

Package ‘CAST’

April 7, 2021

Type Package

Title 'caret' Applications for Spatial-Temporal Models

Version 0.5.1

Author Hanna Meyer [cre, aut],
Chris Reudenbach [ctb],
Marvin Ludwig [ctb],
Thomas Nauss [ctb],
Edzer Pebesma [ctb]

Maintainer Hanna Meyer <hanna.meyer@uni-muenster.de>

Description Supporting functionality to run 'caret' with spatial or spatial-temporal data. 'caret' is a frequently used package for model training and prediction using machine learning. This package includes functions to improve spatial-temporal modelling tasks using 'caret'. It prepares data for Leave-Location-Out and Leave-Time-Out cross-validation which are target-oriented validation strategies for spatial-temporal models. To decrease overfitting and improve model performances, the package implements a forward feature selection that selects suitable predictor variables in view to their contribution to the target-oriented performance. CAST further includes functionality to estimate the (spatial) area of applicability of prediction models by analysing the similarity between new data and training data.

License GPL (>= 2)

URL <https://github.com/HannaMeyer/CAST>

Encoding UTF-8

LazyData false

Depends R (>= 3.5.0)

Imports caret, stats, utils, ggplot2, graphics, reshape, FNN, plyr,
zoo, methods, grDevices, data.table, lattice

Suggests doParallel, randomForest, lubridate, raster, sp, knitr,
mapview, rmarkdown, sf, scales, parallel, latticeExtra,
virtualspecies, gridExtra, viridis, rgeos, stars, scam

RoxygenNote 7.1.0

VignetteBuilder knitr

NeedsCompilation no

Repository CRAN

Date/Publication 2021-04-07 13:20:17 UTC

R topics documented:

aoa	2
bss	5
calibrate_aoa	6
CAST	9
CreateSpacetimeFolds	9
ffs	11
plot_ffs	14

Index **16**

aoa *Area of Applicability*

Description

This function estimates the Dissimilarity Index (DI) and the derived Area of Applicability (AOA) of spatial prediction models by considering the distance of new data (i.e. a Raster Stack of spatial predictors used in the models) in the predictor variable space to the data used for model training. Predictors can be weighted based on the internal variable importance of the machine learning algorithm used for model training. The AOA is derived by applying a threshold on the DI which is the (outlier-removed) maximum DI of the cross-validated training data.

Usage

```
aoa(
  newdata,
  model = NA,
  cl = NULL,
  train = NULL,
  weight = NA,
  variables = "all",
  folds = NULL,
  returnTrainDI = TRUE
)
```

Arguments

newdata	A RasterStack, RasterBrick or data.frame containing the data the model was meant to make predictions for.
model	A train object created with caret used to extract weights from (based on variable importance) as well as cross-validation folds

cl	A cluster object e.g. created with doParallel. Should only be used if newdata is large.
train	A data.frame containing the data used for model training. Only required when no model is given
weight	A data.frame containing weights for each variable. Only required if no model is given.
variables	character vector of predictor variables. if "all" then all variables of the model are used or if no model is given then of the train dataset.
folds	Numeric or character. Folds for cross validation. E.g. Spatial cluster affiliation for each data point. Should be used if replicates are present. Only required if no model is given.
returnTrainDI	A logical: should the DI value of the cross-validated training data be returned as a attribute?

Details

The Dissimilarity Index (DI) and the corresponding Area of Applicability (AOA) are calculated. If variables are factors, dummy variables are created prior to weighting and distance calculation.

Interpretation of results: If a location is very similar to the properties of the training data it will have a low distance in the predictor variable space (DI towards 0) while locations that are very different in their properties will have a high DI. See Meyer and Pebesma (2020) for the full documentation of the methodology.

Value

A RasterStack or data.frame with the DI and AOA AOA has values 0 (outside AOA) and 1 (inside AOA).

Note

If classification models are used, currently the variable importance can only be automatically retrieved if models were trained via `train(predictors,response)` and not via the formula-interface. Will be fixed.

