Package 'mlr3fselect'

August 23, 2026

Title Feature Selection for 'mlr3'

Version 1.7.0

Description Feature selection package of the 'mlr3' ecosystem. It selects the optimal feature set for any 'mlr3' learner. The package works with several optimization algorithms e.g. Random Search, Recursive Feature Elimination, and Genetic Search. Moreover, it can automatically optimize learners and estimate the performance of optimized feature sets with nested resampling.

License LGPL-3

URL https://mlr3fselect.mlr-org.com, https://github.com/mlr-org/mlr3fselect

BugReports https://github.com/mlr-org/mlr3fselect/issues

Depends mlr3 (>= 1.8.0), R (>= 3.1.0)

Imports bbotk (>= 1.9.0), checkmate (>= 2.0.0), cli, data.table, lgr, mlr3misc (>= 0.15.1), paradox (>= 1.0.0), R6, stabm

Suggests e1071, fastVoteR (>= 0.0.3), genalg, mirai, mlr3learners, mlr3pipelines, processx, redux, rpart, rush (>= 1.2.1), testthat (>= 3.0.0)

Config/testthat/edition 3

Config/testthat/parallel false

Encoding UTF-8

Language en-US

NeedsCompilation no

Collate 'ArchiveAsyncFSelect.R' 'ArchiveAsyncFSelectFrozen.R' 'ArchiveBatchFSelect.R' 'AutoFSelector.R' 'CallbackAsyncFSelect.R' 'CallbackBatchFSelect.R' 'ContextAsyncFSelect.R' 'ContextBatchFSelect.R' 'EnsembleFSResult.R' 'FSelectInstanceAsyncSingleCrit.R' 'FSelectInstanceAsyncMultiCrit.R' 'FSelectInstanceBatchSingleCrit.R' 'FSelectInstanceBatchMultiCrit.R' 'mlr_fselectors.R'

'FSelector.R' 'FSelectorAsync.R' 'FSelectorAsyncDesignPoints.R'
'FSelectorAsyncExhaustiveSearch.R'
'FSelectorAsyncFromOptimizerAsync.R'
'FSelectorAsyncRandomSearch.R' 'FSelectorBatch.R'
'FSelectorBatchDesignPoints.R'
'FSelectorBatchExhaustiveSearch.R'
'FSelectorBatchFromOptimizerBatch.R'
'FSelectorBatchGeneticSearch.R' 'FSelectorBatchRFE.R'
'FSelectorBatchRFECV.R' 'FSelectorBatchRandomSearch.R'
'FSelectorBatchSequential.R'
'FSelectorBatchShadowVariableSearch.R' 'ObjectiveFSelect.R'
'ObjectiveFSelectAsync.R' 'ObjectiveFSelectBatch.R'
'assertions.R' 'auto_fselector.R' 'bibentries.R'
'embedded_ensemble_fselect.R' 'ensemble_fselect.R'
'extract_inner_fselect_archives.R'
'extract_inner_fselect_results.R' 'faggregate.R' 'fselect.R'
'fselect_nested.R' 'helper.R' 'mlr_callbacks.R' 'reexports.R'
'sugar.R' 'zzz.R'

Config/roxygen2/version 8.1.0

Author Marc Becker [aut, cre] (ORCID: <https://orcid.org/0000-0002-8115-0400>),
Patrick Schratz [aut] (ORCID: <https://orcid.org/0000-0003-0748-6624>),
Michel Lang [aut] (ORCID: <https://orcid.org/0000-0001-9754-0393>),
Bernd Bischl [aut] (ORCID: <https://orcid.org/0000-0001-6002-6980>),
John Zobolas [aut] (ORCID: <https://orcid.org/0000-0002-3609-8674>)

Maintainer Marc Becker <marcbecker@posteo.de>

Repository CRAN

Date/Publication 2026-08-23 09:00:02 UTC

Contents

mlr3fselect-package 3
ArchiveAsyncFSelect 4
ArchiveAsyncFSelectFrozen 8
ArchiveBatchFSelect 10
assert_async_fselect_callback 14
AutoFSelector 15
auto_fselector 20
CallbackAsyncFSelect 22
CallbackBatchFSelect 23
callback_async_fselect 24
callback_batch_fselect 28
ContextAsyncFSelect 30
ContextBatchFSelect 31
embedded_ensemble_fselect 32
ensemble_fselect 34
ensemble_fs_result 36

extract_inner_fselect_archives . . . 44
extract_inner_fselect_results . . . 45
faggregate . . . 46
fs . . . 47
fselect . . . 48
FSelectInstanceAsyncMultiCrit . . . 51
FSelectInstanceAsyncSingleCrit . . . 54
FSelectInstanceBatchMultiCrit . . . 56
FSelectInstanceBatchSingleCrit . . . 59
FSelector . . . 63
FSelectorAsync . . . 66
FSelectorBatch . . . 67
fselect_nested . . . 69
fsi . . . 71
fsi_async . . . 73
mlr3fselect.async_freeze_archive . . . 76
mlr3fselect.backup . . . 76
mlr3fselect.internal_tuning . . . 77
mlr3fselect.one_se_rule . . . 77
mlr3fselect.svm_rfe . . . 78
mlr_fselectors . . . 79
mlr_fselectors_async_design_points . . . 80
mlr_fselectors_async_exhaustive_search . . . 81
mlr_fselectors_async_random_search . . . 82
mlr_fselectors_design_points . . . 83
mlr_fselectors_exhaustive_search . . . 85
mlr_fselectors_genetic_search . . . 87
mlr_fselectors_random_search . . . 88
mlr_fselectors_rfe . . . 90
mlr_fselectors_rfecv . . . 93
mlr_fselectors_sequential . . . 95
mlr_fselectors_shadow_variable_search . . . 97
ObjectiveFSelect . . . 99
ObjectiveFSelectAsync . . . 101
ObjectiveFSelectBatch . . . 102

Index **104**

`mlr3fselect-package` *mlr3fselect: Feature Selection for 'mlr3'*

Description

Feature selection package of the 'mlr3' ecosystem. It selects the optimal feature set for any 'mlr3' learner. The package works with several optimization algorithms e.g. Random Search, Recursive Feature Elimination, and Genetic Search. Moreover, it can automatically optimize learners and estimate the performance of optimized feature sets with nested resampling.

Author(s)

Maintainer: Marc Becker <marcbecker@posteo.de> (ORCID)

Authors:

- Marc Becker <marcbecker@posteo.de> (ORCID)
- Patrick Schratz <patrick.schratz@gmail.com> (ORCID)
- Michel Lang <michellang@gmail.com> (ORCID)
- Bernd Bischl <bernd_bischl@gmx.net> (ORCID)
- John Zobolas <bblodfon@gmail.com> (ORCID)

See Also

Useful links:

- https://mlr3fselect.mlr-org.com
- https://github.com/mlr-org/mlr3fselect
- Report bugs at https://github.com/mlr-org/mlr3fselect/issues

---

`ArchiveAsyncFSelect` *Rush Data Storage*

Description

The `ArchiveAsyncFSelect` stores all evaluated feature subsets and performance scores in a rush::Rush database.

Details

The ArchiveAsyncFSelect is a connector to a rush::Rush database.

Data Structure

The table (`$data`) has the following columns:

- One column for each feature of the search space (`$search_space`).
- One column for each performance measure (`$codomain`).
- `runtime_learners` (`numeric(1)`)
  Sum of training and predict times logged in learners per mlr3::ResampleResult / evaluation. This does not include potential overhead time.
- `timestamp` (`POSIXct`)
  Time stamp when the evaluation was logged into the archive.

Analysis

For analyzing the feature selection results, it is recommended to pass the ArchiveAsyncFSelect to `as.data.table()`. The returned data table contains the mlr3::ResampleResult for each feature subset evaluation.

S3 Methods

- `as.data.table.ArchiveFSelect(x, unnest = "x_domain", exclude_columns = "uhash", measures = NULL)`
  Returns a tabular view of all evaluated feature subsets.
  ArchiveAsyncFSelect -> `data.table::data.table()`

  - `x` (ArchiveAsyncFSelect)
  - `unnest` (`character()`)
    Transforms list columns to separate columns. Set to `NULL` if no column should be unnested.
  - `exclude_columns` (`character()`)
    Exclude columns from table. Set to `NULL` if no column should be excluded.
  - `measures` (List of mlr3::Measure)
    Score feature subsets on additional measures.

Super classes

`bbotk::Archive` -> `bbotk::ArchiveAsync` -> `ArchiveAsyncFSelect`

Active bindings

`benchmark_result` (mlr3::BenchmarkResult)
Benchmark result.

`ties_method` (`character(1)`)
Method to handle ties in the archive. One of `"least_features"` (default) or `"random"`.

Methods

Public methods:

- `ArchiveAsyncFSelect$new()`
- `ArchiveAsyncFSelect$learner()`
- `ArchiveAsyncFSelect$learners()`
- `ArchiveAsyncFSelect$predictions()`
- `ArchiveAsyncFSelect$resample_result()`
- `ArchiveAsyncFSelect$print()`
- `ArchiveAsyncFSelect$best()`
- `ArchiveAsyncFSelect$push_result()`
- `ArchiveAsyncFSelect$clone()`

`ArchiveAsyncFSelect$new()`: Creates a new instance of this R6 class.

Usage:

```
ArchiveAsyncFSelect$new(
  search_space,
  codomain,
  rush,
  ties_method = "least_features"
)
```

Arguments:

`search_space` (paradox::ParamSet)
: Search space. Internally created from provided mlr3::Task by instance.

`codomain` (paradox::ParamSet)
: Specifies codomain of function. Most importantly the tags of each output "Parameter" define whether it should be minimized or maximized. The default is to minimize each component.

`rush` (Rush)
: If a rush instance is supplied, the optimization runs without batches.

`ties_method` (`character(1)`)
: The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`ArchiveAsyncFSelect$learner()`: Retrieve mlr3::Learner of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive. Learner does not contain a model. Use `$learners()` to get learners with models.

Usage:

```
ArchiveAsyncFSelect$learner(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
: The iteration value to filter for.

`uhash` (`logical(1)`)
: The `uhash` value to filter for.

`ArchiveAsyncFSelect$learners()`: Retrieve list of trained mlr3::Learner objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelect$learners(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
: The iteration value to filter for.

`uhash` (`logical(1)`)
: The `uhash` value to filter for.

`ArchiveAsyncFSelect$predictions()`: Retrieve list of mlr3::Prediction objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelect$predictions(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The uhash value to filter for.

`ArchiveAsyncFSelect$resample_result()`: Retrieve mlr3::ResampleResult of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelect$resample_result(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The uhash value to filter for.

`ArchiveAsyncFSelect$print()`: Printer.

Usage:

```
ArchiveAsyncFSelect$print(...)
```

Arguments:

`...` (ignored).

`ArchiveAsyncFSelect$best()`: Returns the best scoring feature set(s). For single-crit optimization, the solution that minimizes / maximizes the objective function. For multi-crit optimization, the Pareto set / front.

Usage:

```
ArchiveAsyncFSelect$best(n_select = 1, ties_method = NULL)
```

Arguments:

`n_select (integer(1L))`
Amount of points to select. Ignored for multi-crit optimization.

`ties_method (character(1L))`
Method to break ties when multiple points have the same score. Either `"least_features"` or `"random"`. If `NULL` (default), the global ties method set during initialization is used. Ignored for multi-crit optimization. If `n_select > 1L`, the tie method is ignored and the first point is returned.

Returns: `data.table::data.table()`

`ArchiveAsyncFSelect$push_result()`: Push result to the archive.

Usage:

```
ArchiveAsyncFSelect$push_result(key, ys, x_domain, extra = NULL)
```

Arguments:

`key` (`character()`)
Key of the point.

`ys` (`list()`)
Named list of results.

`x_domain` (`list()`)
Is ignored for feature selection.

`extra` (`list()`)
Named list of additional information.

`ArchiveAsyncFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ArchiveAsyncFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`ArchiveAsyncFSelectFrozen`
Frozen Rush Data Storage

Description

Freezes the Redis data base of an ArchiveAsyncFSelect to a `data.table::data.table()`. No further points can be added to the archive but the data can be accessed and analyzed. Useful when the Redis data base is not permanently available. Use the callback mlr3fselect.async_freeze_archive to freeze the archive after the optimization has finished.

S3 Methods

- `as.data.table(archive)`
  ArchiveAsyncFSelectFrozen -> `data.table::data.table()`
  Returns a tabular view of all performed function calls of the Objective.

Super classes

`bbotk::Archive` -> `bbotk::ArchiveAsync` -> `bbotk::ArchiveAsyncFrozen` -> `ArchiveAsyncFSelectFrozen`

Active bindings

`benchmark_result` (mlr3::BenchmarkResult)
Benchmark result.

Methods

Public methods:

- ArchiveAsyncFSelectFrozen$new()
- ArchiveAsyncFSelectFrozen$learner()
- ArchiveAsyncFSelectFrozen$learners()
- ArchiveAsyncFSelectFrozen$predictions()
- ArchiveAsyncFSelectFrozen$resample_result()
- ArchiveAsyncFSelectFrozen$print()
- ArchiveAsyncFSelectFrozen$clone()

`ArchiveAsyncFSelectFrozen$new()`: Creates a new instance of this R6 class.

Usage:

```
ArchiveAsyncFSelectFrozen$new(archive)
```

Arguments:

`archive` (ArchiveAsyncFSelect)
The archive to freeze.

`ArchiveAsyncFSelectFrozen$learner()`: Retrieve mlr3::Learner of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive. Learner does not contain a model. Use `$learners()` to get learners with models.

Usage:

```
ArchiveAsyncFSelectFrozen$learner(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
The iteration value to filter for.

`uhash` (`logical(1)`)
The `uhash` value to filter for.

`ArchiveAsyncFSelectFrozen$learners()`: Retrieve list of trained mlr3::Learner objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelectFrozen$learners(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
The iteration value to filter for.

`uhash` (`logical(1)`)
The `uhash` value to filter for.

`ArchiveAsyncFSelectFrozen$predictions()`: Retrieve list of mlr3::Prediction objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelectFrozen$predictions(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
The iteration value to filter for.

`uhash` (`logical(1)`)
The uhash value to filter for.

