

Package ‘RXshrink’

September 1, 2020

Title Maximum Likelihood Shrinkage using Generalized Ridge or Least Angle Regression Methods

Version 1.4.2

Date 2020-09-01

Author Bob Obenchain

Maintainer Bob Obenchain <wizbob@att.net>

Depends R (>= 3.5.0), lars, ellipse

Description Functions are provided to calculate and display ridge TRACE diagnostics for a variety of shrinkage Paths. TRACEs identify the m-Extent of shrinkage most likely, under Normal-theory, to produce optimally biased estimates of beta-coefficients with minimum MSE Risk. The `unr.ridge()` function implements the “Unrestricted Path” introduced in Obenchain (2020) <arXiv:2005.14291>. This Shrinkage-Path is more efficient than the Paths used by the `qm.ridge()`, `aug.lars()` and `uc.lars()` functions. Optimally biased predictions can be made using `RXpredict()` for all five types of RXshrink linear model TRACE diagnostics. Functions `MLboot()`, `MLcalc()`, `MLhist()` and `MLtrue()` provide insights into the true bias and MSE risk characteristics of non-linear Shrinkage estimators. Functions `unr.aug()` and `unr.biv()` augment the calculations made by `unr.ridge()` to provide plots of the bivariate confidence ellipses corresponding to any of the $p^*(p-1)$ possible pairs of shrunken regression coefficients. The `correct.signs()` function provides estimates with “correct” numerical signs when ill-conditioned (nearly multicollinear) models yield OLS estimates that disagree with the signs of the observed correlations between the y-outcome and the selected x-predictor variables.

License GPL-2

URL <https://www.R-project.org> , <http://localcontrolstatistics.org>

NeedsCompilation no

Repository CRAN

Date/Publication 2020-09-01 15:20:03 UTC

R topics documented:

RXshrink-package	2
aug.lars	4

correct.signs	6
haldport	8
kofm	9
longley2	10
MLboot	11
MLcalc	12
MLhist	14
MLtrue	15
mofk	17
mpg	18
plot.aug.lars	18
plot.qm.ridge	19
plot.RXpredict	20
plot.uc.lars	21
plot.unr.biv	22
plot.unr.ridge	23
qm.ridge	24
RXpredict	26
tycobb	28
uc.lars	29
unr.aug	31
unr.biv	32
unr.ridge	33

Index	36
--------------	-----------

RXshrink-package	<i>Maximum Likelihood (ML) Shrinkage using Generalized Ridge or Least Angle Regression Methods</i>
------------------	--

Description

The functions in this package augment basic calculations of Generalized Ridge and Least Angle Regression with important visualization tools. Specifically, they display TRACES of (biased) estimates for 5 KEY quantities [fitted coefficients, relative MSE, excess eigenvalues, inferior direction cosines and shrinkage factors] that completely characterize the effects of shrinkage along Paths with 2 or more parameters: (m-Extent plus q-Shape or Unrestricted). All Paths start at the Ordinary Least-Squares estimate [$m = 0$] and end at the origin, $(0, 0, \dots, 0)$, where all coefficient estimates have been shrunken to zero [$m = p = \text{rank}(X)$]. Three different types of Likelihood of minimal MSE risk (Classical Normal-Theory, Empirical Bayes, and Random Coefficients) can be monitored to suggest an optimal m-Extent of shrinkage for the given matrix of X-variables and the observed y-outcome vector.

Details

Package: RXshrink
Type: Package
Version: 1.4.2
Date: 2020-09-01
License: GNU GENERAL PUBLIC LICENSE, Version 2, June 1991

`unr.ridge()` calculates and displays TRACES for an Unrestricted (p-parameter) shrinkage Path that always passes through the Beta coefficient point-estimate that is most likely to achieve optimal MSE risk reductions under Normal distribution-theory.

`MLboot()`, `MLcalc()` and `MLhist()` support use of Bootstrap resampling to study the bias and MSE risk characteristics of non-linear (unrestricted) Generalized Ridge Regression (GRR) estimators.

When true regression parameters have user-specified (KNOWN) numerical values, `MLtrue()` uses this information and generates a new data.frame that contains a y-outcome vector of the expected form with "disturbance" terms that are IID Normal errors-in-measurement. Arguments to `MLtrue()` must include the "formula" for a desired linear model and a data.frame containing the specified X-variables.

`qm.ridge()` calculates and displays TRACES for traditional Paths defined by 2-parameters: **q**-Shape and **m**-Extent of Shrinkage. The restricted Path of most likely **q**-Shape is found via search on a lattice of 21 values within [-5,+5]. Default lattices for both q and m searches are easy to modify using the `qmax`, `qmin`, `nq` and `steps` arguments to `qm.ridge()`. The "ordinary" ridge Path of Hoerl and Kennard has **q**-Shape = 0, while "uniform" shrinkage corresponds to **q**-Shape = +1. None of these `qm`-Paths generally achieve overall minimum MSE risk when $p > 2$ because they restrict attention to Shrinkage-factors that are "monotome" (increasing or decreasing).

`aug.lars()` augments the Efron-Hastie `lars()` R-function to perform Least Angle Regression with MSE risk calculations and Maximum Likelihood TRACE displays ...like those of `qm.ridge()` and `unr.ridge()`.

`uc.lars()` applies Least Angle Regression methods to the **Uncorrelated Components** of a possibly ill-conditioned set of x-variables. Calculations use a closed-form expression for lars/lasso shrinkage delta-factors that apply because NO Ill-Conditioning is present in these "uc" cases.

`correct.signs()` displays the Normal-theory maximum likelihood estimate of the regression coefficient vector that minimizes MSE Risk in the UNKNOWN direction of p-space PARALLEL to the true Beta vector. This estimate corrects "wrong-sign" problems in the sense that its coefficients have the same relative magnitudes and numerical signs as those of the "Correlation Form" of the X'y vector.

`RXpredict()` makes predictions (i.e. computes "fitted.values") for all 5 types of RXshrink estimation ...either at a user-specified m-Extent of Shrinkage or at the Normal-theory "ML" m-Extent.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Efron B, Hastie T, Johnstone I, Tibshirani R. (2004) Least angle regression. *Annals of Statistics* **32**, 407-499.
- Goldstein M, Smith AFM. (1974) Ridge-type estimators for regression analysis. *J. Roy. Stat. Soc. B* **36**, 284-291. (The 2-parameter shrinkage family.)
- Obenchain RL. (1975) Ridge Analysis Following a Preliminary Test of the Shrunken Hypothesis. *Technometrics* **17**, 431-441. <doi:10.1080/00401706.1975.10489369>
- Obenchain RL. (1978) Good and Optimal Ridge Estimators. *Annals of Statistics* **6**, 1111-1121. <doi:10.1214/aos/1176344314>
- Obenchain RL. (2020) Ridge TRACE Diagnostics. <https://arxiv.org/abs/2005.14291>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like document. <http://localcontrolstatistics.org>

Examples

```
demo(longley2)
```

aug.lars

Maximum Likelihood Estimation of Effects in Least Angle Regression

Description

These functions perform calculations that determine whether least angle and lasso regression estimates correspond to generalized ridge regression (GRR) estimates (i.e. whether they use shrinkage delta-factors that are both non-negative and strictly less than 1.0). They also estimate the Normal-theory likelihood that MSE risk is minimized and compute diagnostics for display in ridge TRACE plots.