Author(s)

Hanna Meyer

References

Meyer, H., Pebesma, E. (2020): Predicting into unknown space? Estimating the area of applicability of spatial prediction models. <https://arxiv.org/abs/2005.07939>

See Also

[calibrate_aoa](#)

Examples

```

## Not run:
library(sf)
library(raster)
library(caret)
library(viridis)
library(latticeExtra)

# prepare sample data:
dat <- get(load(system.file("extdata", "Cookfarm.RData", package="CAST")))
dat <- aggregate(dat[,c("VW", "Easting", "Northing")], by=list(as.character(dat$SOURCEID)), mean)
pts <- st_as_sf(dat, coords=c("Easting", "Northing"))
pts$ID <- 1:nrow(pts)
set.seed(100)
pts <- pts[1:30,]
studyArea <- stack(system.file("extdata", "predictors_2012-03-25.grd", package="CAST"))[[1:8]]
trainDat <- extract(studyArea, pts, df=TRUE)
trainDat <- merge(trainDat, pts, by.x="ID", by.y="ID")

# visualize data spatially:
spplot(scale(studyArea))
plot(studyArea$DEM)
plot(pts[,1], add=TRUE, col="black")

# train a model:
set.seed(100)
variables <- c("DEM", "NDRE.Sd", "TWI")
model <- train(trainDat[,which(names(trainDat)%in%variables)],
  trainDat$VW, method="rf", importance=TRUE, tuneLength=1,
  trControl=trainControl(method="cv", number=5, savePredictions=T))
print(model) #note that this is a quite poor prediction model
prediction <- predict(studyArea, model)
plot(varImp(model, scale=FALSE))

#...then calculate the AOA of the trained model for the study area:
AOA <- aoa(studyArea, model)
spplot(AOA$DI, col.regions=viridis(100), main="Dissimilarity Index")
#plot predictions for the AOA only:
spplot(prediction, col.regions=viridis(100), main="prediction for the AOA")+
spplot(AOA$AOA, col.regions=c("grey", "transparent"))

####
# Calculating the AOA might be time consuming. Consider running it in parallel:
####
library(doParallel)
library(parallel)
cl <- makeCluster(4)
registerDoParallel(cl)
AOA <- aoa(studyArea, model, cl=cl)

####
#The AOA can also be calculated without a trained model.

```

```
#All variables are weighted equally in this case:
####
AOA <- aoa(studyArea,train=trainDat,variables=variables)
spplot(AOA$DI, col.regions=viridis(100),main="Dissimilarity Index")
spplot(AOA$AOA,main="Area of Applicability")

## End(Not run)
```

bss

Best subset feature selection

Description

Evaluate all combinations of predictors during model training

Usage

```
bss(
  predictors,
  response,
  method = "rf",
  metric = ifelse(is.factor(response), "Accuracy", "RMSE"),
  maximize = ifelse(metric == "RMSE", FALSE, TRUE),
  trControl = caret::trainControl(),
  tuneLength = 3,
  tuneGrid = NULL,
  seed = 100,
  verbose = TRUE,
  ...
)
```

Arguments

predictors	see train
response	see train
method	see train
metric	see train
maximize	see train
trControl	see train
tuneLength	see train
tuneGrid	see train
seed	A random number
verbose	Logical. Should information about the progress be printed?
...	arguments passed to the classification or regression routine (such as <code>randomForest</code>).

Details

bss is an alternative to [ffs](#) and ideal if the training set is small. Models are iteratively fitted using all different combinations of predictor variables. Hence, 2^X models are calculated. Dont try running bss on very large datasets because the computation time is much higher compared to [ffs](#).

The internal cross validation can be run in parallel. See information on parallel processing of caret's train functions for details.

Value

A list of class train. Beside of the usual train content the object contains the vector "selectedvars" and "selectedvars_perf" that give the best variables selected as well as their corresponding performance. It also contains "perf_all" that gives the performance of all model runs.

Note

This validation is particularly suitable for spatial leave-location-out cross validations where variable selection **MUST** be based on the performance of the model on the hold out station. Note that bss is very slow since all combinations of variables are tested. A more time efficient alternative is the forward feature selection ([ffs](#)) ([ffs](#)).

Author(s)

Hanna Meyer

See Also

[train](#), [ffs](#), [trainControl](#), [CreateSpacetimeFolds](#)

Examples

```
## Not run:
data(iris)
bssmodel <- bss(iris[,1:4],iris$Species)
bssmodel$perf_all

## End(Not run)
```

calibrate_aoa

Calibrate the AOA based on the relationship between the DI and the prediction error

Description

Performance metrics are calculated for moving windows of DI values of cross-validated training data

Usage

```
calibrate_aoa(
  AOA,
  model,
  window.size = 5,
  calib = "scam",
  multiCV = FALSE,
  length.out = 10,
  maskAOA = TRUE,
  showPlot = TRUE,
  k = 6,
  m = 2
)
```