`ArchiveAsyncFSelectFrozen$resample_result()`: Retrieve mlr3::ResampleResult of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveAsyncFSelectFrozen$resample_result(i = NULL, uhash = NULL)
```

Arguments:

`i` (`integer(1)`)
The iteration value to filter for.

`uhash` (`logical(1)`)
The uhash value to filter for.

`ArchiveAsyncFSelectFrozen$print()`: Printer.

Usage:

```
ArchiveAsyncFSelectFrozen$print(...)
```

Arguments:

`...` (ignored).

`ArchiveAsyncFSelectFrozen$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ArchiveAsyncFSelectFrozen$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`ArchiveBatchFSelect` *Class for Logging Evaluated Feature Sets*

Description

The ArchiveBatchFSelect stores all evaluated feature sets and performance scores.

Details

The ArchiveBatchFSelect is a container around a `data.table::data.table()`. Each row corresponds to a single evaluation of a feature set. See the section on Data Structure for more information. The archive stores additionally a mlr3::BenchmarkResult (`$benchmark_result`) that records the resampling experiments. Each experiment corresponds to a single evaluation of a feature set. The table (`$data`) and the benchmark result (`$benchmark_result`) are linked by the `uhash` column. If the archive is passed to `as.data.table()`, both are joined automatically.

Data structure

The table (`$data`) has the following columns:

- One column for each feature of the task (`$search_space`).
- One column for each performance measure (`$codomain`).
- `runtime_learners` (`numeric(1)`)
  Sum of training and predict times logged in learners per mlr3::ResampleResult / evaluation. This does not include potential overhead time.
- `timestamp` (`POSIXct`)
  Time stamp when the evaluation was logged into the archive.
- `batch_nr` (`integer(1)`)
  Feature sets are evaluated in batches. Each batch has a unique batch number.
- `uhash` (`character(1)`)
  Connects each feature set to the resampling experiment stored in the mlr3::BenchmarkResult.

Analysis

For analyzing the feature selection results, it is recommended to pass the archive to `as.data.table()`. The returned data table is joined with the benchmark result which adds the mlr3::ResampleResult for each feature set.

The archive provides various getters (e.g. `$learners()`) to ease the access. All getters extract by position (`i`) or unique hash (`uhash`). For a complete list of all getters see the methods section.

The benchmark result (`$benchmark_result`) allows to score the feature sets again on a different measure. Alternatively, measures can be supplied to `as.data.table()`.

S3 Methods

- `as.data.table.ArchiveBatchFSelect(x, exclude_columns = "uhash", measures = NULL)`
  Returns a tabular view of all evaluated feature sets.
  ArchiveBatchFSelect -> `data.table::data.table()`

  - `x` (ArchiveBatchFSelect)
  - `exclude_columns` (`character()`)
    Exclude columns from table. Set to `NULL` if no column should be excluded.
  - `measures` (list of mlr3::Measure)
    Score feature sets on additional measures.

Super classes

`bbotk::Archive` -> `bbotk::ArchiveBatch` -> `ArchiveBatchFSelect`

Public fields

`benchmark_result` (mlr3::BenchmarkResult)
Benchmark result.

Active bindings

`ties_method (character(1))`
Method to handle ties.

Methods

Public methods:

- `ArchiveBatchFSelect$new()`
- `ArchiveBatchFSelect$add_evals()`
- `ArchiveBatchFSelect$learner()`
- `ArchiveBatchFSelect$learners()`
- `ArchiveBatchFSelect$predictions()`
- `ArchiveBatchFSelect$resample_result()`
- `ArchiveBatchFSelect$print()`
- `ArchiveBatchFSelect$best()`
- `ArchiveBatchFSelect$clone()`

`ArchiveBatchFSelect$new()`: Creates a new instance of this R6 class.

Usage:

```
ArchiveBatchFSelect$new(
  search_space,
  codomain,
  check_values = TRUE,
  ties_method = "least_features"
)
```

Arguments:

`search_space` (paradox::ParamSet)
Search space. Internally created from provided mlr3::Task by instance.

`codomain` (bbotk::Codomain)
Specifies codomain of objective function i.e. a set of performance measures. Internally created from provided mlr3::Measures by instance.

`check_values (logical(1))`
If `TRUE` (default), hyperparameter configurations are check for validity.

`ties_method (character(1))`
The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`ArchiveBatchFSelect$add_evals()`: Adds function evaluations to the archive table.

Usage:

```
ArchiveBatchFSelect$add_evals(xdt, xss_trafoed = NULL, ydt)
```

Arguments:

`xdt (data.table::data.table())`
x values as `data.table`. Each row is one point. Contains the value in the *search space* of the FSelectInstanceBatchMultiCrit object. Can contain additional columns for extra information.

`xss_trafoed (list())`
Ignored in feature selection.

`ydt (data.table::data.table())`
Optimal outcome.

`ArchiveBatchFSelect$learner()`: Retrieve mlr3::Learner of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive. Learner does not contain a model. Use `$learners()` to get learners with models.

Usage:

```
ArchiveBatchFSelect$learner(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The `uhash` value to filter for.

`ArchiveBatchFSelect$learners()`: Retrieve list of trained mlr3::Learner objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveBatchFSelect$learners(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The `uhash` value to filter for.

`ArchiveBatchFSelect$predictions()`: Retrieve list of mlr3::Prediction objects of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveBatchFSelect$predictions(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The `uhash` value to filter for.

`ArchiveBatchFSelect$resample_result()`: Retrieve mlr3::ResampleResult of the i-th evaluation, by position or by unique hash `uhash`. `i` and `uhash` are mutually exclusive.

Usage:

```
ArchiveBatchFSelect$resample_result(i = NULL, uhash = NULL)
```

Arguments:

`i (integer(1))`
The iteration value to filter for.

`uhash (logical(1))`
The uhash value to filter for.

`ArchiveBatchFSelect$print()`: Printer.

Usage:

```
ArchiveBatchFSelect$print(...)
```

Arguments:

`...` (ignored).

`ArchiveBatchFSelect$best()`: Returns the best scoring feature sets.

Usage:

```
ArchiveBatchFSelect$best(batch = NULL, ties_method = NULL)
```

Arguments:

`batch (integer())`
The batch number(s) to limit the best results to. Default is all batches.

`ties_method (character(1))`
Method to handle ties. If `NULL` (default), the global ties method set during initialization is used. The default global ties method is `least_features` which selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets.

Returns: `data.table::data.table()`

`ArchiveBatchFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ArchiveBatchFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`assert_async_fselect_callback` *Assertions for Callbacks*

Description

Assertions for CallbackAsyncFSelect class.

Usage

```
assert_async_fselect_callback(callback, null_ok = FALSE)

assert_async_fselect_callbacks(callbacks)
```

Arguments

callback	(CallbackAsyncFSelect).
null_ok	(logical(1)) If TRUE, NULL is allowed.
callbacks	(list of CallbackAsyncFSelect).

Value

[CallbackAsyncFSelect | List of CallbackAsyncFSelects.

AutoFSelector *Class for Automatic Feature Selection*

Description

The AutoFSelector wraps a mlr3::Learner and augments it with an automatic feature selection. The auto_fselector() function creates an AutoFSelector object.

Details

The AutoFSelector is a mlr3::Learner which wraps another mlr3::Learner and performs the following steps during $train():

1. The wrapped (inner) learner is trained on the feature subsets via resampling. The feature selection can be specified by providing a FSelector, a bbotk::Terminator, a mlr3::Resampling, and a mlr3::Measure.
2. A final model is fit on the complete training data with the best-found feature subset.

During $predict() the AutoFSelector just calls the predict method of the wrapped (inner) learner.

Resources

There are several sections about feature selection in the mlr3book.

- Estimate Model Performance with nested resampling.

The gallery features a collection of case studies and demos about optimization.

Nested Resampling

Nested resampling can be performed by passing an AutoFSelector object to mlr3::resample() or mlr3::benchmark(). To access the inner resampling results, set store_fselect_instance = TRUE and execute mlr3::resample() or mlr3::benchmark() with store_models = TRUE (see examples). The mlr3::Resampling passed to the AutoFSelector is meant to be the inner resampling, operating on the training set of an arbitrary outer resampling. For this reason it is not feasible to pass an instantiated mlr3::Resampling here.

Super class

`mlr3::Learner` -> `AutoFSelector`

Public fields

`instance_args` (`list()`)
: All arguments from construction to create the FSelectInstanceBatchSingleCrit.

`fselector` (FSelector)
: Optimization algorithm.

Active bindings

`archive` ([ArchiveBatchFSelect)
: Returns FSelectInstanceBatchSingleCrit archive.

`learner` (mlr3::Learner)
: Trained learner.

`fselect_instance` (FSelectInstanceBatchSingleCrit)
: Internally created feature selection instance with all intermediate results.

`fselect_result` (data.table::data.table)
: Short-cut to `$result` from FSelectInstanceBatchSingleCrit.

`predict_type` (`character(1)`)
: Stores the currently active predict type, e.g. `"response"`. Must be an element of `$predict_types`.

`hash` (`character(1)`)
: Hash (unique identifier) for this object. Covers the id, the parameter values, the predict type, the fallback learner, the parallel predict flag, the FSelector, the arguments of the FSelectInstanceBatchSingleCrit and the store fselect instance flag.

`phash` (`character(1)`)
: Hash (unique identifier) for this partial object. The AutoFSelector has no components that are varied systematically during tuning or feature selection, because the search space is created internally from the task. The partial hash is therefore deliberately identical to `$hash`.

Methods

Public methods:

- `AutoFSelector$new()`
- `AutoFSelector$base_learner()`
- `AutoFSelector$importance()`
- `AutoFSelector$selected_features()`
- `AutoFSelector$oob_error()`
- `AutoFSelector$loglik()`
- `AutoFSelector$print()`
- `AutoFSelector$clone()`

`AutoFSelector$new()`: Creates a new instance of this R6 class.

Usage:

```
AutoFSelector$new(
  fselector,
  learner,
  resampling,
  measure = NULL,
  terminator,
  store_fselect_instance = TRUE,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features",
  rush = NULL,
  id = NULL
)
```

Arguments:

`fselector` (FSelector)
: Optimization algorithm.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measure` (mlr3::Measure)
: Measure to optimize. If `NULL`, default measure is used.

`terminator` (bbotk::Terminator)
: Stop criterion of the feature selection.

`store_fselect_instance (logical(1))`
: If `TRUE` (default), stores the internally created FSelectInstanceBatchSingleCrit with all intermediate results in slot `$fselect_instance`. Is set to `TRUE`, if `store_models = TRUE`

`store_benchmark_result (logical(1))`
: Store benchmark result in archive?

`store_models (logical(1))`. Store models in benchmark result?

`check_values (logical(1))`
: Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
: List of callbacks.

`ties_method (character(1))`
: The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`rush (Rush)`
: If a rush instance is supplied, the optimization runs without batches.

`id (character(1))`
Identifier for the new instance.

`AutoFSelector$base_learner()`: Extracts the base learner from nested learner objects like `GraphLearner` in **mlr3pipelines**. If `recursive = 0`, the (tuned) learner is returned.

Usage:

```
AutoFSelector$base_learner(recursive = Inf)
```

Arguments:

`recursive (integer(1))`
Depth of recursion for multiple nested objects.

Returns: mlr3::Learner.

`AutoFSelector$importance()`: The importance scores of the final model.

Usage:

```
AutoFSelector$importance()
```

Returns: Named `numeric()`.

`AutoFSelector$selected_features()`: The selected features of the final model. These features are selected internally by the learner.

Usage:

```
AutoFSelector$selected_features()
```

Returns: `character()`.

`AutoFSelector$oob_error()`: The out-of-bag error of the final model.

Usage:

```
AutoFSelector$oob_error()
```

Returns: `numeric(1)`.

`AutoFSelector$loglik()`: The log-likelihood of the final model.

Usage:

```
AutoFSelector$loglik()
```

Returns: `logLik`.

`AutoFSelector$print()`: Printer.

Usage:

```
AutoFSelector$print(...)
```

Arguments:

`...` (ignored).

`AutoFSelector$clone()`: The objects of this class are cloneable with this method.

Usage:

```
AutoFSelector$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Examples

```
# Automatic Feature Selection


# split to train and external set
task = tsk("penguins")
split = partition(task, ratio = 0.8)

# create auto fselector
afs = auto_fselector(
  fselector = fs("random_search"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("holdout"),
  measure = msr("classif.ce"),
  term_evals = 4)

# optimize feature subset and fit final model
afs$train(task, row_ids = split$train)

# predict with final model
afs$predict(task, row_ids = split$test)

# show result
afs$fselect_result

# model slot contains trained learner and fselect instance
afs$model

# shortcut trained learner
afs$learner

# shortcut fselect instance
afs$fselect_instance


# Nested Resampling

afs = auto_fselector(
  fselector = fs("random_search"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("holdout"),
  measure = msr("classif.ce"),
  term_evals = 4)

resampling_outer = rsmp("cv", folds = 3)
rr = resample(task, afs, resampling_outer, store_models = TRUE)

# retrieve inner feature selection results.
extract_inner_fselect_results(rr)

# performance scores estimated on the outer resampling
rr$score()
```

```
# unbiased performance of the final model trained on the full data set
rr$aggregate()
```

`auto_fselector` *Function for Automatic Feature Selection*

Description

The AutoFSelector wraps a mlr3::Learner and augments it with an automatic feature selection. The `auto_fselector()` function creates an AutoFSelector object.

Usage

```
auto_fselector(
  fselector,
  learner,
  resampling,
  measure = NULL,
  term_evals = NULL,
  term_time = NULL,
  terminator = NULL,
  store_fselect_instance = TRUE,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features",
  rush = NULL,
  id = NULL
)
```

Arguments

`fselector`	(FSelector) Optimization algorithm.
`learner`	(mlr3::Learner) Learner to optimize the feature subset for.
`resampling`	(mlr3::Resampling) Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.
`measure`	(mlr3::Measure) Measure to optimize. If `NULL`, default measure is used.

`term_evals` (integer(1))
Number of allowed evaluations. Ignored if `terminator` is passed.

`term_time` (integer(1))
Maximum allowed time in seconds. Ignored if `terminator` is passed.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_fselect_instance`
(logical(1))
If `TRUE` (default), stores the internally created FSelectInstanceBatchSingleCrit with all intermediate results in slot `$fselect_instance`. Is set to `TRUE`, if `store_models = TRUE`

`store_benchmark_result`
(logical(1))
Store benchmark result in archive?

`store_models` (logical(1)). Store models in benchmark result?

`check_values` (logical(1))
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ties_method` (character(1))
The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`rush` (Rush)
If a rush instance is supplied, the optimization runs without batches.

`id` (character(1))
Identifier for the new instance.

Details

The AutoFSelector is a mlr3::Learner which wraps another mlr3::Learner and performs the following steps during `$train()`:

1. The wrapped (inner) learner is trained on the feature subsets via resampling. The feature selection can be specified by providing a FSelector, a bbotk::Terminator, a mlr3::Resampling, and a mlr3::Measure.
2. A final model is fit on the complete training data with the best-found feature subset.

During `$predict()` the AutoFSelector just calls the predict method of the wrapped (inner) learner.

Value

AutoFSelector.

Resources

There are several sections about feature selection in the mlr3book.

- Estimate Model Performance with nested resampling.

The gallery features a collection of case studies and demos about optimization.

Nested Resampling

Nested resampling can be performed by passing an AutoFSelector object to `mlr3::resample()` or `mlr3::benchmark()`. To access the inner resampling results, set `store_fselect_instance = TRUE` and execute `mlr3::resample()` or `mlr3::benchmark()` with `store_models = TRUE` (see examples). The mlr3::Resampling passed to the AutoFSelector is meant to be the inner resampling, operating on the training set of an arbitrary outer resampling. For this reason it is not feasible to pass an instantiated mlr3::Resampling here.

Examples

