Package 'sgd'

October 21, 2025

```
Version 1.1.3
Maintainer Junhyung Lyle Kim <jlylekim@gmail.com>
Description A fast and flexible set of tools for large scale estimation. It
     features many stochastic gradient methods, built-in models, visualization
     tools, automated hyperparameter tuning, model checking, interval estimation,
     and convergence diagnostics.
URL https://github.com/airoldilab/sgd
BugReports https://github.com/airoldilab/sgd/issues
License GPL-2
Imports ggplot2, MASS, methods, Rcpp (>= 0.11.3), stats
Suggests bigmemory, glmnet, gridExtra, R.rsp, testthat, microbenchmark
LinkingTo BH, bigmemory, Rcpp, RcppArmadillo
LazyData yes
VignetteBuilder R.rsp
Encoding UTF-8
RoxygenNote 7.3.3
NeedsCompilation yes
Author Junhyung Lyle Kim [cre, aut],
     Dustin Tran [aut],
     Panos Toulis [aut],
     Tian Lian [ctb],
     Ye Kuang [ctb],
     Edoardo Airoldi [ctb]
Repository CRAN
Date/Publication 2025-10-21 11:00:02 UTC
```

Title Stochastic Gradient Descent for Scalable Estimation

Type Package

2 fitted.sgd

Contents

coef.sgd	 	 	 	2
$fitted.sgd \dots \dots \\$	 	 	 	2
plot.sgd	 	 	 	3
predict.sgd	 	 	 	4
•				
•				
winequality	 	 	 	-

Index 10

 $\mathsf{coef}.\mathsf{sgd}$

Extract Model Coefficients

Description

Extract model coefficients from sgd objects. coefficients is an *alias* for it.

Usage

```
## S3 method for class 'sgd'
coef(object, ...)
```

Arguments

object of class sgd.

... some methods for this generic require additional arguments. None are used in

this method.

Value

Coefficients extracted from the model object object.

fitted.sgd

Extract Model Fitted Values

Description

Extract fitted values from from sgd objects. fitted.values is an alias for it.

Usage

```
## S3 method for class 'sgd'
fitted(object, ...)
```

plot.sgd 3

Arguments

object of class sgd.

... some methods for this generic require additional arguments. None are used in

this method.

Value

Fitted values extracted from the object object.

plot.sgd

Plot objects of class sgd.

Description

Plot objects of class sgd.

Usage

```
## S3 method for class 'sgd'
plot(x, ..., type = "mse", xaxis = "iteration")
## S3 method for class 'list'
plot(x, ..., type = "mse", xaxis = "iteration")
```

Arguments

x object of class sgd.

additional arguments used for each type of plot. See 'Details'.

type character specifying the type of plot: "mse", "clf", "mse-param". See 'De-

tails'. Default is "mse".

xaxis character specifying the x-axis of plot: "iteration" plots the y values over the

log-iteration of the algorithm; "runtime" plots the y values over the time in

seconds to reach them. Default is "iteration".

Details

Types of plots available:

mse Mean squared error in predictions, which takes the following arguments:

x_test_test_set

y_test test responses to compare predictions to

clf Classification error in predictions, which takes the following arguments:

x_test_test_set

y_test test responses to compare predictions to

mse-param Mean squared error in parameters, which takes the following arguments:

true_param true vector of parameters to compare to

4 print.sgd

predict.sgd

Model Predictions

Description

Form predictions using the estimated model parameters from stochastic gradient descent.

Usage

```
## S3 method for class 'sgd'
predict(object, newdata, type = "link", ...)
predict_all(object, newdata, ...)
```

Arguments

object

object of class sgd.

newdata

design matrix to form predictions on

type

the type of prediction required. The default "link" is on the scale of the linear predictors; the alternative '"response"' is on the scale of the response variable. Thus for a default binomial model the default predictions are of log-odds (probabilities on logit scale) and 'type = "response"' gives the predicted probabilities. The '"terms"' option returns a matrix giving the fitted values of each term in the

model formula on the linear predictor scale.

. . .

further arguments passed to or from other methods.

Details

A column of 1's must be included to newdata if the parameters include a bias (intercept) term.

print.sgd

Print objects of class sgd.

Description

Print objects of class sgd.

