

Package 'h2o'

March 5, 2014

R topics documented:

h2o-package	1
apply	2
as.data.frame.H2OParsedData	3
as.factor	4
as.h2o	5
colnames	6
Extremes	6
h2o.anyFactor	7
h2o.assign	8
h2o.clusterInfo	9
h2o.clusterStatus	9
h2o.confusionMatrix	10
h2o.cut	11
h2o.downloadAllLogs	12
h2o.downloadCSV	13
h2o.exportHDFS	14
h2o.gbm	14
h2o.glm	16
h2o.importFile	18
h2o.importFolder	20
h2o.importHDFS	22
h2o.importURL	23
h2o.init	25
h2o.kmeans	26
h2o.logAndEcho	28
h2o.ls	29
h2o.nn	29
h2o.parseRaw	31
h2o.pcr	32
h2o.prcomp	34
h2o.predict	35
h2o.randomForest	36
h2o.rm	38

h2o.runif	39
h2o.shutdown	40
h2o.table	41
h2o.uploadFile	42
h2o.__changeLogPath	43
h2o.__clearLogs	44
h2o.__getLogPath	45
h2o.__openLog	45
h2o.__startLogging	46
h2o.__stopLogging	47
H2OClient-class	47
H2ODRFGrid-class	48
H2ODRFModel-class	49
H2OGBMGrid-class	50
H2OGBMModel-class	51
H2OGLMGrid-class	52
H2OGLMGridVA-class	53
H2OGLMModel-class	54
H2OGLMModelVA-class	55
H2OGrid-class	56
H2OGridVA-class	56
H2OKMeansGrid-class	57
H2OKMeansModel-class	58
H2OKMeansModelVA-class	59
H2OModel-class	60
H2OModelVA-class	60
H2ONNGrid-class	61
H2ONNModel-class	62
H2OParsedData-class	63
H2OParsedDataVA-class	66
H2OPCAModel-class	67
H2ORawData-class	68
H2ORawDataVA-class	68
H2ORFModelVA-class	69
head	70
ifelse	71
is.factor	72
levels	72
mean.H2OParsedData	73
nrow	74
quantile.H2OParsedData	74
screeplot.H2OPCAModel	75
sd	76
str	77
sum	77
summary	78
summary.H2OPCAModel	79

h2o-package*H2O R Interface*

Description

This is a package for running H2O via its REST API from within R. To communicate with a H2O instance, the version of the R package must match the version of H2O. When connecting to a new H2O cluster, it is necessary to re-run the initializer.

Details

Package: h2o
Type: Package
Version: 2.2.1.3
Date: 2013-12-12
License: Apache License (== 2.0)
Depends: R (>= 2.13.0), RCurl, rjson, statmod, tools, methods, utils

This package allows the user to run basic H2O commands using R commands. In order to use it, you must first have H2O running (See [How to Start H2O](#)). To run H2O on your local machine, call `h2o.init` without any arguments, and H2O will be automatically launched on <http://127.0.0.1:54321>, where the IP is "127.0.0.1" and the port is 54321. If H2O is running on a cluster, you must provide the IP and port of the remote machine as arguments to the `h2o.init()` call.

H2O supports a number of standard statistical models, such as GLM, K-means, and Random Forest classification. For example, to run GLM, call `h2o.glm` with the H2O parsed data and parameters (response variable, error distribution, etc...) as arguments. (The operation will be done on the server associated with the data object where H2O is running, not within the R environment).

Note that no actual data is stored in the R workspace; and no actual work is carried out by R. R only saves the named objects, which uniquely identify the data set, model, etc on the server. When the user makes a request, R queries the server via the REST API, which returns a JSON file with the relevant information that R then displays in the console.

Author(s)

Anqi Fu, Tom Kraljevic and Petr Maj, with contributions from the 0xdata team

Maintainer: Anqi Fu <anqi@0xdata.com>

References

- [0xdata Homepage](#)
- [H2O Documentation](#)
- [H2O on Github](#)

Examples

```
# Check connection with H2O and ensure local H2O R package matches server version.
# Optionally, ask for startH2O to start H2O if it's not already running.
# Note that for startH2O to work, the IP must be 127.0.0.1 or localhost with port 54321.
library(h2o)
localH2O = h2o.init(ip = "127.0.0.1", port = 54321, startH2O = TRUE)

# Import iris dataset into H2O and print summary
irisPath = system.file("extdata", "iris.csv", package = "h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
summary(iris.hex)

# Attach H2O R package and run GLM demo
??h2o
demo(package = "h2o")
demo(h2o.glm)

# Shutdown local H2O instance when finished
h2o.shutdown(localH2O)
```

apply	<i>Applies a function over an H2O parsed data object.</i>
-------	---

Description

Applies a function over an H2O parsed data object (an array).

Usage

```
apply(X, MARGIN, FUN, ...)
```

Arguments

X	An H2OParsedData object.
MARGIN	The margin along which the function should be applied
FUN	The function to be applied by H2O.
...	Optional arguments to FUN. (Currently unimplemented).

Value

Produces a new [H2OParsedData](#) of the output of the applied function. The output is stored in H2O so that it can be used in subsequent H2O processes.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
summary(apply(iris.hex, 1, sum))
h2o.shutdown(localH2O)
```

```
as.data.frame.H2OParsedData
```

Converts a parsed H2O object to a data frame.

Description

Convert an [H2OParsedData](#) object to a data frame, which allows subsequent data frame operations within the R environment.

Usage

```
## S3 method for class 'H2OParsedData'
as.data.frame(x, ...)
```

Arguments

x	An H2OParsedData object.
...	Additional arguments to be passed to or from methods.

Value

Returns a data frame in the R environment. Note that this call establishes the data set in the R environment, and subsequent operations on the data frame take place within R, not H2O. When data are large, users may experience

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
prostate.data.frame <- as.data.frame(prostate.hex)
summary(prostate.data.frame)
head(prostate.data.frame)
h2o.shutdown(localH2O)
```

as.factor	<i>Converts a column from numeric to factor</i>
-----------	---

Description

Specify a column type to be factor (also called categorical or enumerative), rather than numeric.

Usage

```
as.factor(x)
```

Arguments

x	A column in an object of class H2OParsedData , or data frame.
---	---

Value

Returns the original object of class [H2OParsedData](#), with the requested column specified as a factor, rather than numeric.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
prostate.hex[,4] = as.factor(prostate.hex[,4])
summary(prostate.hex)
h2o.shutdown(localH2O)
```

as.h2o	<i>Converts an R object to an H2O object</i>
--------	--

Description

Convert an R object to an H2O object, copy contents of the object to the running instance of H2O

Usage

```
as.h2o(client, object, key = "", header, sep = "")
```

Arguments

client	The h2o.init object that facilitates communication between R and H2O.
object	The object in the R environment to be converted to an H2O object.
key	(Optional) A reference assigned to the object in the instance of H2O (the key part of the key-value store, where the value is the R object to be converted.)
header	(Optional) A logical value indicating whether the first line of the file contains column headers. If left empty, the parser will try to automatically detect this.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If sep = "", the parse

Details

The R object to be converted to an H2O object should be named so that it can be used in subsequent analysis. Also note that the R object is converted to a parsed H2O data object, and will be treated as a data frame by H2O in subsequent analysis.

Value

Converts an R object to an H2O Parsed data object.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)

data(iris)
summary(iris)
iris.r <- iris
iris.h2o <- as.h2o(localH2O, iris.r, key="iris.h2o")
class(iris.h2o)
h2o.shutdown(localH2O)
```

colnames	<i>Returns column names for a parsed H2O data object.</i>
----------	---

Description

Returns column names for an [H2OParsedData](#) object.

Usage

```
colnames(x, do.NULL = TRUE, prefix = "col")
```

Arguments

x	An H2OParsedData object.
do.NULL	Logical value. If FALSE and names are NULL, names are created.
prefix	Character string denoting prefix for created column names.

Value

Returns a vector of column names.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
summary(iris.hex)
colnames(iris.hex)
h2o.shutdown(localH2O)
```

Extremes

Maxima and Minima

Description

Calculates the (parallel) minimum of the input values. This method extends the [min](#) generic to deal with [H2OParsedData](#) objects.

Usage

```
max(..., na.rm = FALSE)
min(..., na.rm = FALSE)
```

Arguments

<code>...</code>	Numeric, character or H2OParsedData arguments.
<code>na.rm</code>	Logical value where FALSE does not remove NA's in the calculation and TRUE removes NA's in the calculation.

Value

Returns the maximum or minimum over all the input arguments. For a [H2OParsedData](#) object, the function is calculated over all entries in the dataset. An error will occur if any of those entries is non-numeric.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath, key = "australia.hex")
min(australia.hex)
min(c(-1, 0.5, 0.2), FALSE, australia.hex[,1:4])
h2o.shutdown(localH2O)
```

h2o.anyFactor	<i>Determine if an H2O parsed data object contains categorical data.</i>
---------------	--

Description

Checks if an H2O parsed data object has any columns of categorical data.

Usage

```
h2o.anyFactor(x)
```

Arguments

x An [H2OParsedData](#) object.

Value

Returns a logical value indicating whether any of the columns in x are factors.

See Also

[H2OParsedData](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris_wheader.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath)
h2o.anyFactor(iris.hex)
h2o.shutdown(localH2O)
```

h2o.assign	<i>Assigns an H2O hex.key to an H2O object so that it can be utilized in subsequent calls</i>
------------	---

Description

Allows users to assign H2O hex.keys to objects in their R environment so that they can manipulate H2O data frames and parsed data objects.

Usage

```
h2o.assign(data, key)
```

Arguments

data	An H2OParsedData object
key	The hex key to be associated with the H2O parsed data object

Value

The function returns an object of class [H2OParsedData](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package = "h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
psa.qs = quantile(prostate.hex$PSA)
PSA.outliers = prostate.hex[prostate.hex$PSA <= psa.qs[2] | prostate.hex$PSA >= psa.qs[10],]
PSA.outliers = h2o.assign(PSA.outliers, "PSA.outliers")
summary(PSA.outliers)
head(prostate.hex)
head(PSA.outliers)
h2o.shutdown(localH2O)
```

h2o.clusterInfo

Get Information on H2O Cluster

Description

Display the name, version, uptime, total nodes, total memory, total cores and health of a cluster running H2O.

Usage

```
h2o.clusterInfo(client)
```

Arguments

client	An H2OClient object containing the IP address and port of the server running H2O.
--------	---

See Also

[H2OClient](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.clusterInfo(localH2O)
h2o.shutdown(localH2O)
```

h2o.clusterStatus	<i>Retrieve Status of H2O Cluster</i>
-------------------	---------------------------------------

Description

Retrieve information on the status of the cluster running H2O.

Usage

```
h2o.clusterStatus(client)
```

Arguments

client	An H2OClient object containing the IP address and port of the server running H2O.
--------	---

Details

This method prints the status of the H2O cluster represented by `client`, consisting of the following information:

- Version: The version of H2O running on the cluster.
- Cloud Name: Name of the cluster.
- Node Name: Name of the node. (Defaults to the HTTP address).
- Cloud Size: Number of nodes in the cluster.

Furthermore, for each node, this function displays:

- name: Name of the node.
- value_size_bytes: Amount of data stored on the node.
- free_mem_bytes: Amount of free memory on the JVM.
- max_mem_bytes: Maximum amount of memory that the JVM will attempt to use.
- free_disk_bytes: Amount of free space on the disk that launched H2O.
- max_disk_bytes: Size of disk that launched H2O.
- num_cpus: Number of CPUs reported by JVM.
- system_load: Average system load.
- rpcs: Number of remote procedure calls.
- last_contact: Number of seconds since last heartbeat.

See Also

[H2OClient](#), [h2o.init](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.clusterStatus(localH2O)
h2o.shutdown(localH2O)
```

h2o.confusionMatrix	<i>Build a Confusion Matrix from H2O Classification Predictions</i>
---------------------	---

Description

Constructs a confusion matrix from a column of predicted responses and a column of actual (reference) responses in H2O. Note that confusion matrices describe prediction errors for classification data only.

Usage

```
h2o.confusionMatrix(data, reference)
```

Arguments

data	An H2OParsedData object that represents the predicted response values. (Must be a single column).
reference	An H2OParsedData object that represents the actual response values. Must have the same dimensions as data.