```
afs = auto_fselector(
  fselector = fs("random_search"),
  learner = lrn("classif.rpart"),
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  term_evals = 4)

afs$train(tsk("diabetes"))
```

CallbackAsyncFSelect *Asynchronous Feature Selection Callback*

Description

Specialized bbotk::CallbackAsync for asynchronous feature selection. Callbacks allow to customize the behavior of processes in mlr3fselect. The `callback_async_fselect()` function creates a CallbackAsyncFSelect. Predefined callbacks are stored in the dictionary mlr_callbacks and can be retrieved with `clbk()`. For more information on feature selection callbacks see `callback_async_fselect()`.

Super classes

`mlr3misc::Callback` -> `bbotk::CallbackAsync` -> `CallbackAsyncFSelect`

Public fields

`on_eval_after_xs (function())`
: Stage called after xs is passed. Called in `ObjectiveFSelectAsync$eval()`.

`on_resample_begin (function())`
: Stage called at the beginning of an evaluation. Called in `workhorse()` (internal).

`on_resample_before_train (function())`
: Stage called before training the learner. Called in `workhorse()` (internal).

`on_resample_before_predict (function())`
: Stage called before predicting. Called in `workhorse()` (internal).

`on_resample_end (function())`
: Stage called at the end of an evaluation. Called in `workhorse()` (internal).

`on_eval_after_resample (function())`
: Stage called after feature subsets are evaluated. Called in `ObjectiveFSelectAsync$eval()`.

`on_eval_before_archive (function())`
: Stage called before performance values are written to the archive. Called in `ObjectiveFSelectAsync$eval()`.

`on_fselect_result_begin (function())`
: Stage called before the results are written. Called in `FSelectInstance*$assign_result()`.

Methods

Public methods:

- `CallbackAsyncFSelect$clone()`

`CallbackAsyncFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
CallbackAsyncFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`CallbackBatchFSelect` *Create Feature Selection Callback*

Description

Specialized bbotk::CallbackBatch for feature selection. Callbacks allow customizing the behavior of processes in mlr3fselect. The `callback_batch_fselect()` function creates a CallbackBatchFSelect. Predefined callbacks are stored in the dictionary mlr_callbacks and can be retrieved with `clbk()`. For more information on callbacks see `callback_batch_fselect()`.

Super classes

`mlr3misc::Callback` -> `bbotk::CallbackBatch` -> `CallbackBatchFSelect`

Public fields

`on_eval_after_design (function())`
: Stage called after design is created. Called in `ObjectiveFSelectBatch$eval_many()`.

`on_eval_after_benchmark (function())`
: Stage called after feature sets are evaluated. Called in `ObjectiveFSelectBatch$eval_many()`.

`on_eval_before_archive (function())`
: Stage called before performance values are written to the archive. Called in `ObjectiveFSelectBatch$eval_many()`.

`on_auto_fselector_before_final_model (function())`
: Stage called before the final model is trained. Called in `AutoFSelector$train()`. This stage is called after the optimization has finished and the final model is trained with the best feature set found.

`on_auto_fselector_after_final_model (function())`
: Stage called after the final model is trained. Called in `AutoFSelector$train()`. This stage is called after the final model is trained with the best feature set found.

Methods

Public methods:

- `CallbackBatchFSelect$clone()`

`CallbackBatchFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
CallbackBatchFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Examples

```
# Write archive to disk
callback_batch_fselect("mlr3fselect.backup",
  on_optimization_end = function(callback, context) {
    saveRDS(context$instance$archive, "archive.rds")
  }
)
```

`callback_async_fselect` *Create Asynchronous Feature Selection Callback*

Description

Function to create a CallbackAsyncFSelect. Predefined callbacks are stored in the dictionary mlr_callbacks and can be retrieved with `clbk()`.

Feature selection callbacks can be called from different stages of the feature selection process. The stages are prefixed with `on_*`.

```
Start Feature Selection
     - on_optimization_begin
    Start Worker
         - on_worker_begin
         Start Optimization on Worker
           - on_optimizer_before_eval
             Start Evaluation
               - on_eval_after_xs
                 Start Resampling Iteration
                   - on_resample_begin
                   - on_resample_before_train
                   - on_resample_before_predict
                   - on_resample_end
                 End Resampling Iteration
               - on_eval_after_resample
               - on_eval_before_archive
             End Evaluation
          - on_optimizer_after_eval
         End Optimization on Worker
         - on_worker_end
    End Worker
     - on_fselect_result_begin
     - on_result_begin
     - on_result_end
     - on_optimization_end
End Feature Selection
```

See also the section on parameters for more information on the stages. A feature selection callback works with ContextAsyncFSelect.

Usage

```
callback_async_fselect(
  id,
  label = NA_character_,
  man = NA_character_,
  on_optimization_begin = NULL,
  on_worker_begin = NULL,
  on_optimizer_before_eval = NULL,
  on_eval_after_xs = NULL,
  on_resample_begin = NULL,
  on_resample_before_train = NULL,
```

```
  on_resample_before_predict = NULL,
  on_resample_end = NULL,
  on_eval_after_resample = NULL,
  on_eval_before_archive = NULL,
  on_optimizer_after_eval = NULL,
  on_worker_end = NULL,
  on_fselect_result_begin = NULL,
  on_result_begin = NULL,
  on_result_end = NULL,
  on_result = NULL,
  on_optimization_end = NULL
)
```

Arguments

`id` (`character(1)`)
Identifier for the new instance.

`label` (`character(1)`)
Label for the new instance.

`man` (`character(1)`)
String in the format `[pkg]::[topic]` pointing to a manual page for this object. The referenced help package can be opened via method `$help()`.

`on_optimization_begin` (`function()`)
Stage called at the beginning of the optimization. Called in `Optimizer$optimize()`. The functions must have two arguments named `callback` and `context`.

`on_worker_begin` (`function()`)
Stage called at the beginning of the optimization on the worker. Called in the worker loop. The functions must have two arguments named `callback` and `context`.

`on_optimizer_before_eval` (`function()`)
Stage called after the optimizer proposes points. Called in `OptimInstance$.eval_point()`. The functions must have two arguments named `callback` and `context`. The argument of `instance$.eval_point(xs)` and `xs_trafoed` and `extra` are available in the `context`. Or `xs` and `xs_trafoed` of `instance$.eval_queue()` are available in the `context`.

`on_eval_after_xs` (`function()`)
Stage called after xs is passed to the objective. Called in `ObjectiveFSelectAsync$eval()`. The functions must have two arguments named `callback` and `context`. The argument of `$.eval(xs)` is available in the `context`.

`on_resample_begin` (`function()`)
Stage called at the beginning of a resampling iteration. Called in `workhorse()` (internal). See also `mlr3::callback_resample()`. The functions must have two arguments named `callback` and `context`.

on_resample_before_train
: (function())
Stage called before training the learner. Called in workhorse() (internal). See also mlr3::callback_resample(). The functions must have two arguments named callback and context.

on_resample_before_predict
: (function())
Stage called before predicting. Called in workhorse() (internal). See also mlr3::callback_resample(). The functions must have two arguments named callback and context.

on_resample_end
: (function())
Stage called at the end of a resampling iteration. Called in workhorse() (internal). See also mlr3::callback_resample(). The functions must have two arguments named callback and context.

on_eval_after_resample
: (function())
Stage called after a feature subset is evaluated. Called in ObjectiveFSelectAsync$eval(). The functions must have two arguments named callback and context. The resample_result is available in the context.

on_eval_before_archive
: (function())
Stage called before performance values are written to the archive. Called in ObjectiveFSelectAsync$eval(). The functions must have two arguments named callback and context. The aggregated_performance is available in context.

on_optimizer_after_eval
: (function())
Stage called after points are evaluated. Called in OptimInstance$.eval_point(). The functions must have two arguments named callback and context.

on_worker_end
: (function())
Stage called at the end of the optimization on the worker. Called in the worker loop. The functions must have two arguments named callback and context.

on_fselect_result_begin
: (function())
Stage called at the beginning of the result writing. Called in FSelectInstance*$assign_result(). The functions must have two arguments named callback and context. The arguments of $assign_result(xdt, y, extra) are available in context.

on_result_begin
: (function())
Stage called at the beginning of the result writing. Called in OptimInstance$assign_result(). The functions must have two arguments named callback and context. The arguments of $.assign_result(xdt, y, extra) are available in the context.

on_result_end
: (function())
Stage called after the result is written. Called in OptimInstance$assign_result(). The functions must have two arguments named callback and context. The final result instance$result is available in the context.

`on_result` (`function()`)
Deprecated. Use `on_result_end` instead. Stage called after the result is written. Called in `OptimInstance$assign_result()`.

`on_optimization_end` (`function()`)
Stage called at the end of the optimization. Called in `Optimizer$optimize()`.

Details

When implementing a callback, each function must have two arguments named `callback` and `context`. A callback can write data to the state (`$state`), e.g. settings that affect the callback itself. Feature selection callbacks access ContextAsyncFSelect and mlr3::ContextResample.

`callback_batch_fselect`

Create Feature Selection Callback

Description

Function to create a CallbackBatchFSelect. Predefined callbacks are stored in the dictionary mlr_callbacks and can be retrieved with `clbk()`.

Feature selection callbacks can be called from different stages of feature selection. The stages are prefixed with `on_*`. The `on_auto_fselector_*` stages are only available when the callback is used in an AutoFSelector.

```
Start Automatic Feature Selection
  Start Feature Selection
      - on_optimization_begin
      Start FSelect Batch
          - on_optimizer_before_eval
          Start Evaluation
              - on_eval_after_design
              - on_eval_after_benchmark
              - on_eval_before_archive
          End Evaluation
          - on_optimizer_after_eval
      End FSelect Batch
      - on_result
      - on_optimization_end
  End Feature Selection
  - on_auto_fselector_before_final_model
  - on_auto_fselector_after_final_model
End Automatic Feature Selection
```

See also the section on parameters for more information on the stages. A feature selection callback works with bbotk::ContextBatch and ContextBatchFSelect.

Usage

```
callback_batch_fselect(
  id,
  label = NA_character_,
  man = NA_character_,
  on_optimization_begin = NULL,
  on_optimizer_before_eval = NULL,
  on_eval_after_design = NULL,
  on_eval_after_benchmark = NULL,
  on_eval_before_archive = NULL,
  on_optimizer_after_eval = NULL,
  on_result = NULL,
  on_optimization_end = NULL,
  on_auto_fselector_before_final_model = NULL,
  on_auto_fselector_after_final_model = NULL
)
```

Arguments

id	(character(1)) Identifier for the new instance.
label	(character(1)) Label for the new instance.
man	(character(1)) String in the format [pkg]::[topic] pointing to a manual page for this object. The referenced help package can be opened via method $help().
on_optimization_begin	(function()) Stage called at the beginning of the optimization. Called in Optimizer$optimize().
on_optimizer_before_eval	(function()) Stage called after the optimizer proposes points. Called in OptimInstance$eval_batch().
on_eval_after_design	(function()) Stage called after design is created. Called in ObjectiveFSelectBatch$eval_many().
on_eval_after_benchmark	(function()) Stage called after feature sets are evaluated. Called in ObjectiveFSelectBatch$eval_many().
on_eval_before_archive	(function()) Stage called before performance values are written to the archive. Called in ObjectiveFSelectBatch$eval_many().
on_optimizer_after_eval	(function()) Stage called after points are evaluated. Called in OptimInstance$eval_batch().
on_result	(function()) Stage called after result are written. Called in OptimInstance$assign_result().

on_optimization_end
: (function())
Stage called at the end of the optimization. Called in Optimizer$optimize().

on_auto_fselector_before_final_model
: (function())
Stage called before the final model is trained. Called in AutoFSelector$train().

on_auto_fselector_after_final_model
: (function())
Stage called after the final model is trained. Called in AutoFSelector$train().

Details

When implementing a callback, each function must have two arguments named callback and context. A callback can write data to the state ($state), e.g. settings that affect the callback itself. Avoid writing large data the state.

Examples

```
# Write archive to disk
callback_batch_fselect("mlr3fselect.backup",
  on_optimization_end = function(callback, context) {
    saveRDS(context$instance$archive, "archive.rds")
  }
)
```

ContextAsyncFSelect *Asynchronous Feature Selection Context*

Description

A CallbackAsyncFSelect accesses and modifies data during the optimization via the ContextAsyncFSelect. See the section on active bindings for a list of modifiable objects. See callback_async_fselect() for a list of stages that access ContextAsyncFSelect.

Details

Changes to $instance and $optimizer in the stages executed on the workers are not reflected in the main process.

Super classes

mlr3misc::Context -> bbotk::ContextAsync -> ContextAsyncFSelect

Public fields

auto_fselector (AutoFSelector)
: The AutoFSelector instance.

Active bindings

`xs_objective (list())`
: The feature subset currently evaluated.

`resample_result` (mlr3::BenchmarkResult)
: The resample result of the feature subset currently evaluated.

`aggregated_performance (list())`
: Aggregated performance scores and training time of the evaluated feature subset. This list is passed to the archive. A callback can add additional elements which are also written to the archive.

`result_feature_set` (character())
: The feature set passed to `instance$assign_result()`.

Methods

Public methods:

- `ContextAsyncFSelect$clone()`

`ContextAsyncFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ContextAsyncFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

ContextBatchFSelect *Evaluation Context*

Description

The ContextBatchFSelect allows CallbackBatchFSelects to access and modify data while a batch of feature sets is evaluated. See the section on active bindings for a list of modifiable objects. See `callback_batch_fselect()` for a list of stages that access ContextBatchFSelect.

Details

This context is re-created each time a new batch of feature sets is evaluated. Changes to `$objective_fselect`, `$design` `$benchmark_result` are discarded after the function is finished. Modification on the data table in `$aggregated_performance` are written to the archive. Any number of columns can be added.

Super classes

`mlr3misc::Context` -> `bbotk::ContextBatch` -> `ContextBatchFSelect`

Public fields

`auto_fselector` (AutoFSelector)
: The AutoFSelector instance.

Active bindings

`xss` (list())
: The feature sets of the latest batch.

`design` (data.table::data.table)
: The benchmark design of the latest batch.

`benchmark_result` (mlr3::BenchmarkResult)
: The benchmark result of the latest batch.

`aggregated_performance` (data.table::data.table)
: Aggregated performance scores and training time of the latest batch. This data table is passed to the archive. A callback can add additional columns which are also written to the archive.

Methods

Public methods:

- `ContextBatchFSelect$clone()`

`ContextBatchFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ContextBatchFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`embedded_ensemble_fselect` *Embedded Ensemble Feature Selection*

Description

Ensemble feature selection using multiple learners. The ensemble feature selection method is designed to identify the most predictive features from a given dataset by leveraging multiple machine learning models and resampling techniques. Returns an EnsembleFSResult.

Usage