Usage

```
## S3 method for class 'sgd'
print(x, ...)
```

Arguments

x object of class sgd.

... further arguments passed to or from other methods.

residuals.sgd 5

residuals.sgd

Extract Model Residuals

Description

Extract model residuals from sgd objects. resid is an alias for it.

Usage

```
## S3 method for class 'sgd'
residuals(object, ...)
```

Arguments

object of class sgd.

... some methods for this generic require additional arguments. None are used in

this method.

Value

Residuals extracted from the object object.

sgd

Stochastic gradient descent

Description

Run stochastic gradient descent in order to optimize the induced loss function given a model and data.

Usage

```
sgd(x, ...)
## S3 method for class 'formula'
sgd(formula, data, model, model.control = list(), sgd.control = list(...), ...)
## S3 method for class 'matrix'
sgd(x, y, model, model.control = list(), sgd.control = list(...), ...)
## S3 method for class 'big.matrix'
sgd(x, y, model, model.control = list(), sgd.control = list(...), ...)
```

6 sgd

Arguments

a design matrix and the respective vector of outcomes. x, y

arguments to be used to form the default sgd.control arguments if it is not . . .

supplied directly.

formula an object of class "formula" (or one that can be coerced to that class): a sym-

bolic description of the model to be fitted. The details can be found in "glm".

data an optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables in the model. If not found in data, the

variables are taken from environment(formula), typically the environment from

which glm is called.

character specifying the model to be used: "lm" (linear model), "glm" (generalized linear model), "cox" (Cox proportional hazards model), "gmm" (generalized

method of moments), "m" (M-estimation). See 'Details'.

model.control a list of parameters for controlling the model.

> family ("glm") a description of the error distribution and link function to be used in the model. This can be a character string naming a family function, a family function or the result of a call to a family function. (See family for details of family functions.)

rank ("glm") logical. Should the rank of the design matrix be checked?

fn ("gmm") a function $g(\theta, x)$ which returns a k-vector corresponding to the k moment conditions. It is a required argument if gr not specified.

gr ("gmm") a function to return the gradient. If unspecified, a finite-difference approximation will be used.

nparams ("gmm") number of model parameters. This is automatically determined for other models.

type ("gmm") character specifying the generalized method of moments procedure: "twostep" (Hansen, 1982), "iterative" (Hansen et al., 1996). Defaults to "iterative".

wmatrix ("gmm") weighting matrix to be used in the loss function. Defaults to the identity matrix.

loss ("m") character specifying the loss function to be used in the estimating equation. Default is the Huber loss.

lambda1 L1 regularization parameter. Default is 0.

1ambda2 L2 regularization parameter. Default is 0.

an optional list of parameters for controlling the estimation.

method character specifying the method to be used: "sgd", "implicit", "asgd", "ai-sgd", "momentum", "nesterov". Default is "ai-sgd". See 'Details'.

1r character specifying the learning rate to be used: "one-dim", "one-dim-eigen", "d-dim", "adagrad", "rmsprop". Default is "one-dim". See 'Details'.

1r.control vector of scalar hyperparameters one can set dependent on the learning rate. For hyperparameters aimed to be left as default, specify NA in the corresponding entries. See 'Details'.

start starting values for the parameter estimates. Default is random initialization around zero.

mode1

sgd.control

sgd 7

size number of SGD estimates to store for diagnostic purposes (distributed log-uniformly over total number of iterations)

reltol relative convergence tolerance. The algorithm stops if it is unable to change the relative mean squared difference in the parameters by more than the amount. Default is 1e-05.

npasses the maximum number of passes over the data. Default is 3.

pass logical. Should tol be ignored and run the algorithm for all of npasses? shuffle logical. Should the algorithm shuffle the data set including for each pass?

verbose logical. Should the algorithm print progress?

Details

Models: The Cox model assumes that the survival data is ordered when passed in, i.e., such that the risk set of an observation i is all data points after it.