Value

Returns a confusion matrix with the actual value counts along the rows and the predicted value counts along the columns.

See Also

[H2OParsedData](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
prostate.gbm = h2o.gbm(x = 3:9, y = 2, data = prostate.hex)
prostate.pred = h2o.predict(prostate.gbm)
h2o.confusionMatrix(prostate.pred[,1], prostate.hex[,2])
h2o.shutdown(localH2O)
```

h2o.cut

*Convert H2O Numeric Data to Factor***Description**

Divides the range of the H2O data into intervals and codes the values according to which interval they fall in. The leftmost interval corresponds to level one, the next is level two, etc.

Usage

```
h2o.cut(x, breaks)
```

Arguments

x An [H2OParsedData](#) object with numeric columns.

breaks A numeric vector of two or more unique cut points.

Value

A [H2OParsedData](#) object containing the factored data with intervals as levels.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris_wheader.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
summary(iris.hex)

# Cut sepal length column into intervals determined by min/max/quantiles
sepal_len.cut = h2o.cut(iris.hex$sepal_len, c(4.2, 4.8, 5.8, 6, 8))
head(sepal_len.cut)
summary(sepal_len.cut)
h2o.shutdown(localH2O)
```

h2o.downloadAllLogs

*Download H2O Log Files to Disk***Description**

Download all H2O log files to local disk. Generally used for debugging purposes.

Usage

```
h2o.downloadAllLogs(client, dirname = ".", filename = NULL)
```

Arguments

client	An H2OClient object containing the IP address and port of the server running H2O.
dirname	(Optional) A character string indicating the directory that the log file should be saved in.
filename	(Optional) A character string indicating the name that the log file should be saved to.

See Also

[H2OClient](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.downloadAllLogs(localH2O, dirname = getwd(), filename = "h2o_logs.log")
file.info(paste(getwd(), "h2o_logs.log", sep = .Platform$file.sep))
file.remove(paste(getwd(), "h2o_logs.log", sep = .Platform$file.sep))
h2o.shutdown(localH2O)
```

h2o.downloadCSV	<i>Download H2O Data to Disk</i>
-----------------	----------------------------------

Description

Download a H2O dataset to a CSV file on local disk.

Usage

```
h2o.downloadCSV(data, filename)
```

Arguments

data	An H2OParsedData object to be downloaded.
filename	A character string indicating the name that the CSV file should be saved to.

Details

This method requires wget or curl to be installed on your local system. WARNING: Files located on the H2O server may be very large! Make sure you have enough hard drive space to accommodate the entire file.

See Also

[H2OParsedData](#), [H2OParsedDataVA](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris_wheader.csv", package = "h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath)

myFile = paste(getwd(), "my_iris_file.csv", sep = .Platform$file.sep)
h2o.downloadCSV(iris.hex, myFile)
file.info(myFile)
file.remove(myFile)
h2o.shutdown(localH2O)
```

h2o.exportHDFS	<i>Export a H2O Model to HDFS</i>
----------------	-----------------------------------

Description

Saves a model built from a H2O algorithm to HDFS.

Usage

```
h2o.exportHDFS(object, path)
```

Arguments

object	An H2OModel object representing the model to be exported.
path	The HDFS file path where the model should be saved.

See Also

[H2OModel](#)

Examples

```
## Not run:
# This is an example of how to export H2O models to HDFS.
# The user must modify the path to his or her specific HDFS path for this example to run.
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris_wheader.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath)
iris.gbm = h2o.gbm(x = 1:4, y = 5, data = iris.hex)
h2o.exportHDFS(iris.gbm, path = "hdfs://192.168.1.161/datasets/models")
h2o.shutdown(localH2O)

## End(Not run)
```

h2o.gbm

*H2O: Gradient Boosted Machines***Description**

Builds gradient boosted classification trees, and gradient boosted regression trees on a parsed data set.

Usage

```
h2o.gbm(x, y, distribution = "multinomial", data, n.trees = 10, interaction.depth = 5,
        n.minobsinnode = 10, shrinkage = 0.1, n.bins = 100, validation)
```

Arguments

x	A vector containing the names or indices of the predictor variables to use in building the GBM model.
y	The name or index of the response variable. If the data does not contain a header, this is the column index number starting at 0, and increasing from left to right. (The response must be either an integer or a categorical variable).
distribution	The type of GBM model to be produced: classification is "multinomial" (default), "gaussian" is used for regression.
data	An H2OParsedData object containing the variables in the model.
n.trees	(Optional) Number of trees to grow. Must be a nonnegative integer.
interaction.depth	(Optional) Maximum depth to grow the tree.
n.minobsinnode	(Optional) Minimum number of rows to assign to terminal nodes.
shrinkage	(Optional) A learning-rate parameter defining step size reduction.
n.bins	(Optional) Number of bins to use in building histogram.
validation	(Optional) An H2OParsedData object indicating the validation dataset used to construct confusion matrix. If left blank, this defaults to the training data.

Value

An object of class [H2OGBMModel](#) with slots key, data, and model, where the last is a list of the following components:

type	The type of the tree.
n.trees	Number of trees grown.
oob_err	Out of bag error rate.
forest	A matrix giving the minimum, mean, and maximum of the tree depth and number of leaves.
confusion	Confusion matrix of the prediction when classification model is specified.

References

1. Elith, Jane, John R Leathwick, and Trevor Hastie. "A Working Guide to Boosted Regression Trees." *Journal of Animal Ecology* 77.4 (2008): 802-813
2. Friedman, Jerome, Trevor Hastie, Saharon Rosset, Robert Tibshirani, and Ji Zhu. "Discussion of Boosting Papers." *Ann. Statist* 32 (2004): 102-107
3. Hastie, Trevor, Robert Tibshirani, and J Jerome H Friedman. *The Elements of Statistical Learning*. Vol.1. N.p.: Springer New York, 2001. http://www.stanford.edu/~hastie/local.ftp/Springer/OLD//ESLII_print4.pdf

See Also

For more information see: <http://docs.0xdata.com>

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)

# Run regression GBM on australia.hex data
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
independent <- c("premax", "salmax", "minairtemp", "maxairtemp", "maxsst",
  "maxsoilmoist", "Max_czcs")
dependent <- "runoffnew"
h2o.gbm(y = dependent, x = independent, data = australia.hex, n.trees = 10, interaction.depth = 3,
  n.minobsinnode = 2, shrinkage = 0.2, distribution= "gaussian")

# Run multinomial classification GBM on australia data
h2o.gbm(y = dependent, x = independent, data = australia.hex, n.trees = 15, interaction.depth = 5,
  n.minobsinnode = 2, shrinkage = 0.01, distribution= "multinomial")
h2o.shutdown(localH2O)
```

h2o.glm

H2O: Generalized Linear Models

Description

Fit a generalized linear model, specified by a response variable, a set of predictors, and a description of the error distribution.

Usage

```
## Default method:
h2o.glm(x, y, data, family, nfolds = 10, alpha = 0.5, lambda = 1e-5, epsilon = 1e-4,
  standardize = TRUE, tweedie.p = ifelse(family == 'tweedie', 1.5,
  as.numeric(NA)), thresholds, version = 1)

## Import to a ValueArray object:
```

```
h2o.glm.VA(x, y, data, family, nfolds = 10, alpha = 0.5, lambda = 1e-5, epsilon = 1e-4,
  standardize = TRUE, tweedie.p = ifelse(family == 'tweedie', 1.5,
  as.numeric(NA)), thresholds = ifelse(family=='binomial', seq(0, 1, 0.01),
  as.numeric(NA)))
```

```
## Import to a FluidVecs object:
```

```
h2o.glm.FV(x, y, data, family, nfolds = 10, alpha = 0.5, lambda = 1e-5, epsilon = 1e-4,
  standardize = TRUE, tweedie.p = ifelse(family == 'tweedie', 1.5,
  as.numeric(NA)))
```

Arguments

x	A vector containing the names of the predictors in the model.
y	The name of the response variable in the model.
data	An H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing the variables in the model.
family	A description of the error distribution and corresponding link function to be used in the model. Currently, Gaussian, binomial, Poisson, gamma, and Tweedie are supported. When a model is specified as Tweedie, users must also specify the appropriate Tweedie power.
nfolds	(Optional) Number of folds for cross-validation. The default is 10.
alpha	(Optional) The elastic-net mixing parameter, which must be in [0,1]. The penalty is defined to be

$$P(\alpha, \beta) = (1 - \alpha)/2 \|\beta\|_2^2 + \alpha \|\beta\|_1 = \sum_j [(1 - \alpha)/2 \beta_j^2 + \alpha |\beta_j|]$$

so alpha=1 is the lasso penalty, while alpha=0 is the ridge penalty.

lambda	The shrinkage parameter, which multiples $P(\alpha, \beta)$ in the objective. The larger lambda is, the more the coefficients are shrunk toward zero (and each other).
epsilon	(Optional) Number indicating the cutoff for determining if a coefficient is zero.
standardize	(Optional) Logical value indicating whether the data should be standardized (set to mean = 0, variance = 1) before running GLM.
tweedie.p	(Optional) The index of the power variance function for the tweedie distribution. Only used if family = "tweedie".
thresholds	(Optional) Degree to which to weight the sensitivity (the proportion of correctly classified 1s) and specificity (the proportion of correctly classified 0s). The default option is joint optimization for the overall classification rate. Changing this will alter the confusion matrix and the AUC. Only used if family = "binomial".
version	(Optional) The version of GLM to run. If version = 1, this will run the more stable ValueArray implementation, while version = 2 runs the faster, but still beta stage FluidVecs implementation.

Details

IMPORTANT: Currently, to run GLM with `version = 1`, you must import data to a `ValueArray` object using `h2o.importFile.VA`, `h2o.importFolder.VA` or one of its variants. To run with `version = 2`, you must import data to a `FluidVecs` object using `h2o.importFile.FV`, `h2o.importFolder.FV` or one of its variants.

Value

An object of class `H2OGLMModelVA` (`version = 1`) or `H2OGLMModel` (`version = 2`) with slots `key`, `data`, `model` and `xval`. The slot `model` is a list of the following components:

<code>coefficients</code>	A named vector of the coefficients estimated in the model.
<code>rank</code>	The numeric rank of the fitted linear model.
<code>family</code>	The family of the error distribution.
<code>deviance</code>	The deviance of the fitted model.
<code>aic</code>	Akaike's Information Criterion for the final computed model.
<code>null.deviance</code>	The deviance for the null model.
<code>iter</code>	Number of algorithm iterations to compute the model.
<code>df.residual</code>	The residual degrees of freedom.
<code>df.null</code>	The residual degrees of freedom for the null model.
<code>y</code>	The response variable in the model.
<code>x</code>	A vector of the predictor variable(s) in the model.
<code>auc</code>	Area under the curve.
<code>training.err</code>	Average training error.
<code>threshold</code>	Best threshold.
<code>confusion</code>	Confusion matrix.

The slot `xval` is a list of `H2OGLMModel` or `H2OGLMModelVA` objects representing the cross-validation models. (Each of these objects themselves has `xval` equal to an empty list).

See Also

[h2o.importFile](#), [h2o.importFolder](#), [h2o.importHDFS](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)

# Run GLM of CAPSULE ~ AGE + RACE + PSA + DCAPS
prostate.hex = h2o.importURL(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smalldata/logreg/prostate.csv", sep = "/"), key = "prostate.hex")
h2o.glm(y = "CAPSULE", x = c("AGE", "RACE", "PSA", "DCAPS"), data = prostate.hex, family = "binomial",
  nfolds = 10, alpha = 0.5)
# Run GLM of VOL ~ CAPSULE + AGE + RACE + PSA + GLEASON
myX = setdiff(colnames(prostate.hex), c("ID", "DPROS", "DCAPS", "VOL"))
```

```

h2o.glm(y = "VOL", x = myX, data = prostate.hex, family = "gaussian", nfolds = 5, alpha = 0.1)

# Run GLM of IsDepDelayed ~ Distance + Origin + Dest + UniqueCarrier on airlines dataset in HDFS
airlines.hex = h2o.importURL(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smldata/airlines/AirlinesTrain.csv.zip", sep = "/"), version = 2)
h2o.glm(x = c('Distance', 'Origin', 'Dest', 'UniqueCarrier'), y = 'IsDepDelayed',
  family = 'binomial', data = airlines.hex, version = 2)
h2o.shutdown(localH2O)

```

h2o.importFile	<i>Import Local Data File</i>
----------------	-------------------------------

Description

Imports a file from the local path and parses it, returning an object containing the identifying hex key.

