lambda=0.05))
head(predict(fit, Heart$X, type="class", lambda=0.05))
predict(fit, type="vars", lambda=c(0.05, 0.01))
predict(fit, type="nvars", lambda=c(0.05, 0.01))
```

predict.ncvsurv

Model predictions based on a fitted "ncvsurv" object.

Description

Similar to other predict methods, this function returns predictions from a fitted "ncvsurv" object.

Usage

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

Arguments

| | |
|--------|--|
| object | Fitted "ncvsurv" model object. |
| X | Matrix of values at which predictions are to be made. Not used for type="coefficients" or for some of the type settings in predict. |
| lambda | Values of the regularization parameter lambda at which predictions are requested. For values of lambda not in the sequence of fitted models, linear interpolation is used. |
| 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. |
| type | Type of prediction: "link" returns the linear predictors; "response" gives the risk (i.e., exp(link)); "survival" returns the estimated survival function; "median" estimates median survival times. The other options are all identical to their ncvreg counterparts: "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. |
| ... | Not used. |

Details

Estimation of baseline survival function conditional on the estimated values of beta is carried out according to the method described in Chapter 4.3 of Kalbfleish and Prentice. In particular, it agrees exactly the results returned by `survfit.coxph(..., type='kalbfleisch-prentice')` in the survival package.

Value

The object returned depends on type.

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

- Breheny P and Huang J (2011). Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5: 232-253.
- Kalbfleish JD and Prentice RL (2002). *The Statistical Analysis of Failure Time Data*, 2nd edition. Wiley.

See Also

[ncvsurv](#)

Examples

```

data(Lung)
X <- Lung$X
y <- Lung$y

fit <- ncvsurv(X,y)
coef(fit, lambda=0.05)
head(predict(fit, X, type="link", lambda=0.05))
head(predict(fit, X, type="response", lambda=0.05))

# Survival function
S <- predict(fit, X[1,], type="survival", lambda=0.05)
S(100)
S <- predict(fit, X, type="survival", lambda=0.05)
plot(S, xlim=c(0,200))

# Medians
predict(fit, X[1,], type="median", lambda=0.05)
M <- predict(fit, X, type="median")
M[1:10, 1:10]

# Nonzero coefficients
predict(fit, type="vars", lambda=c(0.1, 0.01))
predict(fit, type="nvars", lambda=c(0.1, 0.01))

```

Prostate

Factors associated with prostate specific antigen

Description

Data from a study by Stamey et al. (1989) to examine the association between prostate specific antigen (PSA) and several clinical measures that are potentially associated with PSA in men who were about to receive a radical prostatectomy. The variables are as follows:

- X: A design matrix with 97 instances (rows) and 8 predictor variables (columns). The columns are:
 - lcavol: Log cancer volume
 - lweight: Log prostate weight
 - age: The man's age
 - lbph: Log of the amount of benign hyperplasia
 - svi: Seminal vesicle invasion; 1=Yes, 0=No
 - lcp: Log of capsular penetration
 - gleason: Gleason score
 - pgg45: Percent of Gleason scores 4 or 5
- y: Log PSA

Usage

```
data(Prostate)
```

Source

<https://web.stanford.edu/~hastie/ElemStatLearn>

References

- Hastie T, Tibshirani R, and Friedman J. (2001). *The Elements of Statistical Learning*. Springer.
- Stamey T, et al. (1989). Prostate specific antigen in the diagnosis and treatment of adenocarcinoma of the prostate. II. Radical prostatectomy treated patients. *Journal of Urology*, **16**: 1076-1083.

std

Standardizes a design matrix

Description

The function `std` accepts a design matrix and returns a standardized version of that matrix (i.e., each column will have mean 0 and mean sum of squares equal to 1).

Usage

```
std(X)
```

Arguments

`X` A matrix (or object that can be coerced to a matrix, such as a data frame).

Details

This function centers and scales each column of `X` so that

$$\sum_{i=1}^n x_{ij} = 0$$

and

$$n^{-1} \sum_{i=1}^n x_{ij}^2 = 1$$

for all `j`. This is usually not necessary to call directly, as `ncvreg` internally standardizes the design matrix, but inspection of the standardized design matrix can sometimes be useful. This differs from the base R function `scale` in two ways: (1) `scale` uses the sample standard deviation $\sqrt{\text{sum}(x^2)/(n-1)}$, while `std` uses the root-mean-square, or population, standard deviation $\sqrt{\text{mean}(\text{sum}(x^2))}$, and (2) `std` is faster.

Value

The standardized design matrix, with attributes "center" and "scale" corresponding to the mean and (population) standard deviation used to scale the columns.

Author(s)

Patrick Breheny

Examples

```
X <- matrix(rnorm(50), 10, 5)
S <- std(X)
apply(S, 2, sum)
apply(S, 2, function(x) mean(x^2))
```

summary.cv.ncvreg

Summarizing inferences based on cross-validation

Description

Summary method for cv.ncvreg objects

Usage

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

Arguments

| | |
|--------|--|
| object | A "cv.ncvreg" object. |
| x | A "summary.cv.ncvreg" 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. |
| ... | Further arguments passed to or from other methods. |

Value

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

| | |
|---------|--|
| penalty | The penalty used by ncvreg. |
| model | Either "linear" or "logistic", depending on the family option in ncvreg. |
| 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.ncvreg. |
| 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)

Patrick Breheny <patrick-breheny@uiowa.edu>

References

Breheny, P. and Huang, J. (2011) Coordinate descent algorithms for nonconvex penalized regression, with applications to biological feature selection. *Ann. Appl. Stat.*, 5: 232-253.

See Also

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

Examples