Arguments

AOA	the result of aoa
model	the model used to get the AOA
window.size	Numeric. Size of the moving window. See rollapply .
calib	Character. Function to model the DI~performance relationship. Currently lm and scam are supported
multiCV	Logical. Re-run model fitting and validation with different CV strategies. See details.
length.out	Numeric. Only used if multiCV=TRUE. Number of cross-validation folds. See details.
maskAOA	Logical. Should areas outside the AOA set to NA?
showPlot	Logical.
k	Numeric. See <code>mgcv::s</code>
m	Numeric. See <code>mgcv::s</code>

Details

If multiCV=TRUE the model is re-fitted and validated by length.out new cross-validations where the cross-validation folds are defined by clusters in the predictor space, ranging from three clusters to LOOCV. If the AOA threshold based on the calibration data from multiple CV is larger than the original AOA threshold, the AOA is updated accordingly. See Meyer and Pebesma (2020) for the full documentation of the methodology.

Value

A list of length 2 with the elements "AOA": rasterStack which contains the original DI and the AOA (which might be updated if new test data indicate this option), as well as the expected performance based on the relationship. Data used for calibration are stored in the attributes. The second element is a plot showing the relationship.

Author(s)

Hanna Meyer

References

Meyer, H., Pebesma, E. (2020): Predicting into unknown space? Estimating the area of applicability of spatial prediction models. <https://arxiv.org/abs/2005.07939>

See Also

[aoa](#)

Examples

```
## Not run:
library(sf)
library(raster)
library(caret)
library(viridis)
library(latticeExtra)

# prepare sample data:
library(sf)
library(raster)
library(caret)
# prepare sample data:
dat <- get(load(system.file("extdata", "Cookfarm.RData", package="CAST")))
dat <- aggregate(dat[,c("VW", "Easting", "Northing")], by=list(as.character(dat$SOURCEID)), mean)
pts <- st_as_sf(dat, coords=c("Easting", "Northing"))
pts$ID <- 1:nrow(pts)
studyArea <- stack(system.file("extdata", "predictors_2012-03-25.grd", package="CAST"))[[1:8]]
dat <- extract(studyArea, pts, df=TRUE)
trainDat <- merge(dat, pts, by.x="ID", by.y="ID")

# train a model:
variables <- c("DEM", "NDRE.Sd", "TWI")
set.seed(100)
model <- train(trainDat[,which(names(trainDat)%in%variables)],
  trainDat$VW, method="rf", importance=TRUE, tuneLength=1,
  trControl=trainControl(method="cv", number=5, savePredictions=TRUE))

#...then calculate the AOA of the trained model for the study area:
AOA <- aoa(studyArea, model)

AOA_new <- calibrate_aoa(AOA, model)
plot(AOA_new$AOA[[3]])

## End(Not run)
```


Description

Supporting functionality to run 'caret' with spatial or spatial-temporal data. 'caret' is a frequently used package for model training and prediction using machine learning. CAST includes functions to improve spatial-temporal modelling tasks using 'caret'. It supports Leave-Location-Out and Leave-Time-Out cross-validation of spatial and spatial-temporal models and allows for spatial variable selection to select suitable predictor variables in view to their contribution to the spatial model performance. CAST further includes functionality to estimate the (spatial) area of applicability of prediction models by analysing the similarity between new data and training data.

Details

'caret' Applications for Spatio-Temporal models

Author(s)

Hanna Meyer

References

- Meyer, H., Pebesma, E. (2020): Predicting into unknown space? Estimating the area of applicability of spatial prediction models. <https://arxiv.org/abs/2005.07939>.
- Meyer, H., Reudenbach, C., Wöllauer, S., Nauss, T. (2019): Importance of spatial predictor variable selection in machine learning applications - Moving from data reproduction to spatial prediction. *Ecological Modelling*. 411, 108815.
- Meyer, H., Reudenbach, C., Hengl, T., Katurji, M., Nauß, T. (2018): Improving performance of spatio-temporal machine learning models using forward feature selection and target-oriented validation. *Environmental Modelling & Software* 101: 1-9.

Description

Create spatial, temporal or spatio-temporal Folds for cross validation

Usage

```
CreateSpacetimeFolds(  
  x,  
  spacevar = NA,  
  timevar = NA,  
  k = 10,  
  class = NA,  
  seed = sample(1:1000, 1)  
)
```

Arguments

x	data.frame containing spatio-temporal data
spacevar	Character indicating which column of x identifies the spatial units (e.g. ID of weather stations)
timevar	Character indicating which column of x identifies the temporal units (e.g. the day of the year)
k	numeric. Number of folds. If spacevar or timevar is NA and a leave one location out or leave one time step out cv should be performed, set k to the number of unique spatial or temporal units.
class	Character indicating which column of x identifies a class unit (e.g. land cover)
seed	numeric. See ?seed

Details

Using "class" is helpful in the case that data are clustered in space and are categorical. E.g This is the case for land cover classifications when training data come as training polygons. In this case the data should be split in a way that entire polygons are held back (spacevar="polygonID") but at the same time the distribution of classes should be similar in each fold (class="LUC").