Usage

```

## Default method:
h2o.importFile(object, path, key = "", parse = TRUE, header, sep = "", col.names,
  version = 1)

## Import to a ValueArray object:
h2o.importFile.VA(object, path, key = "", parse = TRUE, header, sep = "", col.names)

## Import to a FluidVecs object:
h2o.importFile.FV(object, path, key = "", parse = TRUE, header, sep = "", col.names)

```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
path	The path of the file to be imported. Each row of data appears as one line of the file. If it does not contain an absolute path, the file name is relative to the current working directory.
key	(Optional) The unique hex key assigned to the imported file. If none is given, a key will automatically be generated based on the file path.
parse	(Optional) A logical value indicating whether the file should be parsed after import.
header	(Optional) A logical value indicating whether the first row is the column header. If missing, H2O will automatically try to detect the presence of a header.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If <code>sep = ""</code> , the parser will automatically detect the separator.

col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
version	(Optional) If version = 1, the file will be imported to a ValueArray object. Otherwise, if version = 2, the file will be imported as a FluidVecs object.

Details

Calling the method with version = 1 is equivalent to `h2o.importFile.VA`, and version = 2 is equivalent to `h2o.importFile.FV`.

WARNING: In H2O, import is lazy! Do not modify the data on hard disk until after parsing is complete.

Value

If `parse = TRUE`, the function returns an object of class [H2OParsedDataVA](#) when version = 1 and an object of class [H2OParsedData](#) when version = 2. Otherwise, when `parse = FALSE`, it returns an object of class [H2ORawDataVA](#) when version = 1 and an object of class [H2ORawData](#) when version = 2.

See Also

[h2o.importFolder](#), [h2o.importHDFS](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package = "h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
class(iris.hex)
summary(iris.hex)
iris.fv = h2o.importFile(localH2O, path = irisPath, key = "iris.fv", version = 2)
class(iris.fv)
h2o.shutdown(localH2O)
```

`h2o.importFolder`

Import Local Directory of Data Files

Description

Imports all the files in the local directory and parses them, concatenating the data into a single H2O data matrix and returning an object containing the identifying hex key.

Usage

```
## Default method:
h2o.importFolder(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names, version = 1)

## Import to a ValueArray object:
h2o.importFolder.VA(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names)

## Import to a FluidVecs object:
h2o.importFolder.FV(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names)
```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
path	The path of the folder directory to be imported. Each row of data appears as one line of the file. If it does not contain an absolute path, the file name is relative to the current working directory.
key	(Optional) The unique hex key assigned to the imported file. If none is given, a key will automatically be generated based on the file path.
pattern	(Optional) Character string containing a regular expression to match file(s) in the folder.
parse	(Optional) A logical value indicating whether the file should be parsed after import.
header	(Optional) A logical value indicating whether the first row is the column header. If missing, H2O will automatically try to detect the presence of a header.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If <code>sep = ""</code> , the parser will automatically detect the separator.
col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
version	(Optional) If <code>version = 1</code> , the folder of files will be imported to a ValueArray object. Otherwise, if <code>version = 2</code> , the files will be imported as a FluidVecs object.

Details

Calling the method with `version = 1` is equivalent to `h2o.importFolder.VA`, and `version = 2` is equivalent to `h2o.importFolder.FV`.

This method imports all the data files in a given folder and concatenates them together row-wise into a single matrix represented by a [H2OParsedDataVA](#) (version = 1) or [H2OParsedData](#) (version = 2) object. The data files must all have the same number of columns, and the columns must be lined up in the same order, otherwise an error will be returned.

WARNING: In H2O, import is lazy! Do not modify the data files on hard disk until after parsing is complete.

Value

If `parse = TRUE`, the function returns an object of class [H2OParsedDataVA](#) when `version = 1` and an object of class [H2OParsedData](#) when `version = 2`. Otherwise, when `parse = FALSE`, it returns an object of class [H2ORawDataVA](#) when `version = 1` and an object of class [H2ORawData](#) when `version = 2`.

See Also

[h2o.importFile](#), [h2o.importHDFS](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
myPath = system.file("extdata", "prostate_folder", package = "h2o")
prostate_all.hex = h2o.importFolder(localH2O, path = myPath)
class(prostate_all.hex)
summary(prostate_all.hex)
prostate_all.fv = h2o.importFolder(localH2O, path = myPath, version = 2)
class(prostate_all.fv)
h2o.shutdown(localH2O)
```

h2o.importHDFS

Import from HDFS

Description

Imports a HDFS file or set of files in a directory and parses them, returning a object containing the identifying hex key.

Usage

```
## Default method:
h2o.importHDFS(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names, version = 1)

## Import to a ValueArray object:
h2o.importHDFS.VA(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names)

## Import to a FluidVecs object:
h2o.importHDFS.FV(object, path, pattern = "", key = "", parse = TRUE, header,
  sep = "", col.names)
```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
path	The path of the file or folder directory to be imported. If it does not contain an absolute path, the file name is relative to the current working directory.
pattern	(Optional) Character string containing a regular expression to match file(s) in the folder.
key	(Optional) The unique hex key assigned to the imported file. If none is given, a key will automatically be generated based on the file path.
parse	(Optional) A logical value indicating whether the file should be parsed after import.
header	(Optional) A logical value indicating whether the first row is the column header. If missing, H2O will automatically try to detect the presence of a header.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If sep = "", the parser will automatically detect the separator.
col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
version	(Optional) If version = 1, the file will be imported to a ValueArray object. Otherwise, if version = 2, the file will be imported as a FluidVecs object.

Details

Calling the method with version = 1 is equivalent to `h2o.importHDFS.VA`, and version = 2 is equivalent to `h2o.importHDFS.FV`.

When path is a directory, this method acts like [h2o.importFolder](#) and concatenates all data files in the folder into a single ValueArray object.

WARNING: In H2O, import is lazy! Do not modify the data files on hard disk until after parsing is complete.

Value

If parse = TRUE, the function returns an object of class [H2OParsedDataVA](#) when version = 1 and an object of class [H2OParsedData](#) when version = 2. Otherwise, when parse = FALSE, it returns an object of class [H2ORawDataVA](#) when version = 1 and an object of class [H2ORawData](#) when version = 2.

See Also

[h2o.importFile](#), [h2o.importFolder](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
## Not run:
# This is an example of how to import files from HDFS.
# The user must modify the path to his or her specific HDFS path for this example to run.
```



```

library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
iris.hex = h2o.importHDFS(localH2O, path = paste("hdfs://192.168.1.161",
  "datasets/runit/iris_wheader.csv", sep = "/"), parse = TRUE)
class(iris.hex)
summary(iris.hex)
iris.fv = h2o.importHDFS(localH2O, path = paste("hdfs://192.168.1.161",
  "datasets/runit/iris_wheader.csv", sep = "/"), parse = TRUE, version = 2)
class(iris.fv)

iris_folder.hex = h2o.importHDFS(localH2O, path = paste("hdfs://192.168.1.161",
  "datasets/runit/iris_test_train", sep = "/"))
summary(iris_folder.hex)
h2o.shutdown(localH2O)

## End(Not run)

```

h2o.importURL

*Import Data from URL***Description**

Imports a file from the URL and parses it, returning an object containing the identifying hex key.

Usage

```

## Default method:
h2o.importURL(object, path, key = "", parse = TRUE, header,
  sep = "", col.names, version = 1)

## Import to a ValueArray object:
h2o.importURL.VA(object, path, key = "", parse = TRUE, header,
  sep = "", col.names)

## Import to a FluidVecs object:
h2o.importURL.FV(object, path, key = "", parse = TRUE, header,
  sep = "", col.names)

```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
path	The complete URL of the file to be imported. Each row of data appears as one line of the file.
key	(Optional) The unique hex key assigned to the imported file. If none is given, a key will automatically be generated based on the URL path.
parse	(Optional) A logical value indicating whether the file should be parsed after import.

header	(Optional) A logical value indicating whether the first row is the column header. If missing, H2O will automatically try to detect the presence of a header.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If sep = "", the parser will automatically detect the separator.
col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
version	(Optional) If version = 1, the file will be imported to a ValueArray object. Otherwise, if version = 2, the file will be imported as a FluidVecs object.

Details

Calling the method with version = 1 is equivalent to `h2o.importURL.VA`, and version = 2 is equivalent to `h2o.importURL.FV`.

WARNING: In H2O, import is lazy! Do not modify the data on hard disk until after parsing is complete.

Value

If parse = TRUE, the function returns an object of class [H2OParsedDataVA](#) when version = 1 and an object of class [H2OParsedData](#) when version = 2. Otherwise, when parse = FALSE, it returns an object of class [H2ORawDataVA](#) when version = 1 and an object of class [H2ORawData](#) when version = 2.

See Also

[h2o.importFile](#), [h2o.importFolder](#), [h2o.importHDFS](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prostate.hex = h2o.importURL(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smldata/logreg/prostate.csv", sep = "/"), key = "prostate.hex")
class(prostate.hex)
summary(prostate.hex)

prostate.fv = h2o.importURL(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smldata/logreg/prostate.csv", sep = "/"), key = "prostate.hex",
  version = 2)
class(prostate.fv)
h2o.shutdown(localH2O)
```

h2o.init*Connect to H2O and Install R Package*

Description

Connects to a running H2O instance and checks the local H2O R package is the correct version (i.e. that the version of the R package and the version of H2O are the same).

Usage

```
h2o.init(ip = "127.0.0.1", port = 54321, startH2O = TRUE, Xmx = "1g")
```

Arguments

ip	Object of class "character" representing the IP address of the server where H2O is running.
port	Object of class "numeric" representing the port number of the H2O server.
startH2O	(Optional) A logical value indicating whether to start the H2O launcher GUI if no connection with H2O is detected. This is only possible if H2O was installed from the InstallBuilder executable and ip = "localhost" or ip = "127.0.0.1".
Xmx	(Optional) A string specifying the maximum size, in bytes, of the memory allocation pool to H2O. This value must be a multiple of 1024 greater than 2MB. Append the letter m or M to indicate megabytes, or g or G to indicate gigabytes.

Details

This method first checks if H2O is connectable. If it cannot connect and startH2O = TRUE with IP of localhost, it will attempt to start an instance of H2O with IP = localhost, port = 54321. Otherwise, it stops immediately with an error. Once connected, the method checks to see if the local H2O R package version matches the version of H2O running on the server. If there is a mismatch and the user indicates she wishes to upgrade, it will remove the local H2O R package and download/install the H2O R package from the server.

Value

Once the package is successfully installed, this method will load it and return a H2OClient object containing the IP address and port number of the H2O server. See the [H2O R package documentation](#) for more details, or type ??h2o.

Note

Users may wish to manually upgrade their package (rather than waiting until being prompted), which requires that they fully uninstall and reinstall the H2O package, and the H2O client package. You must unload packages running in the environment before upgrading. It's recommended that users restart R or R studio after upgrading.

See Also[h2o.shutdown](#)**Examples**

```
# Try to create a localhost connection to H2O.
localH2O = h2o.init()
localH2O = h2o.init(ip = "localhost", port = 54321)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = FALSE)
h2o.shutdown(localH2O)
```

h2o.kmeans

*H2O: K-Means Clustering***Description**

Performs k-means clustering on a data set.

Usage

```
## Default method:
h2o.kmeans(data, centers, cols = "", iter.max = 10, normalize = FALSE,
  init = "none", seed = 0, version = 1)

## Import to a ValueArray object:
h2o.kmeans.VA(data, centers, cols = "", iter.max = 10, normalize = FALSE,
  init = "none", seed = 0)

## Import to a FluidVecs object:
h2o.kmeans.FV(data, centers, cols = "", iter.max = 10, normalize = FALSE,
  init = "none", seed = 0)
```

Arguments

data	An H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing the variables in the model.
centers	The number of clusters k.
cols	(Optional) A vector containing the names of the data columns on which k-means runs. If blank, k-means clustering will be run on the entire data set.
iter.max	(Optional) The maximum number of iterations allowed.
normalize	(Optional) A logical value indicating whether the data should be normalized before running k-means.
init	(Optional) Method by which to select the k initial cluster centroids. Possible values are "none" for random initialization, "plusplus" for k-means++ initialization, and "furthest" for initialization at the furthest point from each successive centroid. See the H2O K-means documentation for more details.

seed	(Optional) Random seed used to initialize the cluster centroids.
version	(Optional) The version of k-means clustering to run. If version = 1, this will run the more stable ValueArray implementation, while version = 2 selects the faster, but still beta stage FluidVecs implementation.