```
# Linear regression -----
data(Prostate)
cvfit <- cv.ncvreg(Prostate$X, Prostate$y)
summary(cvfit)

# Logistic regression -----
data(Heart)
cvfit <- cv.ncvreg(Heart$X, Heart$y, family="binomial")
summary(cvfit)
```

summary.ncvreg

Summary method for ncvreg objects

Description

Inferential summaries for ncvreg and ncvsurv objects based on local marginal false discovery rates.

Usage

```
## S3 method for class 'ncvreg'
summary(object, lambda, which, number, cutoff, ...)
## S3 method for class 'ncvsurv'
summary(object, lambda, which, number, cutoff, ...)
## S3 method for class 'summary.ncvreg'
print(x, digits, ...)
```

Arguments

| | |
|--------|---|
| object | An ncvreg or ncvsurv object. |
| lambda | The regularization parameter value at which inference should be reported. |
| which | Alternatively, lambda may be specified by index; which=10 means: report inference for the 10th value of lambda along the regularization path. If both lambda and which are specified, lambda takes precedence. |
| number | By default, summary will provide an inferential summary for each variable that has been selected (i.e., each variable with a nonzero coefficient). Specifying number=5, for example, means that the summary table will include the 5 features with the lowest mfdR values, regardless of whether they were selected. To see all features, number=Inf. |
| cutoff | Alternatively, specifying for example cutoff=0.3 will report inference for all features with mfdR under 30%. If both number and cutoff are specified, the intersection between both sets of features is reported. |
| x | A summary.ncvreg 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. |
| ... | Further arguments; in particular, if you have set returnX=FALSE, you will need to supply X and y in order to calculate local mFDRs. |

Value

summary.ncvreg and summary.ncvsurv produce object with S3 class summary.ncvreg. The class has its own print method and contains the following list elements:

| | |
|-------------|---|
| penalty | The penalty used by ncvreg or ncvsurv. |
| model | Either "linear", "logistic", or "Cox". |
| n | Number of instances. |
| p | Number of regression coefficients (not including the intercept). |
| lambda | The lambda value at which inference is being reported. |
| nvars | The number of nonzero coefficients (again, not including the intercept) at that value of lambda. |
| table | A table containing estimates, normalized test statistics (z), and an estimate of the local mfdR for each coefficient. The mfdR may be loosely interpreted, in an empirical Bayes sense, as the probability that the given feature is null. |
| unpen.table | If there are any unpenalized coefficients, a separate inferential summary is given for them. Currently, this is based on lm/glm/coxph using the penalized coefficients to provide an offset. This is useful and more or less accurate, but not ideal; we hope to improve the inferential methods for unpenalized variables in the future. |

Author(s)

Patrick Breheny <patrick-breheny@uiowa.edu>

See Also

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

Examples

```
# Linear regression -----
data(Prostate)
fit <- ncvreg(Prostate$X, Prostate$y)
summary(fit, lambda=0.08)

# Logistic regression -----
data(Heart)
fit <- ncvreg(Heart$X, Heart$y, family="binomial")
summary(fit, lambda=0.05)

# Cox regression -----
data(Lung)
fit <- ncvsvr(Lung$X, Lung$y)
summary(fit, lambda=0.1)

# Options -----
fit <- ncvreg(Heart$X, Heart$y, family="binomial")
summary(fit, lambda=0.08, number=3)
summary(fit, lambda=0.08, number=Inf)
summary(fit, lambda=0.08, cutoff=0.5)

fit <- ncvreg(Heart$X, Heart$y, family="binomial", returnX=FALSE)
summary(fit, X=Heart$X, y=Heart$y, lambda=0.08)
```

Index

*Topic **datasets**

Heart, [9](#)

Lung, [10](#)

Prostate, [30](#)

AUC (AUC.cv.ncvsurv), [3](#)

AUC.cv.ncvsurv, [3](#)

coef.ncvreg (predict.ncvreg), [27](#)

coef.ncvsurv (predict.ncvsurv), [28](#)

cv.ncvreg, [4](#), [15](#), [23](#), [33](#), [35](#)

cv.ncvsurv, [3](#), [6](#), [18](#)

fir, [8](#), [22](#)

Heart, [9](#)

heart (Heart), [9](#)

Lung, [10](#)

mfdr, [8](#), [11](#), [20](#), [24](#)

ncvreg, [5](#), [11](#), [12](#), [20–23](#), [26](#), [28](#), [33](#), [35](#)

ncvreg-package, [2](#)

ncvsurv, [7](#), [10](#), [11](#), [16](#), [26](#), [29](#)

perm.ncvreg, [11](#), [19](#), [21](#), [22](#), [24](#)

permres, [20](#), [21](#)

plot.cv.ncvreg, [5](#), [7](#), [22](#), [33](#), [35](#)

plot.mfdr, [11](#), [20](#), [24](#)

plot.ncvreg, [15](#), [18](#), [25](#)

plot.ncvsurv.func, [26](#)

predict.ncvreg, [27](#)

predict.ncvsurv, [26](#), [28](#)

print.summary.cv.ncvreg

(summary.cv.ncvreg), [32](#)

print.summary.ncvreg (summary.ncvreg),

[33](#)

Prostate, [30](#)

prostate (Prostate), [30](#)

scale, [31](#)

std, [31](#)

summary.cv.ncvreg, [5](#), [7](#), [32](#)

summary.ncvreg, [13](#), [33](#)

summary.ncvsurv, [17](#)

summary.ncvsurv (summary.ncvreg), [33](#)