Value

A list that contains a list for model training and a list for model validation that can directly be used as "index" and "indexOut" in caret's trainControl function

Note

Standard k-fold cross-validation can lead to considerable misinterpretation in spatial-temporal modelling tasks. This function can be used to prepare a Leave-Location-Out, Leave-Time-Out or Leave-Location-and-Time-Out cross-validation as target-oriented validation strategies for spatial-temporal prediction tasks. See Meyer et al. (2018) for further information.

Author(s)

Hanna Meyer

References

Meyer, H., Reudenbach, C., Hengl, T., Katurji, M., Nauß, T. (2018): Improving performance of spatio-temporal machine learning models using forward feature selection and target-oriented validation. *Environmental Modelling & Software* 101: 1-9.

See Also

[trainControl,ffs](#)

Examples

```
dat <- get(load(system.file("extdata", "Cookfarm.RData", package="CAST")))
### Prepare for 10-fold Leave-Location-and-Time-Out cross validation
indices <- CreateSpacetimeFolds(dat, "SOURCEID", "Date")
str(indices)
### Prepare for 10-fold Leave-Location-Out cross validation
indices <- CreateSpacetimeFolds(dat, spacevar="SOURCEID")
str(indices)
### Prepare for leave-One-Location-Out cross validation
indices <- CreateSpacetimeFolds(dat, spacevar="SOURCEID",
k=length(unique(dat$SOURCEID)))
str(indices)
```

ffs

Forward feature selection

Description

A simple forward feature selection algorithm

Usage

```
ffs(
  predictors,
  response,
  method = "rf",
  metric = ifelse(is.factor(response), "Accuracy", "RMSE"),
  maximize = ifelse(metric == "RMSE", FALSE, TRUE),
  withinSE = FALSE,
  minVar = 2,
  trControl = caret::trainControl(),
  tuneLength = 3,
  tuneGrid = NULL,
  seed = sample(1:1000, 1),
  verbose = TRUE,
  ...
)
```

Arguments

predictors	see train
response	see train
method	see train
metric	see train
maximize	see train
withinSE	Logical. Models are only selected if they are better than the currently best models Standard error
minVar	Numeric. Number of variables to combine for the first selection. See Details.
trControl	see train
tuneLength	see train
tuneGrid	see train
seed	A random number used for model training
verbose	Logical. Should information about the progress be printed?
...	arguments passed to the classification or regression routine (such as <code>randomForest</code>).

Details

Models with two predictors are first trained using all possible pairs of predictor variables. The best model of these initial models is kept. On the basis of this best model the predictor variables are iteratively increased and each of the remaining variables is tested for its improvement of the currently best model. The process stops if none of the remaining variables increases the model performance when added to the current best model.

The internal cross validation can be run in parallel. See information on parallel processing of caret's `train` functions for details.

Using `withinSE` will favour models with less variables and probably shorten the calculation time

Per Default, the `ffs` starts with all possible 2-pair combinations. `minVar` allows to start the selection with more than 2 variables, e.g. `minVar=3` starts the `ffs` testing all combinations of 3 (instead of 2) variables first and then increasing the number. This is important for e.g. neural networks that often cannot make sense of only two variables. It is also relevant if it is assumed that the optimal variables can only be found if more than 2 are considered at the same time.

Value

A list of class `train`. Beside of the usual `train` content the object contains the vector `"selectedvars"` and `"selectedvars_perf"` that give the order of the best variables selected as well as their corresponding performance (starting from the first two variables). It also contains `"perf_all"` that gives the performance of all model runs.

Note

This validation is particularly suitable for spatial leave-location-out cross validations where variable selection **MUST** be based on the performance of the model on the hold out station. See Meyer et al. (2018) and Meyer et al. (2019) for further details.