Details

IMPORTANT: Currently, to run k-means with version = 1, you must import data to a ValueArray object using [h2o.importFile.VA](#), [h2o.importFolder.VA](#) or one of its variants. To run with version = 2, you must import data to a FluidVecs object using [h2o.importFile.FV](#), [h2o.importFolder.FV](#) or one of its variants.

Value

An object of class [H2OKMeansModelVA](#) (version = 1) or [H2OKMeansModel](#) (version = 2) with slots key, data, and model, where the last is a list of the following components:

centers	A matrix of cluster centers.
cluster	A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing the vector of integers (from 1 to k), which indicate the cluster to which each point is allocated.
size	The number of points in each cluster.
withinss	Vector of within-cluster sum of squares, with one component per cluster.
tot.withinss	Total within-cluster sum of squares, i.e., sum(withinss).

See Also

[h2o.importFile](#), [h2o.importFolder](#), [h2o.importHDFS](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package = "h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
h2o.kmeans(data = prostate.hex, centers = 10, cols = c("AGE", "RACE", "VOL", "GLEASON"))

covPath = system.file("extdata", "covtype.csv", package = "h2o")
covtype.hex = h2o.importFile(localH2O, path = covPath, version = 2)
h2o.kmeans(data = covtype.hex, centers = 5, cols = c(1, 2, 3), version = 2)
h2o.shutdown(localH2O)
```

h2o.logAndEcho	<i>Write and Echo Message to H2O Log</i>
----------------	--

Description

Write a user-defined message to the H2O Java log file and echo it back to the user.

Usage

```
h2o.logAndEcho(conn, message)
```

Arguments

conn	An H2OClient object containing the IP address and port of the server running H2O.
message	A character string to write to the H2O Java log file.

See Also

[H2OClient](#), [h2o.downloadAllLogs](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.logAndEcho(localH2O, "Test log and echo method.")
h2o.shutdown(localH2O)
```

h2o.ls	<i>Obtain a list of H2O keys from the running instance of H2O</i>
--------	---

Description

Allows users to access a list of object keys in the running instance of H2O

Usage

```
h2o.ls(object, pattern)
```

Arguments

object	An H2OClient object containing the IP address and port number of the H2O server.
pattern	A string indicating the type of key to be returned. When pattern is left is unspecified all keys are returned.

Value

Returns a list of hex keys in the current instance of H2O, and their associated sizes in bytes.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath, key = "prostate.hex")
s = runif(nrow(prostate.hex))
prostate.train = prostate.hex[s <= 0.8,]
prostate.train = h2o.assign(prostate.train, "prostate.train")
h2o.ls(localH2O)
h2o.shutdown(localH2O)
```

h2o.nn

*H2O: Neural Networks***Description**

Performs neural networks on an [H2OParsedData](#) object.

Usage

```
h2o.nn(x, y, data, classification = TRUE, activation = "Tanh", dropout = as.numeric(NA),
  layers = 500, rate = 0.01, annealing_rate = 1e-6, l1_reg = 1e-4, l2_reg = 0.0010,
  mom_start = 0.5, mom_ramp = 1e6, mom_stable = 0.99, epochs = 100, validation)
```

Arguments

- | | |
|----------------|---|
| x | A vector containing the names of the predictors in the model. |
| y | The name of the response variable in the model. |
| data | An H2OParsedData object containing the variables in the model. |
| classification | (Optional) A logical value indicating whether the algorithm should conduct classification. |
| activation | (Optional) A string indicating the activation function to use. Must be either "Tanh", "Rectifier", or "Maxout". |
| dropout | (Optional) Input layer dropout ratio. |
| layers | (Optional) Size of the hidden layers. |
| rate | (Optional) The learning rate. A higher learning rate will cause the algorithm to be less stable, while a lower learning rate will result in slower convergence. |
| annealing_rate | (Optional) The learning annealing rate is equal to |

$$rate / (1 + annealing_rate * samples)$$

.

l1_reg	(Optional) L1 regularization parameter.
l2_reg	(Optional) L2 regularization parameter.
mom_start	(Optional) Initial momentum at the beginning of training.
mom_ramp	(Optional) Number of training samples for which momentum increases.
mom_stable	(Optional) Final momentum after the ramp is over.
epochs	(Optional) Number of times the dataset should be iterated. (Can be less than 1.0).
validation	(Optional) An H2OParsedData object indicating the validation dataset used to construct confusion matrix. If left blank, this defaults to the training data.

Value

An object of class [H2ONNModel](#) with slots key, data, and model, where the last is a list of the following components:

confusion	The confusion matrix of the response, with actual observations as rows and predicted values as columns.
train_class_err	Classification error on the training dataset.
train_sqr_err	Mean-squared error on the training dataset.
valid_class_err	Classification error on the validation dataset.
valid_sqr_err	Mean-squared error on the validation dataset.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package = "h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath)
h2o.nn(x = 1:4, y = 5, data = iris.hex, activation = "Rectifier", rate = 0.05)
h2o.shutdown(localH2O)
```

h2o.parseRaw	<i>Parse Raw Data File</i>
--------------	----------------------------

Description

Parses a raw data file, returning an object containing the identifying hex key.

Usage

```
## Default method:
h2o.parseRaw(data, key = "", header, sep = "", col.names, version = 1)

## Import to a ValueArray object:
h2o.parseRaw.VA(data, key = "", header, sep = "", col.names)

## Import to a FluidVecs object:
h2o.parseRaw.FV(data, key = "", header, sep = "", col.names)
```

Arguments

data	An H2ORawDataVA (version = 1) or H2ORawData (version = 2) object to be parsed.
key	(Optional) The hex key assigned to the parsed file.
header	(Optional) A logical value indicating whether the first row is the column header. If missing, H2O will automatically try to detect the presence of a header.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If sep = "", the parser will automatically detect the separator.
col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
version	(Optional) If version = 1, the file will be parsed to a ValueArray object. Otherwise, if version = 2, the file will be parsed to a FluidVecs object.

Details

Calling the method with version = 1 is equivalent to `h2o.parseRaw.VA`, and version = 2 is equivalent to `h2o.parseRaw.FV`. `h2o.parseRaw.VA` should only be used to parse raw data imported using [h2o.importFile.VA](#), [h2o.importFolder.VA](#), or one of its variants. Similarly, `h2o.parseRaw.FV` should only be used to parse raw data imported using [h2o.importFile.FV](#), [h2o.importFolder.FV](#), or one of its variants.

After the raw data file is parsed, it will be automatically deleted from the H2O server.

Value

An object of class [H2OParsedDataVA](#) (version = 1) or [H2OParsedData](#) (version = 2), representing the dataset that was parsed.

See Also

[h2o.importFile](#), [h2o.importFolder](#), [h2o.importHDFS](#), [h2o.importURL](#), [h2o.uploadFile](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.raw = h2o.importFile(localH2O, path = prosPath, parse = FALSE)
# Do not modify prostate.csv on disk at this point!
prostate.hex = h2o.parseRaw(data = prostate.raw, key = "prostate.hex")
# After parsing, it is okay to modify or delete prostate.csv
h2o.shutdown(localH2O)
```

h2o.pcr

H2O: Principal Components Regression

Description

Runs GLM regression on PCA results, and allows for transformation of test data to match PCA transformations of training data.

Usage

```
h2o.pcr(x, y, data, ncomp, family, nfolds = 10, alpha = 0.5, lambda = 1e-05,
        epsilon = 1e-05, tweedie.p)
```

Arguments

x	A vector containing the names of the predictors in the model.
y	The name of the response variable in the model.
data	An H2OParsedData object containing the variables in the model.
ncomp	A number indicating the number of principal components to use in the regression model.
family	A description of the error distribution and corresponding link function to be used in the model. Currently, Gaussian, binomial, Poisson, gamma, and Tweedie are supported.
nfolds	(Optional) Number of folds for cross-validation. The default is 10.
alpha	(Optional) The elastic-net mixing parameter, which must be in [0,1]. The penalty is defined to be

$$P(\alpha, \beta) = (1 - \alpha)/2 \|\beta\|_2^2 + \alpha \|\beta\|_1 = \sum_j [(1 - \alpha)/2 \beta_j^2 + \alpha |\beta_j|]$$

so alpha=1 is the lasso penalty, while alpha=0 is the ridge penalty.

lambda	(Optional) The shrinkage parameter, which multiplies $P(\alpha, \beta)$ in the objective. The larger lambda is, the more the coefficients are shrunk toward zero (and each other).
epsilon	(Optional) Number indicating the cutoff for determining if a coefficient is zero.
tweedie.p	The index of the power variance function for the tweedie distribution. Only used if family = "tweedie"

Details

This method standardizes the data, obtains the first `ncomp` principal components using PCA (in decreasing order of standard deviation), and then runs GLM with the components as the predictor variables.

Value

An object of class [H2OGLMModel](#) with slots `key`, `data`, `model` and `xval`. The slot `model` is a list of the following components:

<code>coefficients</code>	A named vector of the coefficients estimated in the model.
<code>rank</code>	The numeric rank of the fitted linear model.
<code>family</code>	The family of the error distribution.
<code>deviance</code>	The deviance of the fitted model.
<code>aic</code>	Akaike's Information Criterion for the final computed model.
<code>null.deviance</code>	The deviance for the null model.
<code>iter</code>	Number of algorithm iterations to compute the model.
<code>df.residual</code>	The residual degrees of freedom.
<code>df.null</code>	The residual degrees of freedom for the null model.
<code>y</code>	The response variable in the model.
<code>x</code>	A vector of the predictor variable(s) in the model.
<code>auc</code>	Area under the curve.
<code>training.err</code>	Average training error.
<code>threshold</code>	Best threshold.
<code>confusion</code>	Confusion matrix.

The slot `xval` is a list of [H2OGLMModel](#) objects representing the cross-validation models. (Each of these objects themselves has `xval` equal to an empty list).

See Also

[h2o.prcomp](#), [h2o.glm](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)

# Run PCR on Prostate Data
prostate.hex = h2o.importURL(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smalldata/logreg/prostate.csv", sep = "/"), key = "prostate.hex")
h2o.pcr(x = c("AGE", "RACE", "PSA", "DCAPS"), y = "CAPSULE", data = prostate.hex, family = "binomial",
  nfold = 10, alpha = 0.5, ncomp = 3)
h2o.shutdown(localH2O)
```

h2o.prcomp

*Principal Components Analysis***Description**

Performs principal components analysis on the given data set.

Usage

```
h2o.prcomp(data, tol = 0, ignored_cols = "", standardize = TRUE, retx = FALSE)
```

Arguments

data	An H2OParsedData object on which to run principal components analysis.
tol	(Optional) A value indicating the magnitude below which components should be omitted. (Components are omitted if their standard deviations are less than or equal to tol times the standard deviation of the first component.) With the default setting tol = 0, no components are omitted.
ignored_cols	(Optional) A vector of column names or indices indicating the features to ignore when performing PCA. By default, all columns in the dataset are analyzed.
standardize	(Optional) A logical value indicating whether the variables should be shifted to be zero centered and scaled to have unit variance before the analysis takes place.
retx	(Optional) A logical value indicating whether the rotated variables should be returned.

Details

The calculation is done by a singular value decomposition of the (possibly standardized) data set.

Value

An object of class [H2OPCAModel](#) with slots key, data, and model, where the last is a list of the following components:

standardized	A logical value indicating whether the data was centered and scaled.
sdev	The standard deviations of the principal components (i.e., the square roots of the eigenvalues of the covariance/correlation matrix).
rotation	The matrix of variable loadings (i.e., a matrix whose columns contain the eigenvectors).

Note

The signs of the columns of the rotation matrix are arbitrary, and so may differ between different programs for PCA.