Usage

```
aug.lars(form, data, rscale = 1, type = "lar", trace = FALSE,
         eps = .Machine$double.eps, omdmin = 9.9e-13)
```

Arguments

form	A regression formula [y~x1+x2+...] suitable for use with lm().
data	Data frame containing observations on all variables in the formula.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.
type	One of "lasso", "lar" or "forward.stagewise" for function lars(). Names can be abbreviated to any unique substring. Default in aug.lars() is "lar".
trace	If TRUE, lars() function prints out its progress.

eps	The effective zero for lars().
omdmin	Strictly positive minimum allowed value for one-minus-delta (default = 9.9e-013.)

Details

aug.lars() calls the Efron/Hastie lars() function to perform Least Angle Regression on x-variables that have been centered and possibly rescaled but which may be (highly) correlated. Maximum likelihood TRACE displays paralleling those of qm.ridge() and unr.ridge() are also computed and (optionally) plotted.

Value

An output list object of class aug.lars:

form	The regression formula specified as the first argument.
data	Name of the data.frame object specified as the second argument.
p	Number of regression predictor variables.
n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics.
gmat	Orthogonal matrix of direction cosines for regressor principal axes.
lars	An object of class lars.
coef	Matrix of shrinkage-ridge regression coefficient estimates.
risk	Matrix of MSE risk estimates for fitted coefficients.
exev	Matrix of excess MSE eigenvalues (ordinary least squares minus ridge.)
infd	Matrix of direction cosines for the estimated inferior direction, if any.
spat	Matrix of shrinkage pattern multiplicative delta factors.
mlik	Listing of criteria for maximum likelihood selection of M-extent-of-shrinkage.
sxt	Listing of summary statistics for all M-extents-of-shrinkage.
mClk	Most Likely Extent of Shrinkage Observed: best multiple of (1/steps) <= p.
minC	Minimum Observed Value of Normal-theory $-2*\log(\text{Likelihood-Ratio})$.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Breiman L. (1995) Better subset regression using the non-negative garrote. *Technometrics* **37**, 373-384.
- Efron B, Hastie T, Johnstone I, Tibshirani R. (2004) Least angle regression. *Annals of Statistics* **32**, 407-499.
- Efron B, Hastie T. (2005) Least Angle Regression, Lasso and Forward Stagewise. <https://CRAN.R-project.org/package=lars>
- Obenchain RL. (2005) *Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein*. Electronic book-in-progress (185 pages.) <http://localcontrolstatistics.org>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>
- Tibshirani R. (1996) Regression shrinkage and selection via the lasso. *J. Roy. Stat. Soc. B* **58**, 267-288.

See Also

[uc.lars.](#)

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxlobj <- aug.lars(form, data=longley2)
rxlobj
plot(rxlobj)
str(rxlobj)
```

correct.signs

Normal-Theory Maximum Likelihood Estimation of Beta Coefficients with "Correct" Signs

Description

Obenchain(1978) discussed the risk of linear generalized ridge estimators in individual directions within p -dimensional X -space. While shrinkage to ZERO is clearly optimal for all directions strictly ORTHOGONAL to the true BETA, he showed that optimal shrinkage in the UNKNOWN direction PARALLEL to the true BETA is possible. This optimal BETA estimate is of the form $k * X'y$, where k is the positive scalar given in equation (4.2), page 1118. The correct.signs() function computes this estimate, $B(=)$, that uses GRR delta-shrinkage factors proportional to X -eigenvalues.

Usage

```
correct.signs(form, data)
```

Arguments

form	A regression formula [$y \sim x_1 + x_2 + \dots$] suitable for use with <code>lm()</code> .
data	Data frame containing observations on all variables in the formula.

Details

Ill-conditioned (nearly multi-collinear) regression models can produce Ordinary Least Squares estimates with numerical signs that differ from those of the $X'y$ vector. This is disturbing because $X'y$ contains the sample **correlations** between the X-predictor variables and y-response variable. After all, these variables have been "centered" by subtracting off their mean values and rescaled to vectors of length one. Besides displaying OLS estimates, the `correct.signs()` function also displays the "correlation form" of $X'y$, the estimated delta-shrinkage factors, and the k-rescaled beta-coefficients. Finally, the "Bfit" vector of estimates proportional to $B(=)$ is displayed that minimizes the restricted Residual Sum-of-Squares. This restricted RSS of Bfit cannot, of course, be less than the RSS of OLS, but it can be MUCH less than the RSS of $B(=)$ whenever $B(=)$ shrinkage appears excessive.

Value

An output list object of class "correct.signs":

data	Name of the data.frame object specified as the second argument.
form	The regression formula specified as the first argument.
p	Number of regression predictor variables.
n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics (p by 5) from <code>qm.ridge()</code> .
kpb	Maximum likelihood estimate of k-factor in equation (4.2) of Obenchain(1978).
bmf	Rescaling factor for $B(=)$ to minimize the Residual Sum-of-Squares.
signs	Listing of five Beta coefficient statistics (p by 5): OLS, $X'y$, Delta, $B(=)$ and Bfit.
loff	Lack-of-Fit statistics: Residual Sum-of-Squares for OLS, $X'y$, $B(=)$ and Bfit.
sqcor	Squared Correlation between the y-vector and its predicted values. The two values displayed are for OLS predictions or for predictions using Bfit, $X'y$ or $B(=)$. These two values are the familiar R^2 coefficients of determination for OLS and Bfit.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Obenchain RL. (1978) Good and Optimal Ridge Estimators. *Annals of Statistics* **6**, 1111-1121. <doi:10.1214/aos/1176344314>
- Obenchain RL. (2005) *Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein*. Electronic book-in-progress (185+ pages.) <http://localcontrolstatistics.org>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[qm.ridge](#), [unr.ridge](#) and [MLtrue](#).

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxcsobj <- correct.signs(form, data=longley2)
rxcsobj
str(rxcsobj)
```

haldport

Portland Cement benchmark of Hald(1952)

Description

Heat evolved during setting of 13 cement mixtures of four (or five) ingredients. The first four ingredient percentages appear to be "rounded down" to a full integer. The fifth integer percentage of "other" material assures that the five percentages sum to exactly 100%. However, the "centered" X-matrix resulting from inclusion of all five ingredients would then be Singular (rank=4). In other words, regressing any y-Outcome on only the first four X-variables yields an "ill-conditioned" model that, while having numerical "full rank"=4, actually suffers a effective "rank deficiency" of at least $MCAL=1$.

Usage

```
data(haldport)
```

Format

A data frame with 13 observations on the following 6 variables.

- p3ca Positive Integer percentage of $3CaO \cdot Al_2O_3$ in the mixture.
- p3cs Positive Integer percentage of $3CaO \cdot SiO_2$ in the mixture.
- p4caf Positive Integer percentage of $4CaO \cdot Al_2O_3 \cdot Fe_2O_3$ in the mixture.
- p2cs Positive Integer percentage of $2CaO \cdot SiO_2$ in the mixture.
- other Positive Integer percentage of other ingredients in the mixture.
- heat Heat (cals/gm) evolved in setting, recorded to nearest tenth.

Details

The (RXshrink) haldport data are identical to the (MASS) cement data except for variable names and inclusion of the "other" X-variable.

Source

Woods H, Steinour HH, Starke HR. "Effect of composition of Portland cement on heat evolved during hardening. Industrial Engineering and Chemistry 1932; 24: 1207-1214.

References

Hald A. Statistical Theory with Engineering Applications. 1952 (page 647.) New York; Wiley.

kofm	<i>k-Multipliers and delta-Factors for unr.ridge() Shrinkage.</i>
------	---

Description

The kofm() function computes the "kStar" multipliers and the diagonal "d" Matrices used by the unr.ridge() function. While these calculations are extensive, they are greatly simplified by using interpolation between "knots" on "piecewise linear-spline" functions.