Methods:

```
sgd stochastic gradient descent (Robbins and Monro, 1951)
implicit implicit stochastic gradient descent (Toulis et al., 2014)
asgd stochastic gradient with averaging (Polyak and Juditsky, 1992)
ai-sgd implicit stochastic gradient with averaging (Toulis et al., 2015)
momentum "classical" momentum (Polyak, 1964)
nesterov Nesterov's accelerated gradient (Nesterov, 1983)
```

Learning rates and hyperparameters:

one-dim scalar value prescribed in Xu (2011) as

$$a_n = scale * gamma/(1 + alpha * gamma * n)^{(-c)}$$

where the defaults are lr.control = (scale=1, gamma=1, alpha=1, c) where c is 1 if implemented without averaging, 2/3 if with averaging

```
one-dim-eigen diagonal matrix lr.control = NULL
```

d-dim diagonal matrix lr.control = (epsilon=1e-6)

adagrad diagonal matrix prescribed in Duchi et al. (2011) as lr.control = (eta=1, epsilon=1e-6)

rmsprop diagonal matrix prescribed in Tieleman and Hinton (2012) as lr.control = (eta=1, gamma=0.9, epsilon=1e-6)

Value

An object of class "sgd", which is a list containing the following components:

model name of the model

coefficients a named vector of coefficients

converged logical. Was the algorithm judged to have converged?

estimates estimates from algorithm stored at each iteration specified in pos

8 sgd

fitted.values the fitted mean values

pos vector of indices specifying the iteration number each estimate was stored for

residuals the residuals, that is response minus fitted values

times vector of times in seconds it took to complete the number of iterations specified

in pos

model.out a list of model-specific output attributes

Author(s)

Dustin Tran, Tian Lan, Panos Toulis, Ye Kuang, Edoardo Airoldi

References

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121-2159, 2011.

Yurii Nesterov. A method for solving a convex programming problem with convergence rate $O(1/k^2)$. Soviet Mathematics Doklady, 27(2):372-376, 1983.

Boris T. Polyak. Some methods of speeding up the convergence of iteration methods. *USSR Computational Mathematics and Mathematical Physics*, 4(5):1-17, 1964.

Boris T. Polyak and Anatoli B. Juditsky. Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization*, 30(4):838-855, 1992.

Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400-407, 1951.

Panos Toulis, Jason Rennie, and Edoardo M. Airoldi, "Statistical analysis of stochastic gradient methods for generalized linear models", In *Proceedings of the 31st International Conference on Machine Learning*, 2014.

Panos Toulis, Dustin Tran, and Edoardo M. Airoldi, "Stability and optimality in stochastic gradient descent", arXiv preprint arXiv:1505.02417, 2015.

Wei Xu. Towards optimal one pass large scale learning with averaged stochastic gradient descent. arXiv preprint arXiv:1107.2490, 2011.

Dimensions

Examples

```
## Linear regression
set.seed(42)
N <- 1e4
d <- 5
X <- matrix(rnorm(N*d), ncol=d)
theta <- rep(5, d+1)
eps <- rnorm(N)
y <- cbind(1, X) %*% theta + eps
dat <- data.frame(y=y, x=X)
sgd.theta <- sgd(y ~ ., data=dat, model="lm")
sprintf("Mean squared error: %0.3f", mean((theta - as.numeric(sgd.theta$coefficients))^2))</pre>
```

winequality 9

winequality

Wine quality data of white wine samples from Portugal

Description

This dataset is a collection of white "Vinho Verde" wine samples from the north of Portugal. Due to privacy and logistic issues, only physicochemical (inputs) and sensory (the output) variables are available (e.g. there is no data about grape types, wine brand, wine selling price, etc.).

Usage

winequality

Format

A data frame with 4898 rows and 12 variables

- · fixed acidity.
- volatile acidity.
- citric acid.
- · residual sugar.
- chlorides.
- free sulfur dioxide.
- total sulfur dioxide.
- · density.
- pH.
- sulphates.
- alcohol.
- quality (score between 0 and 10).

Source

https://archive.ics.uci.edu/ml/datasets/Wine+Quality

Index

```
\ast datasets
    winequality, 9
as.data.frame,6
\mathsf{coef.sgd}, 2
family, 6
\texttt{fitted.sgd}, \textcolor{red}{2}
formula, 6
glm, 6
plot.list(plot.sgd), 3
plot.sgd, 3
predict.sgd, 4
predict_all (predict.sgd), 4
print.sgd, 4
residuals.sgd, 5
sgd, 5
winequality, 9
```