See Also[h2o.pcr](#)**Examples**

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
australia.pca = h2o.prcomp(data = australia.hex, standardize = TRUE)
print(australia.pca)
h2o.shutdown(localH2O)
```

h2o.predict*H2O Model Predictions*

Description

Obtains predictions from various fitted H2O model objects.

Usage

```
h2o.predict(object, newdata)
```

Arguments

object	A fitted H2OModel or H2OModelVA object for which prediction is desired.
newdata	(Optional) A H2OParsedData or H2OParsedDataVA object in which to look for variables with which to predict. If omitted, the data used to fit the model <code>object@data</code> are used.

Details

This method dispatches on the type of H2O model to select the correct prediction/scoring algorithm.

Value

A [H2OParsedData](#) or [H2OParsedDataVA](#) object containing the predictions.

See Also

[h2o.glm](#), [h2o.kmeans](#), [h2o.randomForest](#), [h2o.prcomp](#), [h2o.gbm](#), [h2o.nn](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
# Run GLM of CAPSULE ~ AGE + RACE + PSA + DCAPS
prostate.hex = h2o.importURL.VA(localH2O, path = paste("https://raw.githubusercontent.com",
  "0xdata/h2o/master/smalldata/logreg/prostate.csv", sep = "/"), key = "prostate.hex")
prostate.glm = h2o.glm(y = "CAPSULE", x = c("AGE", "RACE", "PSA", "DCAPS"), data = prostate.hex,
  family = "binomial", nfolds = 10, alpha = 0.5)
# Get fitted values of prostate dataset
prostate.fit = h2o.predict(object = prostate.glm, newdata = prostate.hex)
summary(prostate.fit)

# Run K-Means Model on prostate data set
covPath = system.file("extdata", "covtype.csv", package="h2o")
covtype.hex = h2o.importFile(localH2O, path = covPath)
covtype.km = h2o.kmeans(data = covtype.hex, centers = 5, cols = c(1, 2, 3))
# Get predicted cluster information for covtype dataset
covtype.clusters = h2o.predict(object = covtype.km, newdata = covtype.hex)
summary(covtype.clusters)
h2o.shutdown(localH2O)
```

h2o.randomForest

H2O: Random Forest

Description

Performs random forest classification on a data set.

Usage

```
## Default method:
h2o.randomForest(x, y, data, ntree = 50, depth = 50, sample.rate = 2/3, classwt = NULL,
  nbins = 100, seed = -1, importance = FALSE, validation, nodesize = 1,
  use_non_local = TRUE, version = 1)

## Import to a ValueArray object:
h2o.randomForest.VA(x, y, data, ntree = 50, depth = 50, sample.rate = 2/3,
  classwt = NULL, nbins = 100, seed = -1, use_non_local = TRUE)

## Import to a FluidVecs object:
h2o.randomForest.FV(x, y, data, ntree = 50, depth = 50, sample.rate = 2/3, nbins = 100,
  seed = -1, importance = FALSE, validation, nodesize = 1)
```

Arguments

x A vector containing the names or indices of the predictor variables to use in building the random forest model.

y	The name or index of the response variable. If the data does not contain a header, this is the column index, designated by increasing numbers from left to right. (The response must be either an integer or a categorical variable).
data	An H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing the variables in the model.
ntree	(Optional) Number of trees to grow. (Must be a nonnegative integer).
depth	(Optional) Maximum depth to grow the tree.
sample.rate	(Optional) Sampling rate for constructing data from which individual trees are grown.
classwt	(Optional) Numeric vector of class weights for a categorical response.
nbins	(Optional) Build a histogram of this many bins, then split at best point.
seed	(Optional) Seed for building the random forest. If seed = -1, one will automatically be generated by H2O.
importance	(Optional) A logical value indicating whether to calculate variable importance. Set to FALSE to speed up computations.
validation	(Optional) An H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object indicating the validation dataset used to construct confusion matrix. If left blank, this defaults to the training data.
nodesize	(Optional) Number of nodes to use for computation.
use_non_local	(Optional) Logical value indicating whether to use non-local data in building random forest model.
version	(Optional) The version of random forest to run. If version = 1, this will run the single-node ValueArray implementation, while version = 2 selects the distributed, but still beta stage FluidVecs implementation.

Details

IMPORTANT: Currently, to run k-means with version = 1, you must import data to a ValueArray object using [h2o.importFile.VA](#), [h2o.importFolder.VA](#) or one of its variants. To run with version = 2, you must import data to a FluidVecs object using [h2o.importFile.FV](#), [h2o.importFolder.FV](#) or one of its variants.

Value

An object of class [H2ORFModelVA](#) (version = 1) or [H2ODRFModel](#) (version = 2) with slots key, data, and model, where the last is a list of the following components:

ntree	Number of trees grown.
mse	Mean-squared error for each tree.
forest	A matrix giving the minimum, mean, and maximum of the tree depth and number of leaves.
confusion	Confusion matrix of the prediction.

Examples

```
# Run an RF model on iris data
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package = "h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
h2o.randomForest(y = 5, x = c(2,3,4), data = iris.hex, ntree = 50, depth = 100)

# Run an RF model on covtype data
covPath = system.file("extdata", "covtype.csv", package="h2o")
covtype.hex = h2o.importFile(localH2O, path = covPath, key = "covtype.hex", version = 2)
h2o.randomForest(y = "Cover_Type", x = setdiff(colnames(covtype.hex), c("Cover_Type",
  "Aspect", "Hillshade_9am")), data = covtype.hex, ntree = 50, depth = 150, version = 2)
h2o.shutdown(localH2O)
```

h2o.rm

Removes H2O objects from the server where H2O is running.

Description

Allows users to remove H2O objects from the server where the instance of H2O is running. This call acts on the H2O server through the R console, and does NOT remove the associated named object from the R environment.

Usage

```
h2o.rm(object, keys)
```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
keys	the hex key associated with the object to be removed.

Note

Users may wish to remove an H2O object on the server that is associated with an object in the R environment. Recommended behavior is to also remove the object in the R environment. See the second example at the end of this section.

See Also

[h2o.assign](#), [h2o.ls](#)

Examples

```
# Remove an H2O object from the server where H2O is running.
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath, key = "prostate.hex")

# Remove an H2O object from the server and from the R environment
h2o.ls(localH2O)
h2o.rm(object = localH2O, keys = "prostate.hex")
remove(prostate.hex)
h2o.ls(localH2O)
h2o.shutdown(localH2O)
```

h2o.runif	<i>Produces a vector of specified length contain random uniform numbers</i>
-----------	---

Description

Produces a vector of random uniform numbers.

Usage

```
h2o.runif(x, min, max)
```

Arguments

x	An H2OParsedData object with number of rows equal to the number of elements the vector of random numbers should have.
min	An integer specifying the lower bound of the distribution.
max	An integer specifying the upper bound of the distribution.

Details

x must be a [H2OParsedData](#) object so that H2O can generate random numbers aligned with the dataset for efficient large-scale sampling and filtering.

Value

A vector of random, uniformly distributed numbers. The elements are between 0 and 1 unless otherwise specified.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath, key = "prostate.hex")
s = h2o.runif(prostate.hex)
summary(s)

prostate.train = prostate.hex[s <= 0.8,]
prostate.train = h2o.assign(prostate.train, "prostate.train")
prostate.test = prostate.hex[s > 0.8,]
prostate.test = h2o.assign(prostate.test, "prostate.test")
nrow(prostate.train) + nrow(prostate.test)
h2o.shutdown(localH2O)
```

h2o.shutdown	<i>Shutdown H2O server</i>
--------------	----------------------------

Description

Shuts down the specified H2O instance. All data on the server will be lost!

Usage

```
h2o.shutdown(client, prompt = TRUE)
```

Arguments

client	An H2OClient client containing the IP address and port of the server running H2O.
prompt	(Optional) A logical value indicating whether to prompt the user before shutting down the H2O server.

Details

This method checks if H2O is running at the specified IP address and port, and if it is, shuts down that H2O instance. **WARNING:** All data, models, and other values stored on the server will be lost! Only call this function if you and all other clients connected to the H2O server are finished and have saved your work.

Note

Users must call `h2o.shutdown` explicitly in order to shut down the local H2O instance started by R. If R is closed before H2O, then an attempt will be made to automatically shut down H2O. This only applies to local instances started with `h2o.init`, not remote H2O servers.

See Also

[h2o.init](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.shutdown(localH2O)
```

h2o.table

*Cross Tabulation of H2O Data***Description**

Uses the cross-classifying factors to build a table of counts at each combination of factor levels.

Usage

```
h2o.table(x)
```

Arguments

x An [H2OParsedData](#) object with at most two integer or factor columns.

Value

A [H2OParsedData](#) object containing the contingency table. If **x** has a single column, this will just be the counts of each factor level.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath, key = "prostate.hex")
summary(prostate.hex)

# Counts of the ages of all patients
head(h2o.table(prostate.hex[,3]))

# Two-way table of ages (rows) and race (cols) of all patients
head(h2o.table(prostate.hex[,c(3,4)]))
h2o.shutdown(localH2O)
```

h2o.uploadFile

*Upload Local Data File***Description**

Uploads a file from the local drive and parses it, returning an object containing the identifying hex key.

Usage

```
## Default method:
h2o.uploadFile(object, path, key = "", parse = TRUE, header,
  sep = "", col.names, silent = TRUE, version = 1)

## Import to a ValueArray object:
h2o.uploadFile.VA(object, path, key = "", parse = TRUE, header,
  sep = "", col.names, silent = TRUE)

## Import to a FluidVecs object:
h2o.uploadFile.FV(object, path, key = "", parse = TRUE, header,
  sep = "", col.names, silent = TRUE)
```

Arguments

object	An H2OClient object containing the IP address and port of the server running H2O.
path	The complete URL or normalized file path of the file to be imported. Each row of data appears as one line of the file.
key	(Optional) The unique hex key assigned to the imported file. If none is given, a key will automatically be generated based on the URL path.
parse	(Optional) A logical value indicating whether the file should be parsed after import.
header	(Optional) A logical value indicating whether the first line of the file contains column headers. If left empty, the parser will try to automatically detect this.
sep	(Optional) The field separator character. Values on each line of the file are separated by this character. If <code>sep = ""</code> , the parser will automatically detect the separator.
col.names	(Optional) A H2OParsedDataVA (version = 1) or H2OParsedData (version = 2) object containing a single delimited line with the column names for the file.
silent	(Optional) A logical value indicating whether or not to display an upload progress bar.
version	(Optional) If <code>version = 1</code> , the file will be imported to a ValueArray object. Otherwise, if <code>version = 2</code> , the file will be imported as a FluidVecs object.

Details

Calling the method with `version = 1` is equivalent to `h2o.uploadFile.VA`, and `version = 2` is equivalent to `h2o.uploadFile.FV`.

WARNING: In H2O, import is lazy! Do not modify the data on hard disk until after parsing is complete.

Value

If `parse = TRUE`, the function returns an object of class `H2OParsedDataVA` when `version = 1` and an object of class `H2OParsedData` when `version = 2`. Otherwise, when `parse = FALSE`, it returns an object of class `H2OParsedDataVA` when `version = 1` and an object of class `H2ORawData` when `version = 2`.

See Also

`h2o.importFile`, `h2o.importFolder`, `h2o.importHDFS`, `h2o.importURL`

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package = "h2o")
prostate.hex = h2o.uploadFile(localH2O, path = prosPath, key = "prostate.hex")
class(prostate.hex)
summary(prostate.hex)
prostate.fv = h2o.uploadFile(localH2O, path = prosPath, key = "prostate.fv", version = 2)
class(prostate.fv)
h2o.shutdown(localH2O)
```

h2o.__changeLogPath	<i>Change Path Where H2O R Logs are Saved</i>
---------------------	---

Description

Change the file path where H2O R command and error response logs are currently being saved.

Usage

```
h2o.__changeLogPath(path, type)
```

Arguments

path	A character string indicating the new file path where logs should be saved.
type	Which log file's path to modify. Either "Command" for POST commands sent between R and H2O, or "Error" for errors returned by H2O in the HTTP response.