Usage

```
kofm(muobj, p, dMSE, delmax = 0.999999)
```

Arguments

muobj	A "numeric" value that is ≥ 0 and $\leq p$.
p	The integer number of non-constant x-variables used in defining the linear model being fitted to ill-conditioned (intercorrelated, confounded) data. Note that p must also be rank of the given X-matrix.
dMSE	Maximum Likelihood estimates of minimum MSE risk Shrinkage Delta-Factors.
delmax	Maximum allowed value for Shrinkage delta-factors that is strictly less than 1. (default = 0.999999, which prints as 1 when rounded to fewer than 6 decimal places.)

Value

An output list of two objects:

kStar	The k-factor needed to multiply each given dMSE-factor.
d	The (p x p) diagonal matrix of requested shrinkage-factors.

Author(s)

Bob Obenchain <wizbob@att.net>

References

Obenchain RL. (2020) *Ridge TRACE Diagnostics*. <https://arxiv.org/abs/2005.14291>

Obenchain RL. (2020) *The Unrestricted Shrinkage Path: Technical Details*. "unrPath.pdf" <http://localcontrolstatistics.org>

See Also

[unr.ridge](#) and [mofk](#).

longley2

Art Hoerl's update of the infamous Longley(1967) benchmark dataset

Description

Data from the "Employment and Training Report of the President, 1976" compiled by Art Hoerl, University of Delaware. Some slightly different numerical values than those used by Longley(1967) ...plus the added years of 1963-1975. Longley(1967) used only data from the 16 years of 1947 through 1962.

Usage

```
data(longley2)
```

Format

A data frame of 7 variables collected for 29 consecutive years; no NAs.

GNP.deflator GNP price deflation index.

Unemployed Unemployment percentage.

Armed.Forces Size of the Armed Forces.

Population Total U.S. Population.

Year 1947 to 1975.

Employed employment.

GNP Gross National Product.

References

Longley JW. An appraisal of least-squares programs from the point of view of the user. J. Amer. Statist. Assoc. 1967; 62: 819-841.

MLboot	<i>Calculate Bootstrap distribution of Unrestricted Maximum Likelihood (ML) point-estimates for a Linear Model.</i>
--------	---

Description

Resample With-Replacement from a given data.frame and recompute MSE risk-optimal estimates of Beta-Coefficients and their Relative MSE risks using MLcalcs() to compute ML point-estimates.

Usage

```
MLboot(form, data, reps=100, seed, rscale=1)
```

Arguments

form	Regression formula [y~x1+x2+...] suitable for use with lm().
data	data.frame containing observations on all variables in the formula.
reps	Number of Bootstrap replications: Minimum reps = 10, Default is reps = 100. While reps = 10000 is reasonable for bivariate (p=2) linear models, even that many reps could be excessive for models with $p \gg 2$.
seed	Either an Integer between 1 and 999 or else missing to generate a random seed.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.

Details

Ill-conditioned and/or nearly multi-collinear linear regression models are unlikely to yield reasonable ML unbiased (OLS) point-estimates. But more reasonable ML "optimally-biased" point-estimates from generalized ridge regression (GRR) typically have questionable MSE risk characteristics because they are complicated non-linear functions of the observed y-outcome vector. Thus the distribution of bootstrap resamples is of considerable interest in both theory and practice.

Value

An output list object of class MLboot:

data	Name of the data.frame object specified as the second argument.
form	The regression formula specified as the first argument.
reps	Number of Bootstrap replications performed.
seed	Seed used to start random number generator.
n	Number of complete observations after removal of all missing values.
p	Number of beta, rmse or dmse estimates resampled.
ols.beta	OLS matrix (reps x p) of unbiased beta-coefficient estimates.

ols.rmse	OLS matrix (reps x p) of beta-coefficient relative variances.
opt.dmse	OPT matrix (reps x p) of delta shrinkage-factors with minimum MSE risk.
opt.beta	OPT matrix (reps x p) of biased beta-coefficient estimates.
opt.rmse	OPT matrix (reps x p) of beta-coefficient relative MSE risks.

Author(s)

Bob Obenchain <wizbob@att.net>

References

Thompson JR. (1968) Some shrinkage techniques for estimating the mean. *Journal of the American Statistical Association* **63**, 113-122. (The "cubic" estimator.)

Obenchain RL. (2020) Ridge TRACE Diagnostics. <https://arxiv.org/abs/2005.14291>

Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[unr.ridge](#), [correct.signs](#)

MLcalc	<i>Calculate Unrestricted Maximum Likelihood (ML) point-estimates for a Linear Model that are either Unbiased (OLS) or Optimally Biased under Normal-distribution theory.</i>
--------	---

Description

Compute MSE risk-optimal point-estimates of Beta-Coefficients and their Relative MSE risks. Much of the code for this function is identical to that of `unr.ridge()`, which computes multiple points along the "Unrestricted" shrinkage-path. `MLcalc()` restricts attention to only 2 points: Unbiased OLS (BLUE) and the optimally Biased estimate with minimum MSE.

Usage

```
MLcalc(form, data, rscale = 1, delmax = 0.999999)
```

Arguments

form	A regression formula [y~x1+x2+...] suitable for use with <code>lm()</code> .
data	Data frame containing observations on all variables in the formula.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.
delmax	Maximum allowed value for Shrinkage delta-factors that is strictly less than 1. (default = 0.999999, which prints as 1 when rounded to fewer than 6 decimal places.)

Details

Ill-conditioned and/or nearly multi-collinear regression models are unlikely to produce Ordinary Least Squares (OLS) regression coefficient estimates that are very close, numerically, to their unknown true values. Specifically, OLS estimates can then tend to have "wrong" numerical signs and/or unreasonable relative magnitudes, while shrunken (generalized ridge) estimates chosen to maximize their likelihood of reducing Mean Squared Error (MSE) Risk (expected loss) can be much more stable and reasonable numerically. On the other hand, because only OLS estimates are guaranteed to be minimax when risk is matrix valued (truly multivariate), no guarantee of an actual reduction in MSE Risk is necessarily associated with Generalized RR shrinkage.

Value

An output list object of class MLcalc:

data	Name of the data.frame object specified as the second argument.
form	The regression formula specified as the first argument.
p	Number of regression predictor variables.
n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics.
gmat	Orthogonal Matrix of Direction Cosines for Principal Axes [1:p, 1:p].
beta	Numerical shrinkage-ridge regression coefficient estimates [1:2, 1:p].
rmse	Numerical MSE risk estimates for fitted coefficients [1:2, 1:p].
dMSE	Numerical delta-factors for shrinking OLS components [1:p].
ys	Numerical rescaling factor for y-outcome variable [1, 1].
xs	Numerical rescaling factors for given x-variables [1:p].

Author(s)

Bob Obenchain <wizbob@att.net>

References

Thompson JR. (1968) Some shrinkage techniques for estimating the mean. *Journal of the American Statistical Association* **63**, 113-122. (The "cubic" estimator.)

Obenchain RL. (2020) Ridge TRACE Diagnostics. <https://arxiv.org/abs/2005.14291>

Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[unr.ridge](#), [MLboot](#)

MLhist

*Plot method for MLboot objects***Description**

Frequency Histogram displays that use both a specified "Middle Percentage" of a MLboot() distribution that contains outliers as well as a proposed total number of bins.