```
embedded_ensemble_fselect(
  task,
  learners,
  init_resampling,
  measure,
  store_benchmark_result = TRUE
)
```

Arguments

`task` (mlr3::Task)
Task to operate on.

`learners` (list of mlr3::Learner)
The learners to be used for feature selection. All learners must have the `selected_features` property, i.e. implement embedded feature selection (e.g. regularized models).

`init_resampling`
(mlr3::Resampling)
The initial resampling strategy of the data, from which each train set will be passed on to the learners and each test set will be used for prediction. Can only be mlr3::ResamplingSubsampling or mlr3::ResamplingBootstrap.

`measure` (mlr3::Measure)
The measure used to score each learner on the test sets generated by `init_resampling`.

`store_benchmark_result`
(`logical(1)`)
Whether to store the benchmark result in EnsembleFSResult or not.

Details

The method begins by applying an initial resampling technique specified by the user, to create **multiple subsamples** from the original dataset (train/test splits). This resampling process helps in generating diverse subsets of data for robust feature selection.

For each subsample (train set) generated in the previous step, the method applies learners that support **embedded feature selection**. These learners are then scored on their ability to predict on the resampled test sets, storing the selected features during training, for each combination of subsample and learner.

Results are stored in an EnsembleFSResult.

Value

an EnsembleFSResult object.

Source

Meinshausen, Nicolai, Buhlmann, Peter (2010). "Stability Selection." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **72**(4), 417–473. ISSN 1369-7412. doi:10.1111/J.14679868.2010.00740.X. 0809.2932.

Hedou, Julien, Maric, Ivana, Bellan, Gregoire, Einhaus, Jakob, Gaudilliere, K. D, Ladant, Xavier F, Verdonk, Franck, Stelzer, A. I, Feyaerts, Dorien, Tsai, S. A, Ganio, A. E, Sabayev, Maximilian, Gillard, Joshua, Amar, Jonas, Cambriel, Amelie, Oskotsky, T. T, Roldan, Alennie, Golob, L. J, Sirota, Marina, Bonham, A. T, Sato, Masaki, Diop, Maigane, Durand, Xavier, Angst, S. M, Stevenson, K. D, Aghaeepour, Nima, Montanari, Andrea, Gaudilliere, Brice (2024). "Discovery of sparse, reliable omic biomarkers with Stabl." *Nature Biotechnology 2024*, 1–13. ISSN 1546-1696. doi:10.1038/s4158702302033x. https://www.nature.com/articles/s41587-023-02033-x.

Zobolas, John, George, Anne-Marie, Lopez, Alberto, Fischer, Sebastian, Becker, Marc, Aittokallio, Tero (2026). "Prognostic biomarker discovery in pancreatic cancer through hybrid ensemble feature selection and multi-omics data." *BioData Mining*. ISSN 1756-0381. doi:10.1186/s13040026-005460.

Examples

```
eefsr = embedded_ensemble_fselect(
  task = tsk("sonar"),
  learners = lrns(c("classif.rpart", "classif.featureless")),
  init_resampling = rsmp("subsampling", repeats = 5),
  measure = msr("classif.ce")
)
eefsr
```

`ensemble_fselect` *Wrapper-based Ensemble Feature Selection*

Description

Ensemble feature selection using multiple learners. The ensemble feature selection method is designed to identify the most predictive features from a given dataset by leveraging multiple machine learning models and resampling techniques. Returns an EnsembleFSResult.

Usage

```
ensemble_fselect(
  fselector,
  task,
  learners,
  init_resampling,
  inner_resampling,
  inner_measure,
  measure,
  terminator,
  callbacks = NULL,
  store_benchmark_result = TRUE,
  store_models = FALSE
)
```

Arguments

`fselector`	(FSelector) Optimization algorithm.
`task`	(mlr3::Task) Task to operate on.

`learners` (list of mlr3::Learner)
The learners to be used for feature selection.

`init_resampling` (mlr3::Resampling)
The initial resampling strategy of the data, from which each train set will be passed on to the auto_fselector to optimize the learners and perform feature selection. Each test set will be used for prediction on the final models returned by auto_fselector. Can only be mlr3::ResamplingSubsampling or mlr3::ResamplingBootstrap.

`inner_resampling` (mlr3::Resampling)
The inner resampling strategy used by the FSelector.

`inner_measure` (mlr3::Measure)
The inner optimization measure used by the FSelector.

`measure` (mlr3::Measure)
Measure used to score each trained learner on the test sets generated by `init_resampling`.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`callbacks` (Named list of lists of CallbackBatchFSelect)
Callbacks to be used for each learner. The lists must be named by the learner ids.

`store_benchmark_result` (`logical(1)`)
Whether to store the benchmark result in EnsembleFSResult or not.

`store_models` (`logical(1)`)
Whether to store models in auto_fselector or not.

Details

The method begins by applying an initial resampling technique specified by the user, to create **multiple subsamples** from the original dataset (train/test splits). This resampling process helps in generating diverse subsets of data for robust feature selection.

For each subsample (train set) generated in the previous step, the method performs **wrapped-based feature selection** (auto_fselector) using each provided learner, the given inner resampling method, inner performance measure, and optimization algorithm. This process generates 1) the best feature subset and 2) a final trained model using these best features, for each combination of subsample and learner. The final models are then scored on their ability to predict on the resampled test sets.

Results are stored in an EnsembleFSResult.

The result object also includes the performance scores calculated during the inner resampling of the training sets, using models with the best feature subsets. These scores are stored in a column named `{measure_id}_inner`.

Value

an EnsembleFSResult object.

Note

The **active measure** of performance is the one applied to the test sets. This is preferred, as inner resampling scores on the training sets are likely to be overestimated when using the final models. Users can change the active measure by using the `set_active_measure()` method of the EnsembleFSResult.

Source

Saeys, Yvan, Abeel, Thomas, Van De Peer, Yves (2008). "Robust feature selection using ensemble feature selection techniques." *Machine Learning and Knowledge Discovery in Databases*, **5212 LNAI**, 313–325. doi:10.1007/9783540874812_21.

Abeel, Thomas, Helleputte, Thibault, Van de Peer, Yves, Dupont, Pierre, Saeys, Yvan (2010). "Robust biomarker identification for cancer diagnosis with ensemble feature selection methods." *Bioinformatics*, **26**, 392–398. ISSN 1367-4803. doi:10.1093/BIOINFORMATICS/BTP630.

Pes, Barbara (2020). "Ensemble feature selection for high-dimensional data: a stability analysis across multiple domains." *Neural Computing and Applications*, **32**(10), 5951–5973. ISSN 14333058. doi:10.1007/s00521019040823.

Zobolas, John, George, Anne-Marie, Lopez, Alberto, Fischer, Sebastian, Becker, Marc, Aittokallio, Tero (2026). "Prognostic biomarker discovery in pancreatic cancer through hybrid ensemble feature selection and multi-omics data." *BioData Mining*. ISSN 1756-0381. doi:10.1186/s13040026-005460.

Examples

```
efsr = ensemble_fselect(
  fselector = fs("random_search"),
  task = tsk("sonar"),
  learners = lrns(c("classif.rpart", "classif.featureless")),
  init_resampling = rsmp("subsampling", repeats = 2),
  inner_resampling = rsmp("cv", folds = 3),
  inner_measure = msr("classif.ce"),
  measure = msr("classif.acc"),
  terminator = trm("evals", n_evals = 10)
)
efsr
```

---

ensemble_fs_result *Ensemble Feature Selection Result*

Description

The `EnsembleFSResult` stores the results of ensemble feature selection. It includes methods for evaluating the stability of the feature selection process and for ranking the selected features among others.

Both functions `ensemble_fselect()` and `embedded_ensemble_fselect()` return an object of this class.

S3 Methods

- `as.data.table.EnsembleFSResult(x, benchmark_result = TRUE)`
  Returns a tabular view of the ensemble feature selection.
  EnsembleFSResult -> `data.table::data.table()`

  - `x` (EnsembleFSResult)
  - `benchmark_result` (`logical(1)`)
    Whether to add the learner, task and resampling information from the benchmark result.

- `c(...)`
  (EnsembleFSResult, ...) -> EnsembleFSResult
  Combines multiple EnsembleFSResult objects into a new EnsembleFSResult.

Public fields

`benchmark_result` (mlr3::BenchmarkResult)
The benchmark result.

`man` (`character(1)`)
Manual page for this object.

Active bindings

`result` (data.table::data.table)
Returns the result of the ensemble feature selection.

`n_learners` (`numeric(1)`)
Returns the number of learners used in the ensemble feature selection.

`measure` (mlr3::Measure)
Returns the 'active' measure that is used in methods of this object.

`active_measure` (`character(1)`)
Indicates the type of the active performance measure.
During the ensemble feature selection process, the dataset is split into **multiple subsamples** (train/test splits) using an initial resampling scheme. So, performance can be evaluated using one of two measures:

- `"outer"`: measure used to evaluate the performance on the test sets.
- `"inner"`: measure used for optimization and to compute performance during inner resampling on the training sets.

`n_resamples` (`character(1)`)
Returns the number of times the task was initially resampled in the ensemble feature selection process.

Methods

Public methods:

- `EnsembleFSResult$new()`
- `EnsembleFSResult$format()`
- `EnsembleFSResult$print()`

- EnsembleFSResult$help()
- EnsembleFSResult$rm_zero_features()
- EnsembleFSResult$set_active_measure()
- EnsembleFSResult$combine()
- EnsembleFSResult$feature_ranking()
- EnsembleFSResult$stability()
- EnsembleFSResult$pareto_front()
- EnsembleFSResult$knee_points()
- EnsembleFSResult$clone()

`EnsembleFSResult$new()`: Creates a new instance of this R6 class.

Usage:

```
EnsembleFSResult$new(
  result,
  features,
  benchmark_result = NULL,
  measure,
  inner_measure = NULL
)
```

Arguments:

`result` (data.table::data.table)
: The result of the ensemble feature selection. Mandatory column names should include `"resampling_iteration"`, `"learner_id"`, `"features"` and `"n_features"`. A column named as `{measure$id}` (scores on the test sets) must also be always present. The column with the performance scores on the inner resampling of the train sets is not mandatory, but note that it should be named as `{inner_measure$id}_inner` to distinguish from the `{measure$id}`.

`features` (`character()`)
: The vector of features of the task that was used in the ensemble feature selection.

`benchmark_result` (mlr3::BenchmarkResult)
: The benchmark result object.

`measure` (mlr3::Measure)
: The performance measure used to evaluate the learners on the test sets generated during the ensemble feature selection process. By default, this serves as the 'active' measure for the methods of this object. The active measure can be updated using the `$set_active_measure()` method.

`inner_measure` (mlr3::Measure)
: The performance measure used to optimize and evaluate the learners during the inner resampling process of the training sets, generated as part of the ensemble feature selection procedure.

`EnsembleFSResult$format()`: Helper for print outputs.

Usage:

```
EnsembleFSResult$format(...)
```

Arguments:

... (ignored).

`EnsembleFSResult$print()`: Printer.

Usage:

```
EnsembleFSResult$print(...)
```

Arguments:

... (ignored).

`EnsembleFSResult$help()`: Opens the corresponding help page referenced by field `$man`.

Usage:

```
EnsembleFSResult$help()
```

`EnsembleFSResult$rm_zero_features()`: Removes rows from the ensemble feature selection result where no features were selected.

If a benchmark result is stored, the corresponding resampling iterations are removed as well. The stability measures are reset and need to be recalculated after this operation.

This method modifies the object by reference. To preserve the original state, explicitly `$clone()` the object beforehand.

Usage:

```
EnsembleFSResult$rm_zero_features()
```

Returns: Returns the object itself, but modified **by reference**.

`EnsembleFSResult$set_active_measure()`: Use this function to change the active measure.

Usage:

```
EnsembleFSResult$set_active_measure(which = "outer")
```

Arguments:

`which (character(1))`
: Which measure from the ensemble feature selection result to use in methods of this object. Should be either `"inner"` (optimization measure used in training sets) or `"outer"` (measure used in test sets, default value).

`EnsembleFSResult$combine()`: Combines a second EnsembleFSResult into the current object, modifying it **in-place**. If the second EnsembleFSResult (`efsr`) is `NULL`, the method returns the object unmodified.

Both objects must have the same task features and `measure`. If the `inner_measure` differs between the objects or is `NULL` in either, it will be set to `NULL` in the combined object. Additionally, the `importance` column will be removed if it is missing in either object. If both objects contain a `benchmark_result`, these will be combined. Otherwise, the combined object will have a `NULL` value for `benchmark_result`.

This method modifies the object by reference. To preserve the original state, explicitly `$clone()` the object beforehand. Alternatively, you can use the `c()` function, which internally calls this method.

Usage:

```
EnsembleFSResult$combine(efsr)
```

Arguments:

`efsr` (EnsembleFSResult)
: A second EnsembleFSResult object to combine with the current object.

Returns: Returns the object itself, but modified **by reference**.

`EnsembleFSResult$feature_ranking()`: Calculates the feature ranking via `fastVoteR::rank_candidates()`.

Usage:

```
EnsembleFSResult$feature_ranking(
  method = "av",
  use_weights = TRUE,
  committee_size = NULL,
  shuffle_features = TRUE
)
```

Arguments:

`method (character(1))`
: The method to calculate the feature ranking. See `fastVoteR::rank_candidates()` for a complete list of available methods. Approval voting (`"av"`) is the default method.

`use_weights (logical(1))`
: The default value (`TRUE`) uses weights equal to the performance scores of each voter/model (or the inverse scores if the measure is minimized). Note that the performance scores need to be non-negative for the weights to be meaningful. If the scores can be negative, it is recommended to set `use_weights = FALSE`, which treats all voters as equal and assigns them the same weight equal to 1.

`committee_size (integer(1))`
: The number of top-ranked features to return. This can speed up methods that build a committee sequentially (e.g., `"seq_pav"`) by computing only the top N candidates rather than the full ranking.

`shuffle_features (logical(1))`
: Whether to shuffle the task features randomly before computing the ranking. Shuffling ensures consistent random tie-breaking across methods and prevents deterministic biases when features with equal scores are encountered. Default is `TRUE` and it's advised to set a seed before running this function. Set to `FALSE` if deterministic ordering of features is preferred (same as during initialization).

Details: The feature ranking process is built on the following framework: models act as *voters*, features act as *candidates*, and voters select certain candidates (features). The primary objective is to compile these selections into a consensus ranked list of features, effectively forming a committee.
For every feature a score is calculated, which depends on the `"method"` argument. The higher the score, the higher the ranking of the feature. Note that some methods output a feature ranking instead of a score per feature, so we always include **Borda's score**, which is method-agnostic, i.e. it can be used to compare the feature rankings across different methods.
We shuffle the input candidates/features so that we enforce random tie-breaking. Users should set the same `seed` for consistent comparison between the different feature ranking methods and for reproducibility.

Returns: A data.table::data.table listing all the features, ordered by decreasing scores (depends on the `"method"`). Columns are as follows:

- `"feature"`: Feature names.
- `"score"`: Scores assigned to each feature based on the selected method (if applicable).
- `"norm_score"`: Normalized scores (if applicable), scaled to the range $[0, 1]$, which can be loosely interpreted as **selection probabilities** (Meinshausen et al. (2010)).
- `"borda_score"`: Borda scores for method-agnostic comparison, ranging in $[0, 1]$, where the top feature receives a score of 1 and the lowest-ranked feature receives a score of 0. This column is always included so that feature ranking methods that output only rankings have also a feature-wise score.

`EnsembleFSResult$stability()`: Calculates the stability of the selected features with the **stabm** package. The results are cached. The cache key covers the stability measure and the arguments passed to it, so the same measure requested with different arguments is calculated again.

Usage:

```
EnsembleFSResult$stability(
  stability_measure = "jaccard",
  stability_args = NULL,
  global = TRUE,
  reset_cache = FALSE
)
```

Arguments:

`stability_measure (character(1))`
: The stability measure to be used. One of the measures returned by `stabm::listStabilityMeasures()` in lower case. Default is `"jaccard"`.

`stability_args (list)`
: Additional arguments passed to the stability measure function.

`global (logical(1))`
: Whether to calculate the stability globally or for each learner.

`reset_cache (logical(1))`
: If `TRUE`, the cached results are ignored.

Returns: A `numeric()` value representing the stability of the selected features. Or a `numeric()` vector with the stability of the selected features for each learner.

`EnsembleFSResult$pareto_front()`: This function identifies the **Pareto front** of the ensemble feature selection process, i.e., the set of points that represent the trade-off between the number of features and performance (e.g. classification error).

Usage:

```
EnsembleFSResult$pareto_front(type = "empirical", max_nfeatures = NULL)
```

Arguments:

`type (character(1))`
: Specifies the type of Pareto front to return. See details.

`max_nfeatures (integer(1))`
: Specifies the maximum number of features for which the estimated Pareto front is computed. Applicable only when `type = "estimated"`. If `NULL` (default), the maximum number of features is determined by the ensemble feature selection process.

Details: Two options are available for the Pareto front:

- `"empirical"` (default): returns the empirical Pareto front.
- `"estimated"`: the Pareto front points are estimated by fitting a linear model with the inverse of the number of features ($1/x$) as input and the associated performance scores as output.

This method is useful when the Pareto points are sparse and the front assumes a convex shape if better performance corresponds to lower measure values (e.g. classification error), or a concave shape otherwise (e.g. classification accuracy).
When `type = "estimated"`, the estimated Pareto front includes points with the number of features ranging from 1 up to `max_nfeatures`. If `max_nfeatures` is not provided, it defaults to the maximum number of features available in the ensemble feature selection `result`, i.e. the maximum out of all learners and resamplings included.

Returns: A data.table::data.table with columns the number of features and the performance that together form the Pareto front.

`EnsembleFSResult$knee_points()`: This function implements various *knee* point identification (KPI) methods, which select points in the Pareto front, such that an optimal trade-off between performance and number of features is achieved. In most cases, only one such point is returned.

Usage:

```
EnsembleFSResult$knee_points(
  method = "NBI",
  type = "empirical",
  max_nfeatures = NULL
)
```

Arguments:

`method (character(1))`
: Type of method to use to identify the knee point.

`type (character(1))`
: Specifies the type of Pareto front to use for the identification of the knee point.

`max_nfeatures (integer(1))`
: Specifies the maximum number of features for which the estimated Pareto front is computed. Applicable only when `type = "estimated"`. If `NULL` (default), the maximum number of features is determined by the ensemble feature selection process. See `pareto_front()` method for more details.

Details: The available KPI methods are:

- `"NBI"` (default): The **Normal-Boundary Intersection** method is a geometry-based method which calculates the perpendicular distance of each point from the line connecting the first and last points of the Pareto front. The knee point is determined as the Pareto point with the maximum distance from this line, see Das (1999).

Returns: A data.table::data.table with the knee point(s) of the Pareto front. If the Pareto front does not span a range in both the number of features and the performance, no knee point can be identified. In this case a warning is signaled and the first point of the Pareto front is returned.

`EnsembleFSResult$clone()`: The objects of this class are cloneable with this method.

Usage:

```
EnsembleFSResult$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

References

Das, I (1999). "On characterizing the 'knee' of the Pareto curve based on normal-boundary intersection." *Structural Optimization*, **18**(1-2), 107–115. ISSN 09344373.

Meinshausen, Nicolai, Buhlmann, Peter (2010). "Stability Selection." *Journal of the Royal Statistical Society Series B: Statistical Methodology*, **72**(4), 417–473. ISSN 1369-7412. doi:10.1111/J.14679868.2010.00740.X. 0809.2932.

Zobolas, John, George, Anne-Marie, Lopez, Alberto, Fischer, Sebastian, Becker, Marc, Aittokallio, Tero (2026). "Prognostic biomarker discovery in pancreatic cancer through hybrid ensemble feature selection and multi-omics data." *BioData Mining*. ISSN 1756-0381. doi:10.1186/s13040026-005460.

Examples

```
efsr = ensemble_fselect(
  fselector = fs("rfe", n_features = 2, feature_fraction = 0.8),
  task = tsk("sonar"),
  learners = lrns(c("classif.rpart", "classif.featureless")),
  init_resampling = rsmp("subsampling", repeats = 2),
  inner_resampling = rsmp("cv", folds = 3),
  inner_measure = msr("classif.ce"),
  measure = msr("classif.acc"),
  terminator = trm("none")
)

# contains the benchmark result
efsr$benchmark_result

# contains the selected features for each iteration
efsr$result

# returns the stability of the selected features
efsr$stability(stability_measure = "jaccard")

# returns a ranking of all features
head(efsr$feature_ranking())

# returns the empirical pareto front, i.e. n_features vs measure (error)
efsr$pareto_front()

# returns the knee points (optimal trade-off between n_features and performance)
efsr$knee_points()

# change to use the inner optimization measure
efsr$set_active_measure(which = "inner")

# Pareto front is calculated on the inner measure
efsr$pareto_front()
```

`extract_inner_fselect_archives` *Extract Inner Feature Selection Archives*

Description

Extract inner feature selection archives of nested resampling. Implemented for mlr3::ResampleResult and mlr3::BenchmarkResult. The function iterates over the AutoFSelector objects and binds the archives to a `data.table::data.table()`. AutoFSelector must be initialized with `store_fselect_instance = TRUE` and `resample()` or `benchmark()` must be called with `store_models = TRUE`.

Usage

```
extract_inner_fselect_archives(x, exclude_columns = "uhash")
```

Arguments

`x`	(mlr3::ResampleResult | mlr3::BenchmarkResult).
`exclude_columns`	(`character()`) Exclude columns from result table. Set to `NULL` if no column should be excluded.

Value

`data.table::data.table()`.

Data structure

The returned data table has the following columns:

- `experiment` (integer(1))
  Index, giving the according row number in the original benchmark grid.
- `iteration` (integer(1))
  Iteration of the outer resampling.
- One column for each feature of the task.
- One column for each performance measure.
- `runtime_learners` (`numeric(1)`)
  Sum of training and predict times logged in learners per mlr3::ResampleResult / evaluation. This does not include potential overhead time.
- `timestamp` (`POSIXct`)
  Time stamp when the evaluation was logged into the archive.
- `batch_nr` (`integer(1)`)
  Feature sets are evaluated in batches. Each batch has a unique batch number.
- `resample_result` (mlr3::ResampleResult)
  Resample result of the inner resampling.

- task_id (character(1)).
- learner_id (character(1)).
- resampling_id (character(1)).

Examples

```
# Nested Resampling on Palmer Penguins Data Set

# create auto fselector
at = auto_fselector(
  fselector = fs("random_search"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("holdout"),
  measure = msr("classif.ce"),
  term_evals = 4)

resampling_outer = rsmp("cv", folds = 2)
rr = resample(tsk("penguins"), at, resampling_outer, store_models = TRUE)

# extract inner archives
extract_inner_fselect_archives(rr)
```

---

extract_inner_fselect_results

Extract Inner Feature Selection Results

Description

Extract inner feature selection results of nested resampling. Implemented for mlr3::ResampleResult and mlr3::BenchmarkResult.

Usage

```
extract_inner_fselect_results(x, fselect_instance, ...)
```

Arguments

x	(mlr3::ResampleResult | mlr3::BenchmarkResult).
fselect_instance	(logical(1)) If TRUE, instances are added to the table.
...	(any) Additional arguments.

Details

The function iterates over the AutoFSelector objects and binds the feature selection results to a data.table::data.table(). AutoFSelector must be initialized with store_fselect_instance = TRUE and resample() or benchmark() must be called with store_models = TRUE. Optionally, the instance can be added for each iteration.

Value

`data.table::data.table()`.

Data structure

The returned data table has the following columns:

- `experiment` (integer(1))
  Index, giving the according row number in the original benchmark grid.
- `iteration` (integer(1))
  Iteration of the outer resampling.
- One column for each feature of the task.
- One column for each performance measure.
- `features` (character())
  Vector of selected feature set.
- `task_id` (`character(1)`).
- `learner_id` (`character(1)`).
- `resampling_id` (`character(1)`).

Examples

```
# Nested Resampling on Palmer Penguins Data Set

# create auto fselector
at = auto_fselector(
  fselector = fs("random_search"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("holdout"),
  measure = msr("classif.ce"),
  term_evals = 4)

resampling_outer = rsmp("cv", folds = 2)
rr = resample(tsk("iris"), at, resampling_outer, store_models = TRUE)

# extract inner results
extract_inner_fselect_results(rr)
```

---

faggregate *Fast Aggregation of ResampleResults and BenchmarkResults*

Description

Aggregates a mlr3::ResampleResult or mlr3::BenchmarkResult for a single simple measure. Returns the aggregated score for each resample result.

Usage

```
faggregate(obj, measure, conditions = FALSE)
```

Arguments

obj	(mlr3::ResampleResult | mlr3::BenchmarkResult).
measure	(mlr3::Measure).
conditions	(logical(1)) If TRUE, the function returns the number of warnings and the number of errors.

Details

This function is faster than $aggregate() because it does not reassemble the resampling results. It only works on simple measures which do not require the task, learner, model or train set to be available.

Value

(data.table::data.table())

---

fs *Syntactic Sugar for Feature Selection Objects Construction*

Description

Functions to retrieve objects, set parameters and assign to fields in one go. Relies on mlr3misc::dictionary_sugar_get() to extract objects from the respective mlr3misc::Dictionary:

- fs() for a FSelector from mlr_fselectors.
- fss() for a list of FSelectors from mlr_fselectors.
- trm() for a bbotk::Terminator from mlr_terminators.
- trms() for a list of Terminators from mlr_terminators.

Usage

```
fs(.key, ...)

fss(.keys, ...)
```

Arguments

.key	(character(1)) Key passed to the respective dictionary to retrieve the object.
...	(any) Additional arguments.
.keys	(character()) Keys passed to the respective dictionary to retrieve multiple objects.

Value

R6::R6Class object of the respective type, or a list of R6::R6Class objects for the plural versions.

Examples

```
# random search fselector with batch size of 5
fs("random_search", batch_size = 5)

# run time terminator with 20 seconds
trm("run_time", secs = 20)
```

---

fselect *Function for Feature Selection*

Description

Function to optimize the features of a mlr3::Learner. The function internally creates a FSelectInstanceBatchSingleCrit or FSelectInstanceBatchMultiCrit which describes the feature selection problem. It executes the feature selection with the FSelector (`fselector`) and returns the result with the feature selection instance (`$result`). The ArchiveBatchFSelect and ArchiveAsyncFSelect (`$archive`) stores all evaluated feature subsets and performance scores.

You can find an overview of all feature selectors on our website.

Usage

```
fselect(
  fselector,
  task,
  learner,
  resampling,
  measures = NULL,
  term_evals = NULL,
  term_time = NULL,
  terminator = NULL,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features",
  rush = NULL
)
```

Arguments

`fselector`	(FSelector) Optimization algorithm.

`task` (mlr3::Task)
Task to operate on.

`learner` (mlr3::Learner)
Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (mlr3::Measure or list of mlr3::Measure)
A single measure creates a FSelectInstanceBatchSingleCrit and multiple measures a FSelectInstanceBatchMultiCrit. If `NULL`, default measure is used.

`term_evals` (`integer(1)`)
Number of allowed evaluations. Ignored if `terminator` is passed.

`term_time` (`integer(1)`)
Maximum allowed time in seconds. Ignored if `terminator` is passed.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_benchmark_result`
(`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ties_method` (`character(1)`)
The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`rush` (`Rush`)
Rush instance for the asynchronous feature selection. Ignored by batch fselectors, which warn when it is set.

Details

The mlr3::Task, mlr3::Learner, mlr3::Resampling, mlr3::Measure and bbotk::Terminator are used to construct a FSelectInstanceBatchSingleCrit. If multiple performance mlr3::Measures are supplied, a FSelectInstanceBatchMultiCrit is created. The parameter `term_evals` and `term_time` are shortcuts to create a bbotk::Terminator. If both parameters are passed, a bbotk::TerminatorCombo is constructed. For other Terminators, pass one with `terminator`. If no termination criterion is needed, set `term_evals`, `term_time` and `terminator` to `NULL`.

Value

FSelectInstanceBatchSingleCrit | FSelectInstanceBatchMultiCrit

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
"classif"	"classif.ce"	**mlr3**
"regr"	"regr.mse"	**mlr3**
"surv"	"surv.cindex"	**mlr3proba**
"dens"	"dens.logloss"	**mlr3proba**
"classif_st"	"classif.ce"	**mlr3spatial**
"regr_st"	"regr.mse"	**mlr3spatial**
"clust"	"clust.dunn"	**mlr3cluster**

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Analysis

For analyzing the feature selection results, it is recommended to pass the archive to `as.data.table()`. The returned data table is joined with the benchmark result which adds the mlr3::ResampleResult for each feature set.