See Also

[h2o.__startLogging](#), [h2o.__stopLogging](#), [h2o.__clearLogs](#), [h2o.__openLog](#), [h2o.__getLogPath](#)

Examples

```
library(h2o)
h2o.__getLogPath("Command")
h2o.__changeLogPath(getwd(), "Command")
h2o.__getLogPath("Command")
```

h2o.__clearLogs	<i>Delete All H2O R Logs</i>
-----------------	------------------------------

Description

Clear all H2O R command and error response logs from local disk. Used primarily for debugging purposes.

Usage

```
h2o.__clearLogs()
```

See Also

[h2o.__startLogging](#), [h2o.__stopLogging](#), [h2o.__openLog](#), [h2o.__getLogPath](#), [h2o.__changeLogPath](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.__startLogging()
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
h2o.__stopLogging()
h2o.__clearLogs()
h2o.shutdown(localH2O)
```

h2o.__getLogPath	<i>Get Path Where H2O R Logs are Saved</i>
------------------	--

Description

Get the file path where H2O R command and error response logs are currently being saved.

Usage

```
h2o.__getLogPath(type)
```

Arguments

type	Which log file's path to get. Either "Command" for POST commands sent between R and H2O, or "Error" for errors returned by H2O in the HTTP response.
------	--

See Also

[h2o.__startLogging](#), [h2o.__stopLogging](#), [h2o.__clearLogs](#), [h2o.__openLog](#), [h2o.__changeLogPath](#)

Examples

```
library(h2o)
h2o.__getLogPath("Command")
h2o.__getLogPath("Error")
```

h2o.__openLog	<i>View H2O R Logs</i>
---------------	------------------------

Description

Open existing logs of H2O R POST commands and error responses on local disk. Used primarily for debugging purposes.

Usage

```
h2o.__openLog(type)
```

Arguments

type	Which log file to open. Either "Command" for POST commands sent between R and H2O, or "Error" for errors returned by H2O in the HTTP response.
------	--

See Also

[h2o.__startLogging](#), [h2o.__stopLogging](#), [h2o.__clearLogs](#), [h2o.__getLogPath](#), [h2o.__changeLogPath](#)

Examples

```
## Not run:
# Skip running this to avoid windows being opened during R CMD check
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)

h2o.__startLogging()
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
h2o.__stopLogging()

h2o.__openLog("Command")
h2o.__openLog("Error")
h2o.shutdown(localH2O)

## End(Not run)
```

h2o.__startLogging	<i>Start Writing H2O R Logs</i>
--------------------	---------------------------------

Description

Begin logging H2O R POST commands and error responses to local disk. Used primarily for debugging purposes.

Usage

```
h2o.__startLogging()
```

See Also

[h2o.__stopLogging](#), [h2o.__clearLogs](#), [h2o.__openLog](#), [h2o.__getLogPath](#), [h2o.__changeLogPath](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.__startLogging()
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
h2o.__stopLogging()
h2o.shutdown(localH2O)
```

h2o.__stopLogging	<i>Stop Writing H2O R Logs</i>
-------------------	--------------------------------

Description

Halt logging of H2O R POST commands and error responses to local disk. Used primarily for debugging purposes.

Usage

```
h2o.__stopLogging()
```

See Also

[h2o.__startLogging](#), [h2o.__clearLogs](#), [h2o.__openLog](#), [h2o.__getLogPath](#), [h2o.__changeLogPath](#)

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
h2o.__startLogging()
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
h2o.__stopLogging()
h2o.shutdown(localH2O)
```

H2OClient-class	<i>Class "H2OClient"</i>
-----------------	--------------------------

Description

An object representing the server/local machine on which H2O is running.

Objects from the Class

Objects can be created by calls of the form `new("H2OClient", ...)`

Slots

ip: Object of class "character" representing the IP address of the H2O server.

port: Object of class "numeric" representing the port number of the H2O server.

Methods

```
h2o.importFile signature(object = "H2OClient", path = "character",    + key = "character", parse = "logical", ...
...
h2o.importFolder signature(object = "H2OClient", path = "character",    + parse = "logical"):
...
h2o.importURL signature(object = "H2OClient", path = "character",    + key = "character", parse = "logical", ...
...
show signature(object = "H2OClient"): ...
```

Examples

```
showClass("H2OClient")
```

H2ODRFGrid-class	Class "H2ODRFGrid"
------------------	--------------------

Description

Object representing the models built by a H2O distributed random forest grid search on FluidVecs.

Objects from the Class

Objects can be created by calls of the form `new("H2ODRFGrid", ...)`.

Slots

- key: Object of class "character", representing the unique hex key that identifies the model.
- data: Object of class "H2OParsedData", which is the input data used to build the model.
- model: Object of class "list" containing "H2ODRFModel" objects representing the models returned by the distributed random forest grid search.
- sumtable: Object of class "list" containing summary statistics of all the models returned by the distributed random forest grid search.

Extends

Class "[H2OGrid](#)", directly.

Methods

No methods defined with class "H2ODRFGrid" in the signature.

See Also

[H2ODRFModel](#), [h2o.randomForest](#)

Examples

```
showClass("H2ODRFGrid")
```

H2ODRFModel-class	Class "H2ODRFModel"
-------------------	---------------------

Description

A class for representing random forest ensembles built on FluidVecs data.

Objects from the Class

Objects can be created by calls of the form `new("H2ODRFModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **type:** The type of the tree, which at this point must be classification.
- **ntree:** Number of trees grown.
- **oob_err:** Out of bag error rate.
- **forest:** A matrix giving the minimum, mean, and maximum of the tree depth and number of leaves.
- **confusion:** Confusion matrix of the prediction.

valid: Object of class "H2OParsedData", which is the data used for validating the model.

Extends

Class "[H2OModel](#)", directly.

Methods

`show signature(object = "H2ODRFModel"): ...`

See Also

[h2o.randomForest](#)

Examples

```
showClass("H2ODRFModel")
```

H2OGBMGrid-class	Class "H2OGBMGrid"
------------------	--------------------

Description

Object representing the models built by a H2O GBM grid search.

Objects from the Class

Objects can be created by calls of the form `new("H2OGBMGrid", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing "H2OGBMModel" objects representing the models returned by the GBM grid search.

sumtable: Object of class "list" containing summary statistics of all the models returned by the GBM grid search.

Extends

Class "[H2OGrid](#)", directly.

Methods

No methods defined with class "H2OGBMGrid" in the signature.

See Also

[H2OGBMModel](#), [h2o.gbm](#)

Examples

```
showClass("H2OGBMGrid")
```

H2OGBMModel-class	Class "H2OGBMModel"
-------------------	---------------------

Description

A class for representing generalized boosted classification/regression models.

Objects from the Class

Objects can be created by calls of the form `new("H2OGBMModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedData](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **err:** The mean-squared error in each tree.
- **cm:** (Only for classification). The confusion matrix of the response, with actual observations as rows and predicted values as columns.

valid: Object of class [H2OParsedData](#), which is the dataset used to validate the model.

Extends

Class "[H2OModel](#)", directly.

Methods

show signature(object = "H2OGBMModel"): ...

See Also

[h2o.gbm](#)

Examples

```
showClass("H2OGBMModel")
```

H2OGLMGrid-class	Class "H2OGLMGrid"
------------------	--------------------

Description

Object representing the models built by a H2O GLM grid search on FluidVecs.

Objects from the Class

Objects can be created by calls of the form `new("H2OGLMGrid", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing "H2OGLMModel" objects representing the models returned by the GLM (FluidVecs) grid search.

sumtable: Object of class "list" containing summary statistics of all the models returned by the GLM (FluidVecs) grid search.

Extends

Class "[H2OGrid](#)", directly.

Methods

No methods defined with class "H2OGLMGrid" in the signature.

See Also

[H2OGLMModel](#), [h2o.glm](#)

Examples

```
showClass("H2OGLMGrid")
```

H2OGLMGridVA-class	Class "H2OGLMGridVA"
--------------------	----------------------

Description

Object representing the models built by a H2O GLM grid search on ValueArray.

Objects from the Class

Objects can be created by calls of the form `new("H2OGLMGridVA", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedDataVA", which is the input data used to build the model.

model: Object of class "list" containing "H2OGLMModelVA" objects representing the models returned by the GLM (ValueArray) grid search.

sumtable: Object of class "list" containing summary statistics of all the models returned by the GLM (ValueArray) grid search.

Extends

Class "[H2OGridVA](#)", directly.

Methods

show `signature(object = "H2OGLMGridVA"): ...`

See Also

[H2OGLMModelVA](#), [h2o.glm](#)

Examples

```
showClass("H2OGLMGridVA")
```

H2OGLMMModel-class	Class "H2OGLMMModel"
--------------------	----------------------

Description

A class for representing generalized linear models.

Objects from the Class

Objects can be created by calls of the form `new("H2OGLMMModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedData](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **coefficients:** A named vector of the coefficients estimated in the model.
- **rank:** The numeric rank of the fitted linear model.
- **family:** The family of the error distribution.
- **deviance:** The deviance of the fitted model.
- **aic:** Akaike's Information Criterion for the final computed model.
- **null.deviance:** The deviance for the null model.
- **iter:** Number of algorithm iterations to compute the model.
- **df.residual:** The residual degrees of freedom.
- **df.null:** The residual degrees of freedom for the null model.
- **y:** The response variable in the model.
- **x:** A vector of the predictor variable(s) in the model.

xval: List of objects of class "H2OGLMMModel", representing the n-fold cross-validation models.

Extends

Class "[H2OModel](#)", directly.

Methods

show signature(object = "H2OGLMMModel"): ...

See Also

[h2o.glm](#)

Examples

```
showClass("H2OGLMMModel")
```

H2OGLMMModelVA-class	Class "H2OGLMMModelVA"
----------------------	------------------------

Description

A class for representing generalized linear models built on ValueArray data.

Objects from the Class

Objects can be created by calls of the form `new("H2OGLMMModelVA", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedDataVA](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **coefficients:** A named vector of the coefficients estimated in the model.
- **rank:** The numeric rank of the fitted linear model.
- **family:** The family of the error distribution.
- **deviance:** The deviance of the fitted model.
- **aic:** Akaike's Information Criterion for the final computed model.
- **null.deviance:** The deviance for the null model.
- **iter:** Number of algorithm iterations to compute the model.
- **df.residual:** The residual degrees of freedom.
- **df.null:** The residual degrees of freedom for the null model.
- **y:** The response variable in the model.
- **x:** A vector of the predictor variable(s) in the model.

xval: List of objects of class "H2OGLMMModelVA", representing the n-fold cross-validation models.

Extends

Class "[H2OModelVA](#)", directly.

Methods

show signature(object = "H2OGLMMModelVA"): ...

See Also

[h2o.glm](#)

Examples

```
showClass("H2OGLMMModelVA")
```

H2OGrid-class

Class "H2OGrid"

Description

Object representing the models built by a H2O grid search algorithm on a FluidVecs dataset.

Objects from the Class

A virtual Class: No objects may be created from it.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing "H2OModel" objects representing the models returned by the grid search algorithm.

sumtable: Object of class "list" containing summary statistics of all the models returned by the grid search algorithm.

Methods

show signature(object = "H2OGrid"): ...

See Also

[H2OGLMGrid](#), [H2OKMeansGrid](#), [H2ODRFGrid](#), [H2OGBMGrid](#), [H2ONNGrid](#)

Examples

```
showClass("H2OGrid")
```

H2OGridVA-class

Class "H2OGridVA"

Description

Object representing the models built by a H2O grid search algorithm on a ValueArray dataset.

Objects from the Class

A virtual Class: No objects may be created from it.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.
data: Object of class "H2OParsedDataVA", which is the input data used to build the model.
model: Object of class "list" containing "H2OModelVA" objects representing the models returned by the grid search algorithm.
sumtable: Object of class "list" containing summary statistics of all the models returned by the grid search algorithm.

Methods

show signature(object = "H2OGridVA"): ...

See Also

[H2OGLMGridVA](#)

Examples

```
showClass("H2OGridVA")
```

H2OKMeansGrid-class	Class "H2OKMeansGrid"
---------------------	-----------------------

Description

Object representing the models built by a H2O K-Means grid search on FluidVecs.

Objects from the Class

Objects can be created by calls of the form `new("H2OKMeansGrid", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.
data: Object of class "H2OParsedData", which is the input data used to build the model.
model: Object of class "list" containing "H2OKMeansModel" objects representing the models returned by the K-Means (FluidVecs) grid search.
sumtable: Object of class "list" containing summary statistics of all the models returned by the K-Means (FluidVecs) grid search.

Extends

Class "[H2OGrid](#)", directly.

Methods

No methods defined with class "H2OKMeansGrid" in the signature.

See Also

[H2OKMeansModel](#), [h2o.kmeans](#)

Examples