Usage

```
MLhist(x, comp="opt.beta", xvar=1, npct = 95, bins = 50 )
```

Arguments

x	An Output-list object of class "MLboot".
comp	One of five possible choices of MLboot() resampled estimates: "ols.beta", "ols.rmse", "opt.dmse", "opt.beta" or "opt.rmse".
xvar	Column number of estimates to be displayed: $1 \leq xvar \leq x\$p$.
npct	An integer percentage of the simulated scalar estimates to be displayed. This percentage should be at least 66 and at most 100.
bins	This proposed number of histogram "breaks" is only a suggestion; actual break-points will be set to "pretty" values.

Value

An output list object of class MLhist:

x	Character string showing user choices for "comp" and "xvar" arguments.
form	The regression formula specifying the linear model in MLboot().
reps	Number of Bootstrap replications performed.
npct	An integer percentage of the simulated scalar estimates displayed.
rbins	Number of histogram Bins requested.
dbins	Number of histogram Bins actually displayed.
ntot	Total number of rows of scalar estimates.
p	Total number of columns of scalar estimates.
nlo	First (smallest) order statistic displayed in Histogram.
nup	Last (largest) order statistic displayed in Histogram.
noin	Total number of scalar estimates displayed in Histogram.
xmn	Observed Mean Estimate: location of vertical "blue" dashed line on Histogram.

Author(s)

Bob Obenchain <wizbob@att.net>

MLtrue	<i>Simulate data for Linear Models with known Parameter values and Normal Errors</i>
--------	--

Description

Using specified numerical values for Parameters, such as the true error-term variance, that usually are unknown, MLtrue() creates a new data.frame that contains variables "Yhat" (the vector of expected values) and "Yvec" (Yhat + IID Normal error-terms) as well as the p given X-variables. Thus, MLtrue() produces "correct" linear models that are ideal for analysis and bootstrapping via methods based on Normal-theory.

Usage

```
MLtrue(form, data, seed, go=TRUE, truv, trub, truc)
```

Arguments

form	A regression formula [$y \sim x_1 + x_2 + \dots$] suitable for use with lm().
data	Data frame containing observations on all variables in the formula.
seed	Seed for random number generation between 1 and 1000. When this seed number is missing in a call to MLtrue(), a new seed is generated using runif(). Error terms are then generated using calls to mnorm(), and the seed used is reported in the MLtrue output.list to enable reproducible research.
go	Logical value: go=TRUE starts the simulation using OLS estimates or the numerical values from the (optional) truv, trub and truc arguments to define the "Yhat" and "Yvec" variables; go=FALSE causes MLtrue() to compute OLS estimates that ignore the truv, trub and truc arguments. Thus, go=FALSE provides a convenient way for users to Print and Examine the MLtrue() output.list before deciding which parameter-settings they wish to actually USE or MODIFY before making a final run with go=TRUE.
truv	Optional: numerical value for the true error variance, σ^2 .
trub	Optional: column vector of numerical values for the true regression coefficients.
truc	Optional: column vector of numerical values for the true uncorrelated components of the regression coefficient vector.

Details

RXshrink functions like unr.ridge() and qm.ridge() calculate maximum likelihood estimates ...either unbiased (BLUE) or "optimally" biased (minimum MSE risk) along a specified "path"... for typical statistical inference situations where true regression parameters are unknown. Furthermore the specified linear-model may be "incorrect" or the error-terms may not be IID Normal variates. In sharp contrast with this usual situation, the MLtrue() function generates a "Yvec" vector of outcomes from a CORRECT model of the given form that does contain IID Normal error-terms and all true parameter values are KNOWN. This makes it interesting to compare plots and output.lists

from RXshrink GRR-estimation functions on a "real" data.frame with the corresponding outputs from a MLtrue() data.frame with known parameter-setting that are either identical to or "close" to those estimated from "real" data. WARNING: All output x-variables are "centered" and "rescaled" to have mean ~ 0 and variance 1. Yhat expected values are "centered" but usually have variance differing from 1. Yvec values are not "centered" and have variance determined by either the original data.frame or the "truv" setting.

Value

An output list object of class MLtrue:

new	A data.frame containing the Yvec and Yhat vectors as well as centered and rescaled versions of the input X-matrix and original Y-vector.
Yvec	An additional copy of Yvec pseudo-data containing Normal error-terms.
Yhat	An additional copy of Yhat linear fitted-values.
seed	Seed actually used in random-number generation; needed for replication.
tvar	True numerical value of σ^2 ...[1,1] matrix.
tbeta	True numerical values of beta-coefficients ...[p,1] matrix.
tcomp	True numerical values of uncorrelated components ...[p,1] matrix.
useb	Logical: TRUE => tbeta vector was used in simulation. FALSE => tcomp vector was used rather than tbeta.
data	The name of the data.frame object specified in the second argument.
form	The regression formula specified in the first argument. NOTE: The final 11 output items [p, ..., xs] are identical to those in output lists from RXshrink functions like unr.ridge() and qm.ridge().

Author(s)

Bob Obenchain <wizbob@att.net>

See Also

[qm.ridge](#), [unr.ridge](#) and [MLboot](#).

Examples

```
# require(RXshrink)
data(mpg)
form <- mpg~cylnds+cubins+hpower+weight
MLtest <- MLtrue(form, mpg, go=FALSE) # all other potential arguments "missing"...
MLtest # print current parameter estimates...
cvec <- c(-0.5, 0.15, -0.16, -0.6) # define alternative "true" components...
MLout <- MLtrue(form, mpg, trunc = cvec) # Use "trunc" input with default go=TRUE...
str(MLout)
formY <- Yhat~cylnds+cubins+hpower+weight # Formula for true expected Y-outcomes...
lmobj <- lm(formY, MLout$new)
max(abs(lmobj$residuals)) # essentially 0 because linear model is "correct"...
# unrobj <- unr.ridge(formY, MLout$new) ...generates "Error" because RSQUARE=1.
```

mofk

m-Extents of Shrinkage used in unr.ridge() Calculations.

Description

The `mofk()` function computes the numerical *m-Extent of Shrinkage* corresponding to a given value for "k", the scalar that multiplies observed ML estimates of "dMSE"-factors for each of the "p" uncorrelated components in GRR estimation using `unr.ridge()`. These calculations are quite simple; the function is many-to-one in the sense that all truly large values of k imply that $m = 0.0$, and k values that are either 0.0 or negative imply that $m = p$.

Usage

```
mofk(p, k, dMSE)
```

Arguments

p	The integer number of non-constant x-variables used in defining the linear model being fitted to ill-conditioned (intercorrelated, confounded) data. Note that p must also be rank of the given X-matrix.
k	The given k-factor that multiplies each given dMSE-factor.
dMSE	Maximum Likelihood estimates of Shrinkage Delta-Factors with minimum MSE risk.

Value

The appropriate scalar value for m:

m	The computed <i>m-Extent of Shrinkage</i> ...a scalar within $[0, p]$.
---	---

Author(s)

Bob Obenchain <wizbob@att.net>

References

Obenchain RL. (2020) *Ridge TRACE Diagnostics*. <https://arxiv.org/abs/2005.14291>

Obenchain RL. (2020) *The Unrestricted Shrinkage Path: Technical Details*. "unrPath.pdf" <http://localcontrolstatistics.org>

See Also

[unr.ridge](#) and [kofm](#).

 mpg

Hocking(1976) Miles Per Gallon benchmark dataset

Description

Performance data on 32 autos from Motor Trends magazine.

Usage

```
data(mpg)
```

Format

A data frame of 11 variables collected for 32 automobiles; no NAs.

cylnds number of cylinders.

cubins cubic inches of engine displacement.

hpower engine horsepower.

weight total auto weight in pounds.

mpg miles per gallon.

shape engine shape (1=Straight, 0=V).

transm transmission type.

speeds number of forward speeds.

carbs number of carburetors.

fratio final drive ratio.

qmilt quarter mile time.