Author(s)

Hanna Meyer

References

- Gasch, C.K., Hengl, T., Gräler, B., Meyer, H., Magney, T., Brown, D.J. (2015): Spatio-temporal interpolation of soil water, temperature, and electrical conductivity in 3D+T: the Cook Agronomy Farm data set. *Spatial Statistics* 14: 70-90.
- Meyer, H., Reudenbach, C., Hengl, T., Katurji, M., Nauß, T. (2018): Improving performance of spatio-temporal machine learning models using forward feature selection and target-oriented validation. *Environmental Modelling & Software* 101: 1-9. doi: [10.1016/j.envsoft.2017.12.001](https://doi.org/10.1016/j.envsoft.2017.12.001)
- Meyer, H., Reudenbach, C., Wöllauer, S., Nauss, T. (2019): Importance of spatial predictor variable selection in machine learning applications - Moving from data reproduction to spatial prediction. *Ecological Modelling*. 411, 108815. doi: [10.1016/j.ecolmodel.2019.108815](https://doi.org/10.1016/j.ecolmodel.2019.108815)

See Also[train,bss](#), [trainControl](#), [CreateSpacetimeFolds](#)**Examples**

```
## Not run:
data(iris)
ffsmodel <- ffs(iris[,1:4],iris$Species)
ffsmodel$selectedvars
ffsmodel$selectedvars_perf

## End(Not run)

# or perform model with target-oriented validation (LLO CV)
#the example is described in Gasch et al. (2015). The ffs approach for this dataset is described in
#Meyer et al. (2018). Due to high computation time needed, only a small and thus not robust example
#is shown here.

## Not run:
#run the model on three cores:
library(doParallel)
cl <- makeCluster(3)
registerDoParallel(cl)

#load and prepare dataset:
dat <- get(load(system.file("extdata", "Cookfarm.RData", package="CAST")))
trainDat <- dat[dat$altitude==0.3&year(dat$Date)==2012&week(dat$Date)%in%c(13:14),]

#visualize dataset:
ggplot(data = trainDat, aes(x=Date, y=VW)) + geom_line(aes(colour=SOURCEID))

#create folds for Leave Location Out Cross Validation:
set.seed(10)
indices <- CreateSpacetimeFolds(trainDat,spacevar = "SOURCEID",k=3)
```

```

ctrl <- trainControl(method="cv",index = indices$index)

#define potential predictors:
predictors <- c("DEM","TWI","BLD","Precip_cum","cday","MaxT_wrcc",
"Precip_wrcc","NDRE.M","Bt","MinT_wrcc","Northing","Easting")

#run ffs model with Leave Location out CV
set.seed(10)
ffsmodel <- ffs(trainDat[,predictors],trainDat$VW,method="rf",
tuneLength=1,trControl=ctrl)
ffsmodel

#compare to model without ffs:
model <- train(trainDat[,predictors],trainDat$VW,method="rf",
tuneLength=1, trControl=ctrl)
model
stopCluster(cl)

## End(Not run)

```

plot_ffs

Plot results of a Forward feature selection or best subset selection

Description

A plotting function for a forward feature selection result. Each point is the mean performance of a model run. Error bars represent the standard errors from cross validation. Marked points show the best model from each number of variables until a further variable could not improve the results. If `type=="selected"`, the contribution of the selected variables to the model performance is shown.

Usage

```

plot_ffs(
  ffs_model,
  plotType = "all",
  palette = rainbow,
  reverse = FALSE,
  marker = "black",
  size = 1.5,
  lwd = 0.5,
  pch = 21,
  ...
)

```

Arguments

ffs_model	Result of a forward feature selection see ffs
plotType	character. Either "all" or "selected"

<code>palette</code>	A color palette
<code>reverse</code>	Character. Should the palette be reversed?
<code>marker</code>	Character. Color to mark the best models
<code>size</code>	Numeric. Size of the points
<code>lwd</code>	Numeric. Width of the error bars
<code>pch</code>	Numeric. Type of point marking the best models
<code>...</code>	Further arguments for base plot if <code>type="selected"</code>

Author(s)

Marvin Ludwig and Hanna Meyer

See Also

[ffs](#), [bss](#)

Examples

```
## Not run:  
data(iris)  
ffsmodel <- ffs(iris[,1:4],iris$Species)  
plot_ffs(ffsmodel)  
#plot performance of selected variables only:  
plot_ffs(ffsmodel,plotType="selected")  
  
## End(Not run)
```

Index

*** package**

CAST, 9

aoa, 2, 7, 8

bss, 5, 13, 15

calibrate_aoa, 3, 6

CAST, 9

CreateSpacetimeFolds, 6, 9, 13

ffs, 6, 11, 11, 14, 15

plot_bss (plot_ffs), 14

plot_ffs, 14

rollapply, 7

train, 5, 6, 12, 13

trainControl, 6, 11, 13