```
showClass("H2OKMeansGrid")
```

H2OKMeansModel-class *Class* "H2OKMeansModel"

Description

A class for representing k-means models.

Objects from the Class

Objects can be created by calls of the form `new("H2OKMeansModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedData](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **centers:** A matrix of cluster centers.
- **cluster:** A [H2OParsedData](#) object containing the vector of integers (from 1:k), which indicate the cluster to which each point is allocated.
- **size:** The number of points in each cluster.
- **withinss:** Vector of within-cluster sum of squares, with one component per cluster.
- **tot.withinss:** Total within-cluster sum of squares, i.e., `sum(withinss)`.

Methods

show `signature(object = "H2OKMeansModel")`: ...

See Also

[h2o.kmeans](#)

Examples

```
showClass("H2OKMeansModel")
```

`H2OKMeansModelVA-class``Class "H2OKMeansModelVA"`

Description

A class for representing k-means clustering models built on ValueArray data.

Objects from the Class

Objects can be created by calls of the form `new("H2OKMeansModelVA", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedDataVA](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **cluster:** A [H2OParsedDataVA](#) object, which contains the cluster assignment for each observation in the input data.
- **centers:** A matrix of cluster centers.
- **withinss:** Within-cluster sum of squared errors for each cluster.
- **tot.withinss:** Sum total within-cluster sum of squared errors.
- **size:** Number of observations in each cluster.

Extends

Class "[H2OModelVA](#)", directly.

Methods

`show` `signature(object = "H2OKMeansModelVA")`: ...

See Also

[h2o.kmeans](#)

Examples

```
showClass("H2OKMeansModelVA")
```

H2OModel-class	Class "H2OModel"
----------------	------------------

Description

Object representing the model built by an H2O algorithm on a FluidVecs dataset.

Objects from the Class

A virtual Class: No objects may be created from it.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing the characteristics of the model returned by the algorithm.

Methods

No methods defined with class "H2OModel" in the signature.

See Also

[H2OGLMModel](#), [H2OKMeansModel](#), [H2ODRFModel](#), [H2OGBMModel](#), [H2OPCAModel](#), [H2ONNModel](#)

Examples

```
showClass("H2OModel")
```

H2OModelVA-class	Class "H2OModelVA"
------------------	--------------------

Description

Object representing the model built by an H2O algorithm on a ValueArray dataset.

Objects from the Class

A virtual Class: No objects may be created from it.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedDataVA", which is the input data used to build the model.

model: Object of class "list" containing the characteristics of the model returned by the algorithm.

Methods

No methods defined with class "H2OModelVA" in the signature.

See Also

[H2OGLMModelVA](#), [H2OKMeansModelVA](#), [H2ORFModelVA](#)

Examples

```
showClass("H2OModelVA")
```

H2ONNGrid-class

Class "H2ONNGrid"

Description

Object representing the models built by a H2O neural networks grid search.

Objects from the Class

Objects can be created by calls of the form `new("H2ONNGrid", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class "H2OParsedData", which is the input data used to build the model.

model: Object of class "list" containing "H2ONNModel" objects representing the models returned by the neural networks grid search.

sumtable: Object of class "list" containing summary statistics of all the models returned by the neural networks grid search.

Extends

Class "[H2OGrid](#)", directly.

Methods

No methods defined with class "H2ONNGrid" in the signature.

See Also

[H2ONNModel](#), [h2o.nn](#)

Examples

```
showClass("H2ONNGrid")
```

H2ONNModel-class	Class "H2ONNModel"
------------------	--------------------

Description

A class for representing neural network models.

Objects from the Class

Objects can be created by calls of the form `new("H2ONNModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedData](#), which is the input data used to build the model.

valid: Object of class "H2OParsedData", representing the validation data set.

model: Object of class "list" containing the following elements:

- **confusion:** The confusion matrix of the response, with actual observations as rows and predicted values as columns.
- **train_class_err:** Classification error on the training dataset.
- **train_sqr_err:** Mean-squared error on the training dataset.
- **train_cross_entropy:** Cross-entropy on the training dataset.
- **valid_class_err:** Classification error on the validation dataset.
- **valid_sqr_err:** Mean-squared error on the validation dataset.
- **valid_cross_entropy:** Cross-entropy on the validation dataset.

Extends

Class "[H2OModel](#)", directly.

Methods

show `signature(object = "H2ONNModel")`: ...

See Also

[h2o.nn](#)

Examples

```
showClass("H2ONNModel")
```

H2OParsedData-class *Class "H2OParsedData"*

Description

A class for representing imported FluidVecs data sets that have been parsed.

Objects from the Class

Objects can be created by calls of the form `new("H2OParsedData", ...)`.

Slots

h2o: Object of class "H2OClient", which is the client object that was passed into the function call.

key: Object of class "character", which is the hex key assigned to the imported data.

logic: Object of class "logical", indicating whether the "H2OParsedData" object represents logical data

Methods

```
- signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
- signature(e1 = "H2OParsedData", e2 = "numeric"): ...
- signature(e1 = "numeric", e2 = "H2OParsedData"): ...
! signature(x = "H2OParsedData"): ...
!= signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
!= signature(e1 = "H2OParsedData", e2 = "numeric"): ...
!= signature(e1 = "numeric", e2 = "H2OParsedData"): ...
[ signature(x = "H2OParsedData"): ...
[<- signature(x = "H2OParsedData"): ...
* signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
* signature(e1 = "H2OParsedData", e2 = "numeric"): ...
* signature(e1 = "numeric", e2 = "H2OParsedData"): ...
/ signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
/ signature(e1 = "H2OParsedData", e2 = "numeric"): ...
/ signature(e1 = "numeric", e2 = "H2OParsedData"): ...
& signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
& signature(e1 = "H2OParsedData", e2 = "numeric"): ...
& signature(e1 = "numeric", e2 = "H2OParsedData"): ...
%% signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
%% signature(e1 = "H2OParsedData", e2 = "numeric"): ...
```

```

%% signature(e1 = "numeric", e2 = "H2OParsedData"): ...
+ signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
+ signature(e1 = "H2OParsedData", e2 = "numeric"): ...
+ signature(e1 = "numeric", e2 = "H2OParsedData"): ...
< signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
< signature(e1 = "H2OParsedData", e2 = "numeric"): ...
< signature(e1 = "numeric", e2 = "H2OParsedData"): ...
<= signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
<= signature(e1 = "H2OParsedData", e2 = "numeric"): ...
<= signature(e1 = "numeric", e2 = "H2OParsedData"): ...
== signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
== signature(e1 = "H2OParsedData", e2 = "numeric"): ...
== signature(e1 = "numeric", e2 = "H2OParsedData"): ...
> signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
> signature(e1 = "H2OParsedData", e2 = "numeric"): ...
> signature(e1 = "numeric", e2 = "H2OParsedData"): ...
>= signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
>= signature(e1 = "H2OParsedData", e2 = "numeric"): ...
>= signature(e1 = "numeric", e2 = "H2OParsedData"): ...
| signature(e1 = "H2OParsedData", e2 = "H2OParsedData"): ...
| signature(e1 = "H2OParsedData", e2 = "numeric"): ...
| signature(e1 = "numeric", e2 = "H2OParsedData"): ...
$ signature(x = "H2OParsedData"): ...
$<- signature(x = "H2OParsedData"): ...
abs signature(x = "H2OParsedData"): ...
apply signature(X = "H2OParsedData"): ...
as.data.frame signature(x = "H2OParsedData"): ...
as.factor signature(x = "H2OParsedData"): ...
ceiling signature(x = "H2OParsedData"): ...
colMeans signature(x = "H2OParsedData"): ...
colnames signature(x = "H2OParsedData"): ...
colnames<- signature(x = "H2OParsedData", value = "character"): ...
colnames<- signature(x = "H2OParsedData", value = "H2OParsedData"): ...
dim signature(x = "H2OParsedData"): ...
dim<- signature(x = "H2OParsedData"): ...
exp signature(x = "H2OParsedData"): ...
floor signature(x = "H2OParsedData"): ...

```

```
h2o.cut signature(x = "H2OParsedData", breaks = "numeric"): ...  
h2o<- signature(x = "H2OParsedData", value = "H2OParsedData"): ...  
h2o<- signature(x = "H2OParsedData", value = "numeric"): ...  
head signature(x = "H2OParsedData"): ...  
histograms signature(object = "H2OParsedData"): ...  
ifelse signature(test = "H2OParsedData"): ...  
is.factor signature(x = "H2OParsedData"): ...  
is.na signature(x = "H2OParsedData"): ...  
length signature(x = "H2OParsedData"): ...  
levels signature(x = "H2OParsedData"): ...  
log signature(x = "H2OParsedData"): ...  
names signature(x = "H2OParsedData"): ...  
names<- signature(x = "H2OParsedData"): ...  
ncol signature(x = "H2OParsedData"): ...  
nrow signature(x = "H2OParsedData"): ...  
quantile signature(x = "H2OParsedData"): ...  
range signature(x = "H2OParsedData"): ...  
sd signature(x = "H2OParsedData"): ...  
show signature(object = "H2OParsedData"): ...  
sign signature(x = "H2OParsedData"): ...  
sqrt signature(x = "H2OParsedData"): ...  
summary signature(object = "H2OParsedData"): ...  
tail signature(x = "H2OParsedData"): ...
```

See Also

[H2ORawData](#), [h2o.parseRaw](#)

Examples

```
showClass("H2OParsedData")
```

H2OParsedDataVA-class *Class "H2OParsedDataVA"*

Description

A class for representing imported ValueArray data sets that have been parsed.

Objects from the Class

Objects can be created by calls of the form `new("H2OParsedDataVA", ...)`.

Slots

h2o: Object of class "H2OClient", which is the client object that was passed into the function call.

key: Object of class "character", which is the hex key assigned to the imported data.

logic: Object of class "logical", indicating whether the "H2OParsedDataVA" object represents logical data

Extends

Class "[H2OParsedData](#)", directly.

Methods

```
colnames signature(x = "H2OParsedDataVA"): ...
colnames<- signature(x = "H2OParsedDataVA", value = "character"): ...
colnames<- signature(x = "H2OParsedDataVA", value = "H2OParsedDataVA"): ...
dim signature(x = "H2OParsedDataVA"): ...
head signature(x = "H2OParsedDataVA"): ...
names signature(x = "H2OParsedDataVA"): ...
names<- signature(x = "H2OParsedDataVA"): ...
ncol signature(x = "H2OParsedDataVA"): ...
nrow signature(x = "H2OParsedDataVA"): ...
show signature(object = "H2OParsedDataVA"): ...
summary signature(object = "H2OParsedDataVA"): ...
tail signature(x = "H2OParsedDataVA"): ...
```

See Also

[H2ORawDataVA](#), [h2o.parseRaw.VA](#)

Examples

```
showClass("H2OParsedDataVA")
```

H2OPCAModel-class	Class "H2OPCAModel"
-------------------	---------------------

Description

A class for representing principal components analysis results.

Objects from the Class

Objects can be created by calls of the form `new("H2OPCAModel", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedData](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **standardized:** A logical value indicating whether the data was centered and scaled.
- **sdev:** The standard deviations of the principal components (i.e., the square roots of the eigenvalues of the covariance/correlation matrix).
- **rotation:** The matrix of variable loadings (i.e., a matrix whose columns contain the eigenvectors).

Extends

Class "[H2OModel](#)", directly.

Methods

show signature(object = "H2OPCAModel"): ...

plot signature(x = "H2OPCAModel", y, ...): ...

summary signature(object = "H2OPCAModel"): ...

See Also

[h2o.prcomp](#)

Examples