The archive provides various getters (e.g. `$learners()`) to ease the access. All getters extract by position (`i`) or unique hash (`uhash`). For a complete list of all getters see the methods section.

The benchmark result (`$benchmark_result`) allows to score the feature sets again on a different measure. Alternatively, measures can be supplied to `as.data.table()`.

Examples

```
# Feature selection on the diabetes data set
task = tsk("diabetes")

# Load learner
learner = lrn("classif.rpart")

# Run feature selection
instance = fselect(
```

```
  fselector = fs("random_search", batch_size = 2),
  task = task,
  learner = learner,
  resampling = rsmp ("holdout"),
  measures = msr("classif.ce"),
  term_evals = 4)

# Subset task to optimized feature set
task$select(instance$result_feature_set)

# Train the learner with optimal feature set on the full data set
learner$train(task)

# Inspect all evaluated feature subsets
as.data.table(instance$archive)
```

`FSelectInstanceAsyncMultiCrit`

Multi-Criteria Feature Selection with Rush

Description

The `FSelectInstanceAsyncMultiCrit` specifies a feature selection problem for a FSelectorAsync. The function `fsi_async()` creates a FSelectInstanceAsyncMultiCrit and the function `fselect()` creates an instance internally.

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
`"classif"`	`"classif.ce"`	**mlr3**
`"regr"`	`"regr.mse"`	**mlr3**
`"surv"`	`"surv.cindex"`	**mlr3proba**
`"dens"`	`"dens.logloss"`	**mlr3proba**
`"classif_st"`	`"classif.ce"`	**mlr3spatial**
`"regr_st"`	`"regr.mse"`	**mlr3spatial**
`"clust"`	`"clust.dunn"`	**mlr3cluster**

Analysis

For analyzing the feature selection results, it is recommended to pass the ArchiveAsyncFSelect to `as.data.table()`. The returned data table contains the mlr3::ResampleResult for each feature subset evaluation.

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Super classes

`bbotk::EvalInstance` -> `bbotk::OptimInstance` -> `bbotk::OptimInstanceAsync` -> `bbotk::OptimInstanceAsyncMul`
-> `FSelectInstanceAsyncMultiCrit`

Methods

Public methods:

- `FSelectInstanceAsyncMultiCrit$new()`
- `FSelectInstanceAsyncMultiCrit$assign_result()`
- `FSelectInstanceAsyncMultiCrit$clone()`

`FSelectInstanceAsyncMultiCrit$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectInstanceAsyncMultiCrit$new(
  task,
  learner,
  resampling,
  measures,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  rush = NULL
)
```

Arguments:

`task` (mlr3::Task)
: Task to operate on.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (list of mlr3::Measure)
Measures to optimize. If `NULL`, **mlr3**'s default measure is used.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_benchmark_result` (`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`rush` (`Rush`)
If a rush instance is supplied, the optimization runs without batches.

`FSelectInstanceAsyncMultiCrit$assign_result()`: The FSelectorAsync object writes the best found points and estimated performance values here (probably the Pareto set / front). For internal use.

Usage:

```
FSelectInstanceAsyncMultiCrit$assign_result(xdt, ydt, extra = NULL, ...)
```

Arguments:

`xdt` (`data.table::data.table()`)
x values as `data.table`. Each row is one point. Contains the value in the *search space* of the FSelectInstanceBatchMultiCrit object. Can contain additional columns for extra information.

`ydt` (`numeric()`)
Optimal outcomes, e.g. the Pareto front.

`extra` (`data.table::data.table()`)
Additional information.

`...` (`any`)
ignored.

`FSelectInstanceAsyncMultiCrit$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectInstanceAsyncMultiCrit$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`FSelectInstanceAsyncSingleCrit`
Single Criterion Feature Selection with Rush

Description

The `FSelectInstanceAsyncSingleCrit` specifies a feature selection problem for a FSelectorAsync. The function `fsi_async()` creates a FSelectInstanceAsyncSingleCrit and the function `fselect()` creates an instance internally.

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
`"classif"`	`"classif.ce"`	**mlr3**
`"regr"`	`"regr.mse"`	**mlr3**
`"surv"`	`"surv.cindex"`	**mlr3proba**
`"dens"`	`"dens.logloss"`	**mlr3proba**
`"classif_st"`	`"classif.ce"`	**mlr3spatial**
`"regr_st"`	`"regr.mse"`	**mlr3spatial**
`"clust"`	`"clust.dunn"`	**mlr3cluster**

Analysis

For analyzing the feature selection results, it is recommended to pass the ArchiveAsyncFSelect to `as.data.table()`. The returned data table contains the mlr3::ResampleResult for each feature subset evaluation.

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Super classes

`bbotk::EvalInstance` -> `bbotk::OptimInstance` -> `bbotk::OptimInstanceAsync` -> `bbotk::OptimInstanceAsyncSi`
-> `FSelectInstanceAsyncSingleCrit`

Methods

Public methods:

- `FSelectInstanceAsyncSingleCrit$new()`
- `FSelectInstanceAsyncSingleCrit$assign_result()`
- `FSelectInstanceAsyncSingleCrit$clone()`

`FSelectInstanceAsyncSingleCrit$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectInstanceAsyncSingleCrit$new(
  task,
  learner,
  resampling,
  measure = NULL,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features",
  rush = NULL
)
```

Arguments:

`task` (mlr3::Task)
: Task to operate on.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measure` (mlr3::Measure)
: Measure to optimize. If `NULL`, default measure is used.

`terminator` (bbotk::Terminator)
: Stop criterion of the feature selection.

`store_benchmark_result` (`logical(1)`)
: Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
: Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
: List of callbacks.

`ties_method` (`character(1)`)
: The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the

same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`rush` (`Rush`)
If a rush instance is supplied, the optimization runs without batches.

`FSelectInstanceAsyncSingleCrit$assign_result()`: The FSelectorAsync object writes the best found point and estimated performance value here. For internal use.

Usage:

```
FSelectInstanceAsyncSingleCrit$assign_result(xdt, y, extra = NULL, ...)
```

Arguments:

`xdt` (`data.table::data.table()`)
x values as `data.table`. Each row is one point. Contains the value in the *search space* of the FSelectInstanceBatchMultiCrit object. Can contain additional columns for extra information.

`y` (`numeric(1)`)
Optimal outcome.

`extra` (`data.table::data.table()`)
Additional information.

`...` (`any`)
ignored.

`FSelectInstanceAsyncSingleCrit$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectInstanceAsyncSingleCrit$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`FSelectInstanceBatchMultiCrit`
Class for Multi Criteria Feature Selection

Description

The FSelectInstanceBatchMultiCrit specifies a feature selection problem for a FSelector. The function `fsi()` creates a FSelectInstanceBatchMultiCrit and the function `fselect()` creates an instance internally.

Resources

There are several sections about feature selection in the mlr3book.

- Learn about multi-objective optimization.

The gallery features a collection of case studies and demos about optimization.

Analysis

For analyzing the feature selection results, it is recommended to pass the archive to `as.data.table()`. The returned data table is joined with the benchmark result which adds the mlr3::ResampleResult for each feature set.

The archive provides various getters (e.g. `$learners()`) to ease the access. All getters extract by position (`i`) or unique hash (`uhash`). For a complete list of all getters see the methods section.

The benchmark result (`$benchmark_result`) allows to score the feature sets again on a different measure. Alternatively, measures can be supplied to `as.data.table()`.

Super classes

`bbotk::EvalInstance` -> `bbotk::OptimInstance` -> `bbotk::OptimInstanceBatch` -> `bbotk::OptimInstanceBatchMul`
-> `FSelectInstanceBatchMultiCrit`

Active bindings

`result_feature_set` (list of `character()`)
Feature sets for task subsetting.

Methods

Public methods:

- `FSelectInstanceBatchMultiCrit$new()`
- `FSelectInstanceBatchMultiCrit$assign_result()`
- `FSelectInstanceBatchMultiCrit$print()`
- `FSelectInstanceBatchMultiCrit$clone()`

`FSelectInstanceBatchMultiCrit$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectInstanceBatchMultiCrit$new(
  task,
  learner,
  resampling,
  measures,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL
)
```

Arguments:

`task` (mlr3::Task)
Task to operate on.

`learner` (mlr3::Learner)
Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (list of mlr3::Measure)
Measures to optimize. If `NULL`, **mlr3**'s default measure is used.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_benchmark_result` (`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`FSelectInstanceBatchMultiCrit$assign_result()`: The FSelector object writes the best found feature subsets and estimated performance values here. For internal use.

Usage:

```
FSelectInstanceBatchMultiCrit$assign_result(xdt, ydt, extra = NULL, ...)
```

Arguments:

`xdt` (`data.table::data.table()`)
x values as `data.table`. Each row is one point. Contains the value in the *search space* of the FSelectInstanceBatchMultiCrit object. Can contain additional columns for extra information.

`ydt` (`data.table::data.table()`)
Optimal outcomes, e.g. the Pareto front.

`extra` (`data.table::data.table()`)
Additional information.

`...` (any)
ignored.

`FSelectInstanceBatchMultiCrit$print()`: Printer.

Usage:

```
FSelectInstanceBatchMultiCrit$print(...)
```

Arguments:

`...` (ignored).

`FSelectInstanceBatchMultiCrit$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectInstanceBatchMultiCrit$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Examples

```
# Feature selection on Palmer Penguins data set


task = tsk("penguins")

# Construct feature selection instance
instance = fsi(
  task = task,
  learner = lrn("classif.rpart"),
  resampling = rsmp("cv", folds = 3),
  measures = msrs(c("classif.ce", "time_train")),
  terminator = trm("evals", n_evals = 4)
)

# Choose optimization algorithm
fselector = fs("random_search", batch_size = 2)

# Run feature selection
fselector$optimize(instance)

# Optimal feature sets
instance$result_feature_set

# Inspect all evaluated sets
as.data.table(instance$archive)
```

`FSelectInstanceBatchSingleCrit`

Class for Single Criterion Feature Selection

Description

The FSelectInstanceBatchSingleCrit specifies a feature selection problem for a FSelector. The function `fsi()` creates a FSelectInstanceBatchSingleCrit and the function `fselect()` creates an instance internally.

The instance contains an ObjectiveFSelectBatch object that encodes the black box objective function a FSelector has to optimize. The instance allows the basic operations of querying the objective at design points (`$eval_batch()`). This operation is usually done by the FSelector. Evaluations of feature subsets are performed in batches by calling `mlr3::benchmark()` internally. The evaluated feature subsets are stored in the Archive (`$archive`). Before a batch is evaluated, the bbotk::Terminator is queried for the remaining budget. If the available budget is exhausted, an exception is raised, and no further evaluations can be performed from this point on. The FSelector is also supposed to store its final result, consisting of a selected feature subset and associated estimated performance values, by calling the method `instance$assign_result()`.

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
"classif"	"classif.ce"	**mlr3**
"regr"	"regr.mse"	**mlr3**
"surv"	"surv.cindex"	**mlr3proba**
"dens"	"dens.logloss"	**mlr3proba**
"classif_st"	"classif.ce"	**mlr3spatial**
"regr_st"	"regr.mse"	**mlr3spatial**
"clust"	"clust.dunn"	**mlr3cluster**

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Analysis

For analyzing the feature selection results, it is recommended to pass the archive to `as.data.table()`. The returned data table is joined with the benchmark result which adds the mlr3::ResampleResult for each feature set.

The archive provides various getters (e.g. `$learners()`) to ease the access. All getters extract by position (`i`) or unique hash (`uhash`). For a complete list of all getters see the methods section.

The benchmark result (`$benchmark_result`) allows to score the feature sets again on a different measure. Alternatively, measures can be supplied to `as.data.table()`.

Super classes

`bbotk::EvalInstance` -> `bbotk::OptimInstance` -> `bbotk::OptimInstanceBatch` -> `bbotk::OptimInstanceBatchSi`
-> `FSelectInstanceBatchSingleCrit`

Active bindings

`result_feature_set (character())`
Feature set for task subsetting.

Methods

Public methods:

- FSelectInstanceBatchSingleCrit$new()
- FSelectInstanceBatchSingleCrit$assign_result()
- FSelectInstanceBatchSingleCrit$print()
- FSelectInstanceBatchSingleCrit$clone()

`FSelectInstanceBatchSingleCrit$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectInstanceBatchSingleCrit$new(
  task,
  learner,
  resampling,
  measure,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features"
)
```

Arguments:

`task` (mlr3::Task)
: Task to operate on.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measure` (mlr3::Measure)
: Measure to optimize. If `NULL`, default measure is used.

`terminator` (bbotk::Terminator)
: Stop criterion of the feature selection.

`store_benchmark_result` (`logical(1)`)
: Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
: Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
: List of callbacks.

`ties_method` (`character(1)`)
: The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the

same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

`FSelectInstanceBatchSingleCrit$assign_result()`: The FSelector writes the best found feature subset and estimated performance value here. For internal use.

Usage:

```
FSelectInstanceBatchSingleCrit$assign_result(xdt, y, extra = NULL, ...)
```

Arguments:

`xdt` (`data.table::data.table()`)
x values as `data.table`. Each row is one point. Contains the value in the *search space* of the FSelectInstanceBatchMultiCrit object. Can contain additional columns for extra information.

`y` (`numeric(1)`)
Optimal outcome.

`extra` (`data.table::data.table()`)
Additional information.

`...` (any)
ignored.

`FSelectInstanceBatchSingleCrit$print()`: Printer.

Usage:

```
FSelectInstanceBatchSingleCrit$print(...)
```

Arguments:

`...` (ignored).

`FSelectInstanceBatchSingleCrit$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectInstanceBatchSingleCrit$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Examples

```
# Feature selection on Palmer Penguins data set


task = tsk("penguins")
learner = lrn("classif.rpart")

# Construct feature selection instance
instance = fsi(
  task = task,
  learner = learner,
  resampling = rsmp("cv", folds = 3),
  measures = msr("classif.ce"),
```

```
  terminator = trm("evals", n_evals = 4)
)

# Choose optimization algorithm
fselector = fs("random_search", batch_size = 2)

# Run feature selection
fselector$optimize(instance)

# Subset task to optimal feature set
task$select(instance$result_feature_set)

# Train the learner with optimal feature set on the full data set
learner$train(task)

# Inspect all evaluated sets
as.data.table(instance$archive)
```

`FSelector` *FSelector*

Description

The 'FSelector" implements the optimization algorithm.

Details

`FSelector` is an abstract base class that implements the base functionality each fselector must provide.