References

Hocking RA. "The Analysis and Selection of Variables in Regression." *Biometrics* 1976; 32: 1-51.

 plot.aug.lars

Plot method for aug.lars objects

Description

Plot TRACE displays for aug.lars regression coefficients. The default is to display all five traces on one page with no legends.

Usage

```
## S3 method for class 'aug.lars'
plot(x, trace = "all", trkey = FALSE, ...)
```

Arguments

x	Output list object of class aug.lars.
trace	One of seven possible options: "all" to display 5 traces in one graph, "seq" to display 5 full-sized traces in sequence in response to user prompts, "coef" to display only the estimated beta coefficient trace, "rmse" to display only the estimated relative MSE risk trace, "exev" to display only the estimated excess MSE eigenvalue (OLS minus larlso), "infd" to display only the estimated inferior direction cosine trace, or "spat" to display only the delta-factor pattern trace.
trkey	If TRUE, display a crude legend at the bottom of each trace plot.
...	Optional argument(s) passed on to plot().

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxlobj <- aug.lars(form, data=longley2)
plot(rxlobj)
```

plot.qm.ridge

Plot method for qm.ridge objects

Description

Plot TRACE displays for 2-parameter (q, m) generalized ridge coefficients. The default is to display all five traces on one page with no legends.

Usage

```
## S3 method for class 'qm.ridge'
plot(x, trace = "all", trkey = FALSE, ... )
```

Arguments

x	Output list object of class qm.ridge.
trace	One of seven possible options: "all" to display 5 traces in one graph, "seq" to display 5 full-sized traces in sequence in response to user prompts, "coef" to display only the estimated beta coefficient trace, "rmse" to display only the estimated relative MSE risk trace, "exev" to display only the estimated excess MSE eigenvalue (OLS minus larlso), "infd" to display only the estimated inferior direction cosine trace, or "spat" to display only the delta-factor pattern trace.

trkey If TRUE, display a crude legend at the bottom of each trace plot.
 ... Optional argument(s) passed on to plot().

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxrobj <- qm.ridge(form, data=longley2)
plot(rxrobj)
```

plot.RXpredict

Plot method for RXpredict objects

Description

Plot Predicted and/or Fitted.Values for all 5 RXshrink regression estimation methods. The default is to plot Predictions for the y-Outcome variable.

Usage

```
## S3 method for class 'RXpredict'
plot(x, fit = "yvecprd", ... )
```

Arguments

x Output list object of class RXpredict.
 fit One of three possible options: "yvecprd" to display Predictions of the Observed y-Outcomes in a single plot, "cryprd" to display Fitted.Values for the Centered and Rescaled y-Outcomes, "both" to display "yvecprd" and "cryprd" plots as rows on one page.
 ... Optional argument(s) passed on to plot().

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```

data(longley2)
form <- Employed~GNP+GNP.deflator+Unemployed+Armed.Forces+Population+Year
rxuobj <- unr.ridge(form, longley2)
rxuprd <- RXpredict(rxuobj, longley2)
plot(rxuprd)
# Clearly Biased predictions can still represent "Optimal" Variance-Bias Trade-Offs...

```

plot.uc.lars

Plot method for uc.lars objects

Description

Plot TRACE displays of lars regression coefficients for uncorrelated components. The default is to display all five traces on one page with no legends.

Usage

```

## S3 method for class 'uc.lars'
plot(x, trace = "all", trkey = FALSE, ...)

```

Arguments

x	Output list object of class uc.lars.
trace	One of seven possible options: "all" to display 5 traces in one graph, "seq" to display 5 full-sized traces in sequence in response to user prompts, "coef" to display only the estimated shrunken beta coefficient trace, "rmse" to display only the estimated relative MSE risk trace, "exev" to display only the estimated excess MSE eigenvalue (OLS minus ridge) trace, "infd" to display only the estimated inferior direction cosine trace, or "spat" to display only the shrinkage (delta) factor pattern trace.
trkey	If TRUE, display a crude legend at the bottom of each trace plot.
...	Optional argument(s) passed on to plot().

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```

data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxucobj <- uc.lars(form, data=longley2)
plot(rxucobj)

```

plot.unr.biv

Plot method for unr.biv objects

Description

Plot Bivariate Confidence Ellipse(s) or a TRACE display for unrestricted ridge shrinkage coefficients. The default is to display two ellipses at confidence levels 0.95 and 0.50 on a single plot for the first 2 of $p \geq 2$ coefficients. The projection of the `unr.ridge()` Path from the OLS solution ($m == 0$) to the Minimum MSE Risk solution ($m == mUnr$) onto this bivariate-plane is also shown in "red".

Usage

```
## S3 method for class 'unr.biv'
plot(x, type = "ellip", ...)
```

Arguments

x	Output list object of class <code>unr.ridge</code> .
type	One of 2 options: "ellip" or "trace". The default option of "ellip" displays the Confidence Ellipse(s) specified by arguments to <code>unr.biv()</code> .
...	Optional argument(s) passed on to <code>plot()</code> .

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```
# Cost-Effectiveness inferences using Linear Models and GRR...
## Not run:
# library(ICEinfer)
data(sepsis)
ndr <- ICEpref(sepsis$qalypres, sepsis$totcost, lambda=50000, beta=0.1)
sndr <- data.frame(cbind(ndr, sepsis)) # ndr: non-linear diminishing returns...
form4 <- ndr ~ icu + age + orgfails + apache
usndra <- unr.aug(unr.ridge(form4, sndr)) # compare ndr of 2 Intensive Care Units...
plot(uobj <- unr.biv(usndra, 2, 4))
uobj # implicit print...
# plot(uobj, type = "tr")

## End(Not run)
```

plot.unr.ridge *Plot method for unr.ridge objects*

Description

Plot TRACE displays for unrestricted ridge shrinkage coefficients. The default is to display all five traces on one page with no legends.

Usage

```
## S3 method for class 'unr.ridge'  
plot(x, trace = "all", trkey = FALSE, ... )
```

Arguments

x	Output list object of class unr.ridge.
trace	One of seven possible options: "all" to display 5 traces in one graph, "seq" to display 5 full-sized traces in sequence in response to user prompts, "coef" to display only the estimated beta coefficient trace, "rmse" to display only the estimated relative MSE risk trace, "exev" to display only the estimated excess MSE eigenvalue (OLS minus lars), "infd" to display only the estimated inferior direction cosine trace, or "spat" to display only the delta-factor pattern trace.
trkey	If TRUE, display a crude legend at the bottom of each trace plot.
...	Optional argument(s) passed on to plot().

Value

NULL

Author(s)

Bob Obenchain <wizbob@att.net>

Examples

```
data(longley2)  
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed  
rxunobj <- unr.ridge(form, data=longley2)  
plot(rxunobj)
```

qm.ridge	<i>Restricted (2-parameter) Maximum Likelihood Shrinkage in Regression</i>
----------	--

Description

Computes and displays TRACES for a **q**-Shaped shrinkage PATH, including the **m**-Extent of shrinkage along that path, most likely under Normal-distribution theory to yield optimal reductions in MSE Risk. When $\text{rank}(X\text{-matrix}) = p > 2$, the most likely solution on the `qm.ridge()` path will be less likely to have minimal MSE risk than the optimal solution on the `p`-parameter `unr.ridge()` path. The Hoerl-Kennard "ordinary" ridge path has Shape $q=0$ within the `qm.ridge()` family.

Usage