```
showClass("H2OPCAModel")
```

H2ORawData-class	Class "H2ORawData"
------------------	--------------------

Description

A class for representing imported FluidVecs data sets that have not been parsed.

Objects from the Class

Objects can be created by calls of the form `new("H2ORawData", ...)`.

Slots

- h2o:** Object of class "H2OClient", which is the client object that was passed into the function call.
- key:** Object of class "character", which is the hex key assigned to the imported data.

Methods

h2o.parseRaw signature(data = "H2OParsedData", key = "character", header = "logical" sep = "character", ...
show signature(object = "H2ORawData"): ...

See Also

[H2OParsedData](#)

Examples

`showClass("H2ORawData")`

H2ORawDataVA-class	Class "H2ORawDataVA"
--------------------	----------------------

Description

A class for representing imported ValueArray data sets that have not been parsed.

Objects from the Class

Objects can be created by calls of the form `new("H2ORawDataVA", ...)`.

Slots

- h2o:** Object of class "H2OClient", which is the client object that was passed into the function call.
- key:** Object of class "character", which is the hex key assigned to the imported data.

Methods

```

h2o.parseRaw.VA signature(data = "H2OParsedDataVA", key = "character", header = "logical" sep = "charac
...
show signature(object = "H2ORawDataVA"): ...

```

See Also

[H2OParsedDataVA](#)

Examples

```
showClass("H2ORawDataVA")
```

H2ORFModelVA-class	Class "H2ORFModelVA"
--------------------	----------------------

Description

A class for representing random forest ensembles built on ValueArray data.

Objects from the Class

Objects can be created by calls of the form `new("H2ORFModelVA", ...)`.

Slots

key: Object of class "character", representing the unique hex key that identifies the model.

data: Object of class [H2OParsedDataVA](#), which is the input data used to build the model.

model: Object of class "list" containing the following elements:

- **type:** The type of the tree, which at this point must be classification.
- **ntree:** Number of trees grown.
- **oob_err:** Out of bag error rate.
- **forest:** A matrix giving the minimum, mean, and maximum of the tree depth and number of leaves.
- **confusion:** Confusion matrix of the prediction.

Extends

Class "[H2OModelVA](#)", directly.

Methods

```
show signature(object = "H2ORFModelVA"): ...
```

See Also

[h2o.randomForest](#)

Examples

```
showClass("H2ORFModelVA")
```

head

Return the First or Last Part of a H2O Dataset

Description

Returns the first or last rows of an H2O parsed data object.

Usage

```
## S3 method for class 'H2OParsedData'
head(x, n = 6L, ...)
## S3 method for class 'H2OParsedData'
tail(x, n = 6L, ...)
## S3 method for class 'H2OParsedDataVA'
head(x, n = 6L, ...)
## S3 method for class 'H2OParsedDataVA'
tail(x, n = 6L, ...)
```

Arguments

x	An H2O parsed data object.
n	(Optional) A single integer. If positive, number of rows in x to return. If negative, all but the n first/last number of rows in x.
...	Arguments to be passed to or from other methods. (Currently unimplemented).

Value

A data frame containing the first or last n rows of an [H2OParsedData](#) object.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
head(australia.hex, 10)
tail(australia.hex, 10)
h2o.shutdown(localH2O)
```

ifelse	<i>Applies conditional statements to an H2OParsedData object.</i>
--------	---

Description

Applies conditional statements to numeric vectors in H2O parsed data objects when the data are numeric.

Usage

```
ifelse(test, yes, no)
```

Arguments

test	A logical description of the condition to be met (>, <, =, etc...)
yes	The value to return if the condition is TRUE.
no	The value to return if the condition is FALSE.

Details

Only numeric values can be tested, and only numeric results can be returned for either condition. Categorical data is not currently supported for this function and returned values cannot be categorical in nature.

Value

Retruns a vector of new values matching the conditions stated in the ifelse call.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
australia.hex[,9] <- ifelse(australia.hex[,3] < 279.9, 1, 0)
summary(australia.hex)
h2o.shutdown(localH2O)
```

<code>is.factor</code>	<i>Tells user if given column is categorical data or not.</i>
------------------------	---

Description

Tells user if given column is categorical data or not.

Usage

```
is.factor(x)
```

Arguments

`x` Columns of an H2O parsed data object.

Value

A logical value TRUE if column contains categorical data, FALSE otherwise.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
prostate.hex[,4]=as.factor(prostate.hex[,4])
is.factor(prostate.hex[,4])
is.factor(prostate.hex[,3])
h2o.shutdown(localH2O)
```

<code>levels</code>	<i>Levels of Categorical Data</i>
---------------------	-----------------------------------

Description

Returns a list of the unique values found in a column of categorical data.

Usage

```
levels(x)
```

Arguments

`x` Column of categorical data in an [H2OParsedData](#) object.

Value

Returns a list containing one entry for each unique value found in the column of categorical data.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
levels(iris.hex[,5])
h2o.shutdown(localH2O)
```

mean.H2OParsedData	<i>Arithmetic Mean of H2O Dataset</i>
--------------------	---------------------------------------

Description

mean.H2OParsedData, a method for the [mean](#) generic. Calculate the mean of each numeric column in a H2O dataset.

Usage

```
## S3 method for class 'H2OParsedData'
mean(x, trim = 0, na.rm = FALSE, ...)
```

Arguments

x	An H2OParsedData object.
trim	(The fraction (0 to 0.5) of observations to trim from each end of x before the mean is computed. (Currently unimplemented).
na.rm	Logical value indicating whether NA or missing values should be stripped before the computation.
...	Potential further arguments. (Currently unimplemented).

Value

An [H2OParsedData](#) object of scalar numeric value representing the arithmetic mean of each numeric column of x. If x is not logical or numeric, then NA_real_ is returned, with a warning.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
mean(prostate.hex$AGE)
h2o.shutdown(localH2O)
```

nrow	<i>The Number of Rows/Columns of a H2O Dataset</i>
------	--

Description

Returns a count of the number of rows in an [H2OParsedData](#) object.

Usage

```
nrow(x)
ncol(x)
```

Arguments

x An [H2OParsedData](#) object.

Value

An integer of length 1 indicating the number of rows or columns in the dataset.

See Also

[dim](#) which returns all dimensions

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
nrow(iris.hex)
ncol(iris.hex)
h2o.shutdown(localH2O)
```

quantile.H2OParsedData	<i>Obtain and display quantiles for H2O parsed data.</i>
------------------------	--

Description

quantile.H2OParsedData, a method for the [quantile](#) generic. Obtain and return quantiles for an [H2OParsedData](#) object.

Usage

```
## S3 method for class 'H2OParsedData'
quantile(x, probs = seq(0, 1, 0.25), na.rm = FALSE, names = TRUE, ...)
```

Arguments

x	An H2OParsedData object with a single numeric column.
probs	numeric vector of probabilities with values in [0,1].
na.rm	logical; if true, any NA and NaN's are removed from x before the quantiles are computed.
names	logical; if true, the result has a names attribute.
...	further arguments passed to or from other methods.

Details

Note that H2O parsed data objects can be quite large, and are therefore often distributed across multiple nodes in an H2O cluster. As a result, percentiles at the 1st, 5th, 10th, 25th, 33, 50, 66, 75, 90, 95, 99th, and other values cannot be returned. This range includes the 1st quantile at the 25th percentile, median at the 50th percentile, and 3rd quantile at the 75th percentile.

Value

A vector describing the percentiles at the given cutoffs for the [H2OParsedData](#) object.

Examples

```
# Request quantiles for an H2O parsed data set:
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)

# Request quantiles for a subset of columns in an H2O parsed data set
quantile(prostate.hex[,3])
for(i in 1:ncol(prostate.hex))
  quantile(prostate.hex[,i])
h2o.shutdown(localH2O)
```

screepLOT.H2OPCAModel *Summarizes the columns of an H2O parsed FluidVecs data set.*

Description

screepLOT.H2OPCAModel, a method for the [screepLOT](#) generic. Plots the variances against the number of the principal component generated by [h2o.prcomp](#).

Usage

```
## S3 method for class 'H2OPCAModel'
screepLOT(x, npcs = min(10, length(x@model$sdev)), type = "barplot",
  main = paste("h2o.prcomp(", x@data@key, ")", sep=""), ...)
```

Arguments

x	An H2OPCAModel object.
npcs	Number of components to be plotted.
type	Type of plot, must be either "barplot" or "lines".
main	Title of the plot.
...	Additional parameters to be passed to the plotting function.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package = "h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
australia.pca = h2o.prcomp(data = australia.hex, standardize = TRUE)
screeplot(australia.pca)
h2o.shutdown(localH2O)
```

sd	<i>Standard Deviation of a Numeric Column of H2O Data</i>
----	---

Description

Calculates the standard deviation of a [H2OParsedData](#) column of continuous real valued data.

Usage

```
sd(x, na.rm = FALSE)
```

Arguments

x	An H2OParsedData object containing numeric data.
na.rm	Logical value where FALSE does not remove NA's in the calculation and TRUE removes NA's in the calculation.

Value

Returns a vector of values of the standard deviations for the requested columns.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
irisPath = system.file("extdata", "iris.csv", package="h2o")
iris.hex = h2o.importFile(localH2O, path = irisPath, key = "iris.hex")
sd(iris.hex[,4])
h2o.shutdown(localH2O)
```

str

*Display the Structure of a H2O Dataset***Description**

A method for the [str](#) generic. Obtain information about H2O parsed data objects and their structure.

Usage

```
## S3 method for class 'H2OParsedData'
str(object, ...)
## S3 method for class 'H2OParsedDataVA'
str(object, ...)
```

Arguments

object An [H2OParsedData](#) object.

... Potential further arguments. (Currently unimplemented).

Value

A table listing summary information including variable names, types (for example, enum or numeric), count of observations and columns.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
str(prostate.hex)
h2o.shutdown(localH2O)
```

sum

*Sum of Numeric Values***Description**

Calculates the sum of all the values present in its arguments. This method extends the [sum](#) generic to deal with [H2OParsedData](#) objects.

Usage

```
sum(..., na.rm = FALSE)
```

Arguments

... Numeric, complex, logical or [H2OParsedData](#) arguments.

na.rm Logical value where FALSE does not remove NA's in the calculation and TRUE removes NA's in the calculation.

Value

Returns the sum over all the input arguments. For a [H2OParsedData](#) object, the sum is taken over all entries in the dataset. An error will occur if any of those entries is non-numeric.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath, key = "australia.hex")
sum(australia.hex)
sum(c(400, 1234, -1250), TRUE, australia.hex[,1:4])
h2o.shutdown(localH2O)
```

summary

*Summarizes the columns of a H2O Dataset***Description**

A method for the [summary](#) generic. Summarizes the columns of an H2O parsed object or subset of columns and rows using vector notation (e.g. dataset[row, col])

Usage

```
## S3 method for class 'H2OParsedData'
summary(object, ...)
## S3 method for class 'H2OParsedDataVA'
summary(object, ...)
```

Arguments

object An [H2OParsedData](#) object.

... Additional arguments affecting the summary produced. (Currently unimplemented).

Value

A matrix displaying the minimum, 1st quartile, median, mean, 3rd quartile and maximum for each numeric column included in the request call, a summary of the levels and member counts for each factor column. and a the levels and member counts of the elements in factor columns for all of the columns specified in the summary call.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
prosPath = system.file("extdata", "prostate.csv", package="h2o")
prostate.hex = h2o.importFile(localH2O, path = prosPath)
summary(prostate.hex)
summary(prostate.hex$GLEASON)
summary(prostate.hex[,4:6])
h2o.shutdown(localH2O)
```

summary.H2OPCAModel	<i>Summarizes the H2O PCA Model</i>
---------------------	-------------------------------------

Description

summary.H2OPCAModel, a method for the [summary](#) generic. Summarizes the importance of each principal component returned by [h2o.prcomp](#).

Usage

```
## S3 method for class 'H2OPCAModel'
summary(object, ...)
```

Arguments

object	An H2OPCAModel object.
...	Additional arguments affecting the summary produced. (Currently unimplemented).

Value

A matrix displaying the standard deviation, proportion of variance explained and cumulative proportion of variance explained by each principal component.

Examples

```
library(h2o)
localH2O = h2o.init(ip = "localhost", port = 54321, startH2O = TRUE)
ausPath = system.file("extdata", "australia.csv", package="h2o")
australia.hex = h2o.importFile(localH2O, path = ausPath)
australia.pca = h2o.prcomp(data = australia.hex, standardize = TRUE)
summary(australia.pca)
h2o.shutdown(localH2O)
```