Resources

There are several sections about feature selection in the mlr3book.

- Learn more about fselectors.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Public fields

`id (character(1))`
: Identifier of the object. Used in tables, plot and text output.

Active bindings

`param_set` paradox::ParamSet
: Set of control parameters.

`properties` (`character()`)
: Set of properties of the fselector. Must be a subset of `mlr_reflections$fselect_properties`.

`packages` (`character()`)
: Set of required packages. Note that these packages will be loaded via `requireNamespace()`, and are not attached.

`label` (`character(1)`)
: Label for this object. Can be used in tables, plot and text output instead of the ID.

`man` (`character(1)`)
: String in the format `[pkg]::[topic]` pointing to a manual page for this object. The referenced help package can be opened via method `$help()`.

Methods

Public methods:

- `FSelector$new()`
- `FSelector$format()`
- `FSelector$print()`
- `FSelector$help()`
- `FSelector$clone()`

`FSelector$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelector$new(
  id = "fselector",
  param_set,
  properties,
  packages = character(),
  label = NA_character_,
  man = NA_character_
)
```

Arguments:

`id` (`character(1)`)
: Identifier for the new instance.

`param_set` paradox::ParamSet
: Set of control parameters.

`properties` (`character()`)
: Set of properties of the fselector. Must be a subset of `mlr_reflections$fselect_properties`.

`packages` (`character()`)
: Set of required packages. Note that these packages will be loaded via `requireNamespace()`, and are not attached.

`label` (`character(1)`)
: Label for this object. Can be used in tables, plot and text output instead of the ID.

`man` (`character(1)`)
String in the format `[pkg]::[topic]` pointing to a manual page for this object. The referenced help package can be opened via method `$help()`.

`FSelector$format()`: Helper for print outputs.

Usage:

```
FSelector$format(...)
```

Arguments:

`...` (ignored).

Returns: (`character()`).

`FSelector$print()`: Print method.

Usage:

```
FSelector$print(...)
```

Arguments:

`...` (ignored).

Returns: (`character()`).

`FSelector$help()`: Opens the corresponding help page referenced by field `$man`.

Usage:

```
FSelector$help()
```

`FSelector$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelector$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

See Also

Other FSelector: `mlr_fselectors`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_genetic_search`, `mlr_fselectors_random_search`, `mlr_fselectors_rfe`, `mlr_fselectors_rfecv`, `mlr_fselectors_sequential`, `mlr_fselectors_shadow_variable_search`

`FSelectorAsync` *Class for Asynchronous Feature Selection Algorithms*

Description

The FSelectorAsync implements the asynchronous optimization algorithm.

Details

FSelectorAsync is an abstract base class that implements the base functionality each asynchronous fselector must provide.

Resources

There are several sections about feature selection in the mlr3book.

- Learn more about fselectors.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Super class

`FSelector` -> `FSelectorAsync`

Methods

Public methods:

- `FSelectorAsync$optimize()`
- `FSelectorAsync$clone()`

`FSelectorAsync$optimize()`: Performs the feature selection on a FSelectInstanceAsyncSingleCrit or FSelectInstanceAsyncMultiCrit until termination. The single evaluations will be written into the ArchiveAsyncFSelect that resides in the FSelectInstanceAsyncSingleCrit/FSelectInstanceAsyncMultiCrit. The result will be written into the instance object.

Usage:

```
FSelectorAsync$optimize(inst)
```

Arguments:

`inst` (FSelectInstanceAsyncSingleCrit | FSelectInstanceAsyncMultiCrit).

Returns: `data.table::data.table()`

`FSelectorAsync$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorAsync$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

FSelectorBatch *Class for Batch Feature Selection Algorithms*

Description

The FSelectorBatch implements the optimization algorithm.

Details

FSelectorBatch is an abstract base class that implements the base functionality each fselector must provide. A subclass is implemented in the following way:

- Inherit from FSelectorBatch.
- Specify the private abstract method `$.optimize()` and use it to call into your optimizer.
- You need to call `instance$eval_batch()` to evaluate design points.
- The batch evaluation is requested at the FSelectInstanceBatchSingleCrit/FSelectInstanceBatchMultiCrit object `instance`, so each batch is possibly executed in parallel via `mlr3::benchmark()`, and all evaluations are stored inside of `instance$archive`.
- Before the batch evaluation, the bbotk::Terminator is checked, and if it is positive, an exception of class `"terminated_error"` is generated. In the latter case the current batch of evaluations is still stored in `instance`, but the numeric scores are not sent back to the handling optimizer as it has lost execution control.
- After such an exception was caught we select the best set from `instance$archive` and return it.
- Note that therefore more points than specified by the bbotk::Terminator may be evaluated, as the Terminator is only checked before a batch evaluation, and not in-between evaluation in a batch. How many more depends on the setting of the batch size.
- Overwrite the private super-method `.assign_result()` if you want to decide how to estimate the final set in the instance and its estimated performance. The default behavior is: We pick the best resample experiment, regarding the given measure, then assign its set and aggregated performance to the instance.

Private Methods

- `.optimize(instance)` -> `NULL`
  Abstract base method. Implement to specify feature selection of your subclass. See technical details sections.
- `.assign_result(instance)` -> `NULL`
  Abstract base method. Implement to specify how the final feature subset is selected. See technical details sections.

Resources

There are several sections about feature selection in the mlr3book.

- Learn more about fselectors.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Super class

`FSelector` -> `FSelectorBatch`

Methods

Public methods:

- `FSelectorBatch$new()`
- `FSelectorBatch$optimize()`
- `FSelectorBatch$clone()`

`FSelectorBatch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorBatch$new(
  id = "fselector_batch",
  param_set,
  properties,
  packages = character(),
  label = NA_character_,
  man = NA_character_
)
```

Arguments:

`id` (`character(1)`)
: Identifier for the new instance.

`param_set` paradox::ParamSet
: Set of control parameters.

`properties` (`character()`)
: Set of properties of the fselector. Must be a subset of `mlr_reflections$fselect_properties`.

`packages` (`character()`)
: Set of required packages. Note that these packages will be loaded via `requireNamespace()`, and are not attached.

`label` (`character(1)`)
: Label for this object. Can be used in tables, plot and text output instead of the ID.

`man` (`character(1)`)
: String in the format `[pkg]::[topic]` pointing to a manual page for this object. The referenced help package can be opened via method `$help()`.

`FSelectorBatch$optimize()`: Performs the feature selection on a FSelectInstanceBatchSingleCrit or FSelectInstanceBatchMultiCrit until termination. The single evaluations will be written into the ArchiveBatchFSelect that resides in the FSelectInstanceBatchSingleCrit / FSelectInstanceBatchMultiCrit. The result will be written into the instance object.

Usage:

```
FSelectorBatch$optimize(inst)
```

Arguments:

`inst` (FSelectInstanceBatchSingleCrit | FSelectInstanceBatchMultiCrit).

Returns: `data.table::data.table()`.

`FSelectorBatch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatch$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`fselect_nested` *Function for Nested Resampling*

Description

Function to conduct nested resampling.

Usage

```
fselect_nested(
  fselector,
  task,
  learner,
  inner_resampling,
  outer_resampling,
  measure = NULL,
  term_evals = NULL,
  term_time = NULL,
  terminator = NULL,
  store_fselect_instance = TRUE,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features"
)
```

Arguments

`fselector` (FSelector)
Optimization algorithm.

`task` (mlr3::Task)
Task to operate on.

`learner` (mlr3::Learner)
Learner to optimize the feature subset for.

`inner_resampling` (mlr3::Resampling)
Resampling used for the inner loop.

`outer_resampling` mlr3::Resampling)
Resampling used for the outer loop.

`measure` (mlr3::Measure)
Measure to optimize. If `NULL`, default measure is used.

`term_evals` (`integer(1)`)
Number of allowed evaluations. Ignored if `terminator` is passed.

`term_time` (`integer(1)`)
Maximum allowed time in seconds. Ignored if `terminator` is passed.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_fselect_instance` (`logical(1)`)
If `TRUE` (default), stores the internally created FSelectInstanceBatchSingleCrit with all intermediate results in slot `$fselect_instance`. Is set to `TRUE`, if `store_models = TRUE`

`store_benchmark_result` (`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ties_method` (`character(1)`)
The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

Value

mlr3::ResampleResult

Examples

```
# Nested resampling on Palmer Penguins data set
rr = fselect_nested(
  fselector = fs("random_search"),
  task = tsk("penguins"),
  learner = lrn("classif.rpart"),
  inner_resampling = rsmp ("holdout"),
  outer_resampling = rsmp("cv", folds = 2),
  measure = msr("classif.ce"),
  term_evals = 4)

# Performance scores estimated on the outer resampling
rr$score()

# Unbiased performance of the final model trained on the full data set
rr$aggregate()
```

---

fsi *Syntactic Sugar for Feature Selection Instance Construction*

Description

Function to construct a FSelectInstanceBatchSingleCrit or FSelectInstanceBatchMultiCrit.

Usage

```
fsi(
  task,
  learner,
  resampling,
  measures = NULL,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features"
)
```

Arguments

task	(mlr3::Task) Task to operate on.
learner	(mlr3::Learner) Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (mlr3::Measure or list of mlr3::Measure)
A single measure creates a FSelectInstanceBatchSingleCrit and multiple measures a FSelectInstanceBatchMultiCrit. If `NULL`, default measure is used.

`terminator` (bbotk::Terminator)
Stop criterion of the feature selection.

`store_benchmark_result`
(`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ties_method` (`character(1)`)
The method to break ties when selecting sets while optimizing and when selecting the best set. Can be `"least_features"` or `"random"`. The option `"least_features"` (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The `random` method returns a random feature set from the best feature sets. Ignored if multiple measures are used.

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
`"classif"`	`"classif.ce"`	**mlr3**
`"regr"`	`"regr.mse"`	**mlr3**
`"surv"`	`"surv.cindex"`	**mlr3proba**
`"dens"`	`"dens.logloss"`	**mlr3proba**

"classif_st"	"classif.ce"	**mlr3spatial**
"regr_st"	"regr.mse"	**mlr3spatial**
"clust"	"clust.dunn"	**mlr3cluster**

Examples

```
# Feature selection on Palmer Penguins data set


task = tsk("penguins")
learner = lrn("classif.rpart")

# Construct feature selection instance
instance = fsi(
  task = task,
  learner = learner,
  resampling = rsmp("cv", folds = 3),
  measures = msr("classif.ce"),
  terminator = trm("evals", n_evals = 4)
)

# Choose optimization algorithm
fselector = fs("random_search", batch_size = 2)

# Run feature selection
fselector$optimize(instance)

# Subset task to optimal feature set
task$select(instance$result_feature_set)

# Train the learner with optimal feature set on the full data set
learner$train(task)

# Inspect all evaluated sets
as.data.table(instance$archive)
```

---

fsi_async *Syntactic Sugar for Asynchronous Feature Selection Instance Construction*

Description

Function to construct a FSelectInstanceAsyncSingleCrit or FSelectInstanceAsyncMultiCrit.

Usage

```
fsi_async(
  task,
```

```
  learner,
  resampling,
  measures = NULL,
  terminator,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  check_values = FALSE,
  callbacks = NULL,
  ties_method = "least_features",
  rush = NULL
)
```

Arguments

task	(mlr3::Task) Task to operate on.
learner	(mlr3::Learner) Learner to optimize the feature subset for.
resampling	(mlr3::Resampling) Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.
measures	(mlr3::Measure or list of mlr3::Measure) A single measure creates a FSelectInstanceAsyncSingleCrit and multiple measures a FSelectInstanceAsyncMultiCrit. If NULL, default measure is used.
terminator	(bbotk::Terminator) Stop criterion of the feature selection.
store_benchmark_result	(logical(1)) Store benchmark result in archive?
store_models	(logical(1)). Store models in benchmark result?
check_values	(logical(1)) Check the parameters before the evaluation and the results for validity?
callbacks	(list of CallbackBatchFSelect) List of callbacks.
ties_method	(character(1)) The method to break ties when selecting sets while optimizing and when selecting the best set. Can be "least_features" or "random". The option "least_features" (default) selects the feature set with the least features. If there are multiple best feature sets with the same number of features, one is selected randomly. The random method returns a random feature set from the best feature sets. Ignored if multiple measures are used.
rush	(Rush) If a rush instance is supplied, the optimization runs without batches.

Resources

There are several sections about feature selection in the mlr3book.

- Getting started with wrapper feature selection.
- Do a sequential forward selection Palmer Penguins data set.

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.
- Run a feature selection with Shadow Variable Search.

Default Measures

If no measure is passed, the default measure is used. The default measure depends on the task type.

Task	Default Measure	Package
"classif"	"classif.ce"	**mlr3**
"regr"	"regr.mse"	**mlr3**
"surv"	"surv.cindex"	**mlr3proba**
"dens"	"dens.logloss"	**mlr3proba**
"classif_st"	"classif.ce"	**mlr3spatial**
"regr_st"	"regr.mse"	**mlr3spatial**
"clust"	"clust.dunn"	**mlr3cluster**

Examples

```
# Feature selection on Palmer Penguins data set


task = tsk("penguins")
learner = lrn("classif.rpart")

# Construct feature selection instance
instance = fsi(
  task = task,
  learner = learner,
  resampling = rsmp("cv", folds = 3),
  measures = msr("classif.ce"),
  terminator = trm("evals", n_evals = 4)
)

# Choose optimization algorithm
fselector = fs("random_search", batch_size = 2)

# Run feature selection
fselector$optimize(instance)

# Subset task to optimal feature set
task$select(instance$result_feature_set)
```

```
# Train the learner with optimal feature set on the full data set
learner$train(task)

# Inspect all evaluated sets
as.data.table(instance$archive)
```

`mlr3fselect.async_freeze_archive` *Freeze Archive Callback*

Description

This CallbackAsyncFSelect freezes the ArchiveAsyncFSelect to ArchiveAsyncFSelectFrozen after the optimization has finished.

Examples

```
clbk("mlr3fselect.async_freeze_archive")
```

`mlr3fselect.backup` *Backup Benchmark Result Callback*

Description

This CallbackBatchFSelect writes the mlr3::BenchmarkResult after each batch to disk. The `path` must be set, e.g. `clbk("mlr3fselect.backup", path = "backup.rds")`. The benchmark result is written to a temporary file first and then renamed, so that a crash while writing does not destroy the backup of the previous batch.

Examples

```
clbk("mlr3fselect.backup", path = "backup.rds")

# Run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("random_search"),
  task = tsk("diabetes"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("holdout"),
  measures = msr("classif.ce"),
  term_evals = 4,
  callbacks = clbk("mlr3fselect.backup", path = tempfile(fileext = ".rds")))
```

mlr3fselect.internal_tuning

Internal Tuning Callback

Description

This callback runs internal tuning alongside the feature selection. The internal tuning values are aggregated and stored in the results. The final model is trained with the best feature set and the tuned value.