```
qm.ridge(form, data, rscale = 1, Q = "qmse", steps = 8, nq = 21,
         qmax = 5, qmin = -5, omdmin = 9.9e-13)
```

Arguments

form	A regression formula [y~x1+x2+...] suitable for use with <code>lm()</code> .
data	Data frame containing observations on all variables in the formula.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.
Q	Shape parameter that controls the curvature of the shrinkage path through regression-coefficient likelihood space (default = "qmse" implies use the value found most likely to be optimal.) Use $Q = 0$ to specify Hoerl-Kennard "ordinary" ridge regression.
steps	Number of equally spaced values per unit change along the horizontal m -Extent-of-shrinkage axis for estimates to be calculated and displayed in TRACES (default = 8.)
nq	Number of equally spaced values on the lattice of all possible values for shrinkage q -Shape between the "qmin" and "qmax" parameter settings (default = 21.)
qmax	Maximum allowed q -Shape (default = +5.)
qmin	Minimum allowed q -Shape (default = -5.)
omdmin	Strictly positive minimum value for one-minus-delta (default = 9.9e-013.)

Details

Traditional `qm.ridge()` paths cannot be overall-optimal when $p > 2$ because they are restricted to using strictly "monotone" (increasing or decreasing) shrinkage factors. Still, the "best" **m**-Extent of **qm**-shrinkage is marked by a vertical dashed-line on all 5-types of `qm.ridge()` TRACE displays. Compared to OLS estimates, these shrunken estimates have higher likelihood of reduced MSE risk and can be much more stable and reasonable, numerically. On the other hand, because only OLS estimates are guaranteed to be minimax when risk is MATRIX valued (truly multivariate), no guarantee of an actual reduction in MSE Risk is necessarily associated with shrinkage.

Value

An output list object of class qm.ridge:

form	The regression formula specified as the first argument.
data	Name of the data.frame object specified as the second argument.
p	Number of regression predictor variables.
n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics.
mx	Matrix containing mean values of X-predictors.
crlqstat	Listing of criteria for maximum likelihood selection of path q-Shape.
qmse	Numerical value of q-Shape most likely to be optimal.
qp	Numerical value of the q-Shape actually used for shrinkage.
coef	Matrix of shrinkage-ridge regression coefficient estimates.
risk	Matrix of MSE risk estimates for fitted coefficients.
exev	Matrix of excess MSE eigenvalues (ordinary least squares minus ridge.)
infd	Matrix of direction cosines for the estimated inferior direction, if any.
spat	Matrix of shrinkage pattern multiplicative delta factors.
mlik	Listing of criteria for maximum likelihood selection of m-Extent-of-shrinkage.
sext	Listing of summary statistics for all m-Extents-of-shrinkage.
mClk	Most-Likely Extent of Shrinkage Observed: best multiple of (1/steps) <= p.
minC	Minimum Observed Value of Normal-theory $-2*\log(\text{Likelihood-Ratio})$.
QS	Was a Mesh-Search for Best q-Shape requested? : 1 => Yes, 0 => No.
qML	Computable only when p=rank=2: True Most-Likely q-Shape.
kML	Computable only when p=rank=2: True Most-Likely k-Factor.
dML1	Computable only when p=rank=2: True Most-Likely Delta[1]-Factor.
dML2	Computable only when p=rank=2: True Most-Likely Delta[2]-Factor.
mML	Computable only when p=rank=2: True Most-Likely m-Extent.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Burr TL, Fry HA. (2005) Biased Regression: The Case for Cautious Application. *Technometrics* **47**, 284-296.
- Goldstein M, Smith AFM. (1974) Ridge-type estimators for regression analysis. *J. Roy. Stat. Soc. B* **36**, 284-291. (The 2-parameter shrinkage family.)
- Obenchain RL. (1975) Ridge Analysis Following a Preliminary Test of the Shrunken Hypothesis. *Technometrics* **17**, 431-441. <doi:10.1080/00401706.1975.10489369>
- Obenchain RL. (1978) Good and Optimal Ridge Estimators. *Annals of Statistics* **6**, 1111-1121. <doi:10.1214/aos/1176344314>
- Obenchain RL. (2005) *Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein*. Electronic book-in-progress (185+ pages.) <http://localcontrolstatistics.org>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like document. <http://localcontrolstatistics.org>

See Also

[unr.ridge](#), [correct.signs](#), [MLtrue](#) and [RXpredict](#).

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxrobj <- qm.ridge(form, data=longley2)
rxrobj
plot(rxrobj)
str(rxrobj)
```

RXpredict

Predictions from Models Fit using RXshrink Generalized Ridge Estimation Methods.

Description

RXpredict() makes in-sample predictions (i.e. computes "fitted.values") for all 5 forms of RXshrink estimation either at user specified m-Extents of shrinkage or at the normal-theory Maximum Likelihood (ML) m-Extent.

Usage

```
RXpredict(x, data, m="ML", rscale=1)
```

Arguments

x	An object output by one of the 5 RXshrink estimation functions. Thus <code>class(x)</code> must be "qm.ridge", "unr.ridge", "aug.lars", "uc.lars" or "correct.signs".
data	Data frame containing observations on all variables used by the RXshrink function used in estimation.
m	Either [i] a "numeric" value ≥ 0 and $\leq \text{rank}(X)$ or [ii] the (default) character string "ML" to request use of the m-Extent of shrinkage most likely to be MSE optimal under Normal distribution-theory.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.

Value

An output list object of class RXpredict:

cryprd	Predicted values for the "centered" and possibly "rescaled" outcome y-vector, cry. These values correspond, for example, to the default "predicted.values" from <code>lm()</code> .
cry	This the "centered" and possibly "rescaled" outcome y-vector from the input data.frame.
yvecprd	Predicted values for the Y-outcome variable, yvec.
yvec	The Y-outcome vector from the input data.frame specified by the "data" argument.
m	"numeric" Value of m-Extent requested in the call to <code>RXpredict()</code> , possibly via a default call with <code>m="ML"</code> . Restriction: $0 \leq m \leq \text{rank}(X)$.
mobs	Observed m-Extent most close to the above m, yet is on the lattice of m-Extents stored within the x-object.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Obenchain RL. (1978) Good and Optimal Ridge Estimators. *Annals of Statistics* **6**, 1111-1121. <[doi:10.1214/aos/1176344314](https://doi.org/10.1214/aos/1176344314)>
- Obenchain RL. (2005) *Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein*. Electronic book-in-progress (185+ pages.) <http://localcontrolstatistics.org>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[qm.ridge](#), [unr.ridge](#) and [MLtrue](#).

Examples

```

data(tycobb)
form <- batavg~atbats+seasons+CMspl
rxuobj <- unr.ridge(form, data=tycobb)
tycfit <- RXpredict(rxuobj, tycobb, m="ML")
plot(tycfit)
tycobb$batavg[18] # Ty Cobb's batavg = 0.401 in 1922
abline(h=tycfit$cry[18], lty=2, lwd=3, col="red")

```

tycobb

Ty Cobb batting statistics for 1905–1928 with Carl Morris' 2-piece Spline term.

Description

Linear Regression models can be used to predict Ty Cobb's Expected true yearly batting averages from his observed yearly "batavg" and the 5 other variables stored in the "tycobb" data.frame. Predictions from such "models" can address the question: "Was Ty Cobb ever a TRUE .400 hitter?" Since a player's seasonal batavg is a "random variable," the fact that Cobb's batavg was 0.4196 in 1911 and 0.4105 in 1912 does not necessarily imply that his Expected Seasonal batavg was truly over .400 in either of those consecutive seasons. For example, his batavg was 0.4011 in 1922 (10 years later). However, his seasonal batavg had dipped to "only" 0.3341 in 1920.

"Cobb lived off the field as though he wished to live forever. He lived on the field as though it was his last day." – Branch Rickey, Major League Baseball Hall of Fame executive.

Usage

```
data(tycobb)
```

Format

A data frame with 24 observations (years) on the following 6 variables.

year Ty Cobb's 24 American League Seasons: 1905 - 1928.

hits Total number of Hits that season.

atbats Total number of times at Bat that season.

CMspl Carl Morris' Piecewise-Linear Spline term with "knot" in season 6 (1910).

seasons Linear "Trend" term: 1, 2, ..., 24.

batavg Cobb's Seasonal Batting Average ... 7 decimal places.

References

Carl Morris. (1982). "Was Ty Cobb ever a TRUE .400 hitter?" One-page Handout for his JSM Lecture on August 18 in Cincinnati, Ohio.

uc.lars	<i>Maximum Likelihood Least Angle Regression on Uncorrelated X-Components</i>
---------	---

Description

Apply least angle regression estimation to the uncorrelated components of a possibly ill-conditioned linear regression model and generate normal-theory maximum likelihood TRACE displays.

Usage

```
uc.lars(form, data, rscale = 1, type = "lar", trace = FALSE,
        eps = .Machine$double.eps, omdmin = 9.9e-13)
```

Arguments

form	A regression formula [$y \sim x_1 + x_2 + \dots$] suitable for use with <code>lm()</code> .
data	Data frame containing observations on all variables in the formula.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.
type	One of "lasso", "lar" or "forward.stagewise" for function <code>lars()</code> . Names can be abbreviated to any unique substring. Default in <code>uc.lars()</code> is "lar".
trace	If TRUE, <code>lars()</code> function prints out its progress.
eps	The effective zero for <code>lars()</code> .
omdmin	Strictly positive minimum allowed value for one-minus-delta (default = $9.9e-013$.)

Details

`uc.lars()` applies Least Angle Regression to the uncorrelated components of a possibly ill-conditioned set of x -variables. A closed-form expression for the lars/lasso shrinkage delta factors exists in this case: $\Delta(i) = \max(0, 1 - k/|\text{PC}(i)|)$, where $\text{PC}(i)$ is the principal correlation between y and the i -th principal coordinates of X . Note that the k -factor in this formulation is limited to a subset of $[0, 1]$. $\text{MCAL} = 0$ occurs at $k = 0$, while $\text{MCAL} = p$ results when k is the maximum absolute principal correlation.

Value

An output list object of class `uc.lars`:

form	The regression formula specified as the first argument.
data	Name of the data.frame object specified as the second argument.
p	Number of regression predictor variables.

n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics.
gmat	Orthogonal matrix of direction cosines for regressor principal axes.
lars	An object of class lars.
coef	Matrix of shrinkage-ridge regression coefficient estimates.
risk	Matrix of MSE risk estimates for fitted coefficients.
exev	Matrix of excess MSE eigenvalues (ordinary least squares minus ridge.)
infd	Matrix of direction cosines for the estimated inferior direction, if any.
spat	Matrix of shrinkage pattern multiplicative delta factors.
mlik	Listing of criteria for maximum likelihood selection of M-extent-of-shrinkage.
sext	Listing of summary statistics for all M-extents-of-shrinkage.
mClk	Most Likely Extent of Shrinkage Observed: best multiple of (1/steps) \leq p.
minC	Minimum Observed Value of Normal-theory $-2*\log(\text{Likelihood})$.

Author(s)

Bob Obenchain <wizbob@att.net>

References

Efron B, Hastie T. (2005) Least Angle Regression, Lasso and Forward Stagewise. <https://CRAN.R-project.org/package=lars>

Obenchain RL. (1994-2005) Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein. <http://localcontrolstatistics.org>

Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[aug.lars](#).

Examples

```
data(longley2)
form <- GNP~GNP.deflator+Unemployed+Armed.Forces+Population+Year+Employed
rxucobj <- uc.lars(form, data=longley2)
rxucobj
plot(rxucobj)
str(rxucobj)
```

unr.aug	<i>Augment calculations performed by <code>unr.ridge()</code> to prepare for display of elliptical confidence regions for pairs of biased coefficient estimates using <code>plot.unr.biv()</code></i>
---------	---

Description

This function makes classical (rather than Bayesian) Normal distribution-theory calculations of the form proposed in Obenchain(1977). Instead of providing "new" confidence regions for estimable linear functions, Generalized Ridge Regression (GRR) can focus interest on estimates that are within traditional confidence intervals and regions but which deviate reasonably from the interval or region centroid.

Usage

```
unr.aug(uobj)
```

Arguments

`uobj` An output object of class "unr.ridge".

Value

An output list object of class "unr.aug"...

<code>p</code>	Number of regression predictor variables.
<code>LModj</code>	The <code>lm()</code> output object for the model fitted using <code>unr.ridge()</code> .
<code>bstar</code>	The <code>p</code> by <code>p+1</code> matrix of shrunken GRR coefficients. The first <code>p</code> correspond to "knots" on piecewise linear splines, and the last column contains minimum MSE risk estimates.
<code>mcal</code>	<code>p+1</code> increasing measures of shrinkage "Extent". The first is 0 for the OLS (BLUE) estimate, while the last denotes the "Extent" most likely to yield minimum MSE risk.
<code>vnames</code>	All <code>p+1</code> names for variables used in the GRR model.

Author(s)

Bob Obenchain <wizbob@att.net>

References

- Obenchain RL. (1977) Classical F-Tests and Confidence Regions for Ridge Regression. *Technometrics* 19, 429-439. <http://doi.org/10.1080/00401706.1977.10489582>
- Obenchain RL. (2020) *Ridge TRACE Diagnostics*. <https://arxiv.org/abs/2005.14291>
- Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[unr.ridge](#), [mofk](#) and [kofm](#)

unr.biv	<i>Specify pairs of GRR Coefficient Estimates for display in Bivariate Confidence Regions</i>
---------	---

Description

This function specifies which Pair of GRR estimates to display and the single (or dual) Confidence Level(s) of the Ellipse(s) displayed. Requested confidence levels must both be equal to or greater than 0.05 and less than or equal to 0.95.

Usage

```
unr.biv(uraug, x1 = 1, x2 = 2, conf1 = 0.95, conf2 = 0.50)
```

Arguments

uraug	An output list object of class "unr.aug" for a GRR model with $p \geq 2$ x-Variables.
x1	Integer index value ≥ 1 and $\leq p$ for the x-Coefficient displayed on the horizontal axis.
x2	Integer index value ≥ 1 and $\leq p$ for the x-Coefficient displayed on the vertical axis. Index x2 must differ from x1 to display a plot.
conf1	This first Confidence level must be ≥ 0.05 and ≤ 0.95 to display an Ellipse.
conf2	When the second Confidence level is ≥ 0.05 and ≤ 0.95 , its Ellipse is displayed. No plot is displayed when both conf1 and conf2 are outside of the [0.05, 0.95] range.

Value

An output list object of class "unr.biv"...

p	Number of regression predictor variables.
LMobj	The <code>lm()</code> output object for the model fitted using <code>unr.ridge()</code> .
bstar	The p by p+1 matrix of shrunken GRR coefficients. The first p correspond to "knots" on piecewise linear splines, and the last column contains minimum MSE risk estimates.
mcal	p+1 increasing measures of shrinkage "Extent". The first value is 0 for the OLS = BLU estimate, plotted as a Point of color "blue". The last value corresponds to the "Extent" most likely to yield minimum MSE risk, plotted as a Point of color "purple".
ellip1	matrix[100, 2] of points on confidence ellipse 1.
conf1	confidence level of ellipse 1 within [0.05, 0.95].

ecor1	Pearson correlation between x1 and x2 in ellipse 1.
ellip2	matrix[100, 2] of points on confidence ellipse 2.
conf2	confidence level of ellipse 2 within [0.05, 0.95].
ecor2	Pearson correlation between x1 and x2 in ellipse 2.

Author(s)

Bob Obenchain <wizbob@att.net>

References

Obenchain RL. (1977) Classical F-Tests and Confidence Regions for Ridge Regression. *Technometrics* 19, 429-439. <http://doi.org/10.1080/00401706.1977.10489582>

Murdoch DJ. and Chow ED. (1996). A graphical display of large correlation matrices. *The American Statistician* 50, 178-180.

Murdoch DJ. ellipse: Functions for Drawing Ellipses and Ellipse-Like Confidence Regions. <https://CRAN.R-project.org/package=ellipse>

See Also

[unr.aug](#), [ellipse.lm](#), [unr.ridge](#)

unr.ridge	<i>Unrestricted Maximum Likelihood (ML) Shrinkage using a Piecewise Linear-Spline PATH</i>
-----------	--

Description

Compute and display TRACES for the p-parameter Shrinkage PATH passing through the (classical) Normal-theory Maximum Likelihood (ML) point-estimate of the Beta coefficient vector. The m-Extent of overall Optimal Shrinkage corresponding to this solution is marked by a vertical dashed-line on all 5-types of unr.ridge TRACE displays.

Usage

```
unr.ridge(form, data, rscale = 1, steps = 8, delmax = 0.999999)
```

Arguments

form	A regression formula [y~x1+x2+...] suitable for use with lm().
data	Data frame containing observations on all variables in the formula.
rscale	One of three possible choices (0, 1 or 2) for "rescaling" of variables (after being "centered") to remove all "non-essential" ill-conditioning: 0 implies no rescaling; 1 implies divide each variable by its standard error; 2 implies rescale as in option 1 but re-express answers as in option 0.

steps	Number of equally spaced values per unit change along the horizontal M-extent-of-shrinkage axis for estimates to be calculated and displayed in TRACES (default = 8.)
delmax	Maximum allowed value for Shrinkage delta-factors that is strictly less than 1. (default = 0.999999, which prints as 1 when rounded to fewer than 6 decimal places.)

Details

Ill-conditioned and/or nearly multi-collinear regression models are unlikely to produce Ordinary Least Squares (OLS) regression coefficient estimates that are very close, numerically, to their unknown true values. Specifically, OLS estimates can then tend to have "wrong" numerical signs and/or unreasonable relative magnitudes, while shrunken (generalized ridge) estimates chosen to maximize their likelihood of reducing Mean Squared Error (MSE) Risk (expected loss) can be much more stable and reasonable, numerically. On the other hand, because only OLS estimates are guaranteed to be minimax when risk is matrix valued (truly multivariate), no guarantee of an actual reduction in MSE Risk is necessarily associated with shrinkage.

Value

An output list object of class unr.ridge:

data	Name of the data.frame object specified as the second argument.
form	The regression formula specified as the first argument.
p	Number of regression predictor variables.
n	Number of complete observations after removal of all missing values.
r2	Numerical value of R-square goodness-of-fit statistic.
s2	Numerical value of the residual mean square estimate of error.
prinstat	Listing of principal statistics.
rscale	Variable re-scaling code of 0, 1 or 2 used in calculations.
data	data.frame containing all variables listed in the formula.
gmat	Orthogonal Matrix of Direction Cosines for Principal Axes.
coef	Matrix of shrinkage-ridge regression coefficient estimates.
rmse	Matrix of MSE risk estimates for fitted coefficients.
exev	Matrix of excess MSE eigenvalues (ordinary least squares minus ridge.)
infd	Matrix of direction cosines for the estimated inferior direction, if any.
spat	Matrix of shrinkage pattern multiplicative delta factors.
mlik	Listing of criteria for maximum likelihood selection of M-extent-of-shrinkage.
sext	Listing of summary statistics for all M-extents-of-shrinkage.
mUnr	Unrestricted m-Extent of Shrinkage corresponding to $k^* == 1$ on TRACE displays.
mClk	Most Likely Extent of Shrinkage Observed: best multiple of $(1/steps) \leq p$.
minC	Minimum Observed Value of Normal-theory $-2*\log(\text{Likelihood-Ratio})$.
dMSE	Most Likely to be Optimal-values for Shrinkage Delta-factors [1:p].

Author(s)

Bob Obenchain <wizbob@att.net>

References

Thompson JR. (1968) Some shrinkage techniques for estimating the mean. *Journal of the American Statistical Association* **63**, 113-122. (The “cubic” estimator.)

Obenchain RL. (1978) Good and Optimal Ridge Estimators. *Annals of Statistics* **6**, 1111-1121. <doi:10.1214/aos/1176344314>

Obenchain RL. (2005) *Shrinkage Regression: ridge, BLUP, Bayes, spline and Stein*. Electronic book-in-progress; 185+ pages. <http://localcontrolstatistics.org>

Obenchain RL. (2020) *Ridge TRACE Diagnostics*. <https://arxiv.org/abs/2005.14291>

Obenchain RL. (2020) *The Unrestricted Shrinkage Path: Technical Details*. "unrPath.pdf" <http://localcontrolstatistics.org>

Obenchain RL. (2020) *RXshrink_in_R.PDF* RXshrink package vignette-like file. <http://localcontrolstatistics.org>

See Also

[mofk](#), [kofm](#), [correct.signs](#), [MLtrue](#) and [RXpredict](#).

Examples

```
data(longley2)
form <- Employed~GNP+GNP.deflator+Unemployed+Armed.Forces+Population+Year
rxuobj <- unr.ridge(form, data=longley2)
rxuobj
plot(rxuobj)
str(rxuobj)
```

Index

- * **datasets**
 - haldport, 8
 - longley2, 10
 - mpg, 18
 - tycobb, 28
 - * **histogram**
 - MLhist, 14
 - * **package**
 - RXshrink-package, 2
 - * **predict**
 - plot.RXpredict, 20
 - * **regression**
 - aug.lars, 4
 - correct.signs, 6
 - kofm, 9
 - MLboot, 11
 - MLcalc, 12
 - MLtrue, 15
 - mofk, 17
 - qm.ridge, 24
 - RXpredict, 26
 - uc.lars, 29
 - unr.aug, 31
 - unr.biv, 32
 - unr.ridge, 33
 - * **trace**
 - plot.aug.lars, 18
 - plot.qm.ridge, 19
 - plot.uc.lars, 21
 - plot.unr.biv, 22
 - plot.unr.ridge, 23
- aug.lars, 4, 30
- correct.signs, 6, 12, 26, 35
- ellipse.lm, 33
- haldport, 8
- kofm, 9, 17, 32, 35
- longley2, 10
- MLboot, 11, 13, 16
- MLcalc, 12
- MLhist, 14
- MLtrue, 8, 15, 26, 27, 35
- mofk, 10, 17, 32, 35
- mpg, 18
- plot.aug.lars, 18
- plot.qm.ridge, 19
- plot.RXpredict, 20
- plot.uc.lars, 21
- plot.unr.biv, 22
- plot.unr.ridge, 23
- qm.ridge, 8, 16, 24, 27
- RXpredict, 26, 26, 35
- RXshrink-package, 2
- tycobb, 28
- uc.lars, 6, 29
- unr.aug, 31, 33
- unr.biv, 32
- unr.ridge, 8, 10, 12, 13, 16, 17, 26, 27, 32, 33, 33