Examples

```
clbk("mlr3fselect.internal_tuning")
```

mlr3fselect.one_se_rule

One Standard Error Rule Callback

Description

Selects the smallest feature set within one standard error of the best as the result. If there are multiple such feature sets with the same number of features, the first one is selected. If the sets have exactly the same performance but different number of features, the one with the smallest number of features is selected.

Source

Kuhn, Max, Johnson, Kjell (2013). "Applied Predictive Modeling." In chapter Over-Fitting and Model Tuning, 61–92. Springer New York, New York, NY. ISBN 978-1-4614-6849-3.

Examples

```
clbk("mlr3fselect.one_se_rule")

# Run feature selection on the diabetes data set with the callback
instance = fselect(
  fselector = fs("random_search"),
  task = tsk("diabetes"),
  learner = lrn("classif.rpart"),
  resampling = rsmp ("cv", folds = 3),
  measures = msr("classif.ce"),
  term_evals = 10,
  callbacks = clbk("mlr3fselect.one_se_rule"))
# Smallest feature set within one standard error of the best
instance$result
```

`mlr3fselect.svm_rfe` *SVM-RFE Callback*

Description

Runs a recursive feature elimination with a mlr3learners::LearnerClassifSVM. The SVM must be configured with `type = "C-classification"` and `kernel = "linear"`.

Details

The importance score of a feature is the squared weight of the hyperplane that separates the two classes. The weights are only defined for a linear kernel and a single separating hyperplane, therefore the callback only works on binary classification tasks.

Source

Guyon I, Weston J, Barnhill S, Vapnik V (2002). "Gene Selection for Cancer Classification using Support Vector Machines." *Machine Learning*, **46**(1), 389–422. ISSN 1573-0565. doi:10.1023/A:1012487302797.

Examples

```
clbk("mlr3fselect.svm_rfe")

library(mlr3learners)

# Create instance with classification svm with linear kernel
instance = fsi(
  task = tsk("sonar"),
  learner = lrn("classif.svm", type = "C-classification", kernel = "linear"),
  resampling = rsmp("cv", folds = 3),
  measures = msr("classif.ce"),
  terminator = trm("none"),
  callbacks = clbk("mlr3fselect.svm_rfe"),
  store_models = TRUE
)

fselector = fs("rfe", feature_number = 5, n_features = 10)

# Run recursive feature elimination on the Sonar data set
fselector$optimize(instance)
```

`mlr_fselectors` *Dictionary of FSelectors*

Description

A mlr3misc::Dictionary storing objects of class FSelector. Each fselector has an associated help page, see `mlr_fselectors_[id]`.

For a more convenient way to retrieve and construct fselectors, see `fs()`/`fss()`.

Format

R6::R6Class object inheriting from mlr3misc::Dictionary.

Methods

See mlr3misc::Dictionary.

S3 methods

- `as.data.table(dict, ..., objects = FALSE)`
  mlr3misc::Dictionary -> `data.table::data.table()`
  Returns a `data.table::data.table()` with fields "key", "label", "properties" and "packages" as columns. If `objects` is set to `TRUE`, the constructed objects are returned in the list column named `object`.

See Also

Sugar functions: `fs()`, `fss()`

Other FSelector: `FSelector`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_genetic_search`, `mlr_fselectors_random_search`, `mlr_fselectors_rfe`, `mlr_fselectors_rfecv`, `mlr_fselectors_sequential`, `mlr_fselectors_shadow_variable_search`

Examples

```
as.data.table(mlr_fselectors)
mlr_fselectors$get("random_search")
fs("random_search")
```

`mlr_fselectors_async_design_points` *Feature Selection with Asynchronous Design Points*

Description

Subclass for asynchronous design points feature selection.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("async_design_points")
```

Parameters

`design` data.table::data.table
Design points to try in search, one per row.

Super classes

`FSelector` -> `FSelectorAsync` -> `FSelectorAsyncFromOptimizerAsync` -> `FSelectorAsyncDesignPoints`

Methods

Public methods:

- `FSelectorAsyncDesignPoints$new()`
- `FSelectorAsyncDesignPoints$clone()`

`FSelectorAsyncDesignPoints$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorAsyncDesignPoints$new()
```

`FSelectorAsyncDesignPoints$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorAsyncDesignPoints$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

See Also

Other FSelectorAsync: `mlr_fselectors_async_exhaustive_search`, `mlr_fselectors_async_random_search`

`mlr_fselectors_async_exhaustive_search` *Feature Selection with Asynchronous Exhaustive Search*

Description

Feature Selection using the Asynchronous Exhaustive Search Algorithm. Exhaustive Search generates all possible feature sets. The feature sets are evaluated asynchronously.

Details

The feature selection terminates itself when all feature sets are evaluated. It is not necessary to set a termination criterion.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("async_exhaustive_search")
```

Control Parameters

`max_features integer(1)`
: Maximum number of features. By default, number of features in mlr3::Task.

Super classes

`FSelector` -> `FSelectorAsync` -> `FSelectorAsyncExhaustiveSearch`

Methods

Public methods:

- `FSelectorAsyncExhaustiveSearch$new()`
- `FSelectorAsyncExhaustiveSearch$optimize()`
- `FSelectorAsyncExhaustiveSearch$clone()`

`FSelectorAsyncExhaustiveSearch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorAsyncExhaustiveSearch$new()
```

`FSelectorAsyncExhaustiveSearch$optimize()`: Starts the asynchronous optimization.

Usage:

```
FSelectorAsyncExhaustiveSearch$optimize(inst)
```

Arguments:

`inst` (FSelectInstanceAsyncSingleCrit | FSelectInstanceAsyncMultiCrit).

Returns: data.table::data.table.

`FSelectorAsyncExhaustiveSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorAsyncExhaustiveSearch$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

See Also

Other FSelectorAsync: `mlr_fselectors_async_design_points`, `mlr_fselectors_async_random_search`

`mlr_fselectors_async_random_search`
Feature Selection with Asynchronous Random Search

Description

Feature selection using Asynchronous Random Search Algorithm.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("async_random_search")
```

Control Parameters

`max_features integer(1)`
Maximum number of features. By default, number of features in mlr3::Task.

Super classes

`FSelector` -> `FSelectorAsync` -> `FSelectorAsyncRandomSearch`

Methods

Public methods:

- `FSelectorAsyncRandomSearch$new()`
- `FSelectorAsyncRandomSearch$clone()`

`FSelectorAsyncRandomSearch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorAsyncRandomSearch$new()
```

`FSelectorAsyncRandomSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorAsyncRandomSearch$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Source

Bergstra J, Bengio Y (2012). "Random Search for Hyper-Parameter Optimization." *Journal of Machine Learning Research*, **13**(10), 281–305. https://jmlr.csail.mit.edu/papers/v13/bergstra12a.html.

See Also

Other FSelectorAsync: `mlr_fselectors_async_design_points`, `mlr_fselectors_async_exhaustive_search`

---

mlr_fselectors_design_points

Feature Selection with Design Points

Description

Feature selection using user-defined feature sets.

Details

The feature sets are evaluated in order as given.

The feature selection terminates itself when all feature sets are evaluated. It is not necessary to set a termination criterion.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("design_points")
```

Parameters

`batch_size integer(1)`
: Maximum number of configurations to try in a batch.

`design` data.table::data.table
: Design points to try in search, one per row.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchFromOptimizerBatch` -> `FSelectorBatchDesignPoints`

Methods

Public methods:

- FSelectorBatchDesignPoints$new()
- FSelectorBatchDesignPoints$clone()

FSelectorBatchDesignPoints$new(): Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchDesignPoints$new()
```

FSelectorBatchDesignPoints$clone(): The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchDesignPoints$clone(deep = FALSE)
```

Arguments:

deep Whether to make a deep clone.

See Also

Other FSelector: FSelector, mlr_fselectors, mlr_fselectors_exhaustive_search, mlr_fselectors_genetic_searc
mlr_fselectors_random_search, mlr_fselectors_rfe, mlr_fselectors_rfecv, mlr_fselectors_sequential,
mlr_fselectors_shadow_variable_search

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("diabetes")
learner = lrn("classif.rpart")

# create design
design = mlr3misc::rowwise_table(
  ~age, ~glucose, ~insulin, ~mass, ~pedigree, ~pregnant, ~pressure, ~triceps,
  TRUE, FALSE,    TRUE,     TRUE,  FALSE,     TRUE,       FALSE,    TRUE,
  TRUE, TRUE,     FALSE,    TRUE,  FALSE,     TRUE,       FALSE,    FALSE,
  TRUE, FALSE,    TRUE,     TRUE,  FALSE,     TRUE,       FALSE,    FALSE,
  TRUE, FALSE,    TRUE,     TRUE,  FALSE,     TRUE,       TRUE,     TRUE
)

# run feature selection on the diabetes data set
instance = fselect(
  fselector = fs("design_points", design = design),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce")
)
```

```
# best performing feature set
instance$result

# all evaluated feature sets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

`mlr_fselectors_exhaustive_search`
Feature Selection with Exhaustive Search

Description

Feature Selection using the Exhaustive Search Algorithm. Exhaustive Search generates all possible feature sets.

Details

The feature selection terminates itself when all feature sets are evaluated. It is not necessary to set a termination criterion.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("exhaustive_search")
```

Control Parameters

`max_features integer(1)`
: Maximum number of features. By default, number of features in mlr3::Task.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchExhaustiveSearch`

Methods

Public methods:

- `FSelectorBatchExhaustiveSearch$new()`
- `FSelectorBatchExhaustiveSearch$clone()`

`FSelectorBatchExhaustiveSearch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchExhaustiveSearch$new()
```

`FSelectorBatchExhaustiveSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchExhaustiveSearch$clone(deep = FALSE)
```

Arguments:

deep Whether to make a deep clone.

See Also

Other FSelector: FSelector, mlr_fselectors, mlr_fselectors_design_points, mlr_fselectors_genetic_search, mlr_fselectors_random_search, mlr_fselectors_rfe, mlr_fselectors_rfecv, mlr_fselectors_sequential, mlr_fselectors_shadow_variable_search

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("exhaustive_search"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  term_evals = 10
)

# best performing feature set
instance$result

# all evaluated feature sets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

`mlr_fselectors_genetic_search` *Feature Selection with Genetic Search*

Description

Feature selection using the Genetic Algorithm from the package **genalg**.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("genetic_search")
```

Control Parameters

For the meaning of the control parameters, see `genalg::rbga.bin()`. `genalg::rbga.bin()` internally terminates after `iters` iteration. We set `iters = 100000` to allow the termination via our terminators. If more iterations are needed, set `iters` to a higher value in the parameter set.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchGeneticSearch`

Methods

Public methods:

- `FSelectorBatchGeneticSearch$new()`
- `FSelectorBatchGeneticSearch$clone()`

`FSelectorBatchGeneticSearch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchGeneticSearch$new()
```

`FSelectorBatchGeneticSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchGeneticSearch$clone(deep = FALSE)
```

Arguments:

deep Whether to make a deep clone.

See Also

Other FSelector: `FSelector`, `mlr_fselectors`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_random_search`, `mlr_fselectors_rfe`, `mlr_fselectors_rfecv`, `mlr_fselectors_sequential`, `mlr_fselectors_shadow_variable_search`

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("genetic_search"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  term_evals = 10
)

# best performing feature set
instance$result

# all evaluated feature sets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

---

mlr_fselectors_random_search

Feature Selection with Random Search

Description

Feature selection using Random Search Algorithm.

Details

The feature sets are randomly drawn. The sets are evaluated in batches of size `batch_size`. Larger batches mean we can parallelize more, smaller batches imply a more fine-grained checking of termination criteria.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("random_search")
```

Control Parameters

`max_features integer(1)`
: Maximum number of features. By default, number of features in mlr3::Task.

`batch_size integer(1)`
: Maximum number of feature sets to try in a batch.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchRandomSearch`

Methods

Public methods:

- `FSelectorBatchRandomSearch$new()`
- `FSelectorBatchRandomSearch$clone()`

`FSelectorBatchRandomSearch$new()`: Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchRandomSearch$new()
```

`FSelectorBatchRandomSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchRandomSearch$clone(deep = FALSE)
```

Arguments:

deep Whether to make a deep clone.

Source

Bergstra J, Bengio Y (2012). "Random Search for Hyper-Parameter Optimization." *Journal of Machine Learning Research*, **13**(10), 281–305. https://jmlr.csail.mit.edu/papers/v13/bergstra12a.html.

See Also

Other FSelector: `FSelector`, `mlr_fselectors`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_genetic_search`, `mlr_fselectors_rfe`, `mlr_fselectors_rfecv`, `mlr_fselectors_sequential`, `mlr_fselectors_shadow_variable_search`

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")
```

```
# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("random_search"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  term_evals = 10
)

# best performing feature subset
instance$result

# all evaluated feature subsets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

`mlr_fselectors_rfe` *Feature Selection with Recursive Feature Elimination*

Description

Feature selection using the Recursive Feature Elimination (RFE) algorithm. Recursive feature elimination iteratively removes features with a low importance score. Only works with mlr3::Learners that can calculate importance scores (see the section on optional extractors in mlr3::Learner).

Details

The learner is trained on all features at the start and importance scores are calculated for each feature. Then the least important feature is removed and the learner is trained on the reduced feature set. The importance scores are calculated again and the procedure is repeated until the desired number of features is reached. The non-recursive option (`recursive = FALSE`) only uses the importance scores calculated in the first iteration.

The feature selection terminates itself when `n_features` is reached. It is not necessary to set a termination criterion.

When using a cross-validation resampling strategy, the importance scores of the resampling iterations are aggregated. The parameter `aggregation` determines how the importance scores are aggregated. By default (`"rank"`), the importance score vector of each fold is ranked and the feature with the lowest average rank is removed. The option `"mean"` averages the score of each feature across the resampling iterations and removes the feature with the lowest average score. Averaging the scores is not appropriate for most importance measures.

Archive

The ArchiveBatchFSelect holds the following additional columns:

- `"importance"` (`numeric()`)
  The importance score vector of the feature subset.

Resources

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("rfe")
```

Control Parameters

`n_features integer(1)`
: The minimum number of features to select, by default half of the features.

`feature_fraction double(1)`
: Fraction of features to retain in each iteration. The default of 0.5 retains half of the features.

`feature_number integer(1)`
: Number of features to remove in each iteration.

`subset_sizes integer()`
: Vector of the number of features to retain in each iteration. Must be sorted in decreasing order.

`recursive logical(1)`
: If `TRUE` (default), the feature importance is calculated in each iteration.

`aggregation character(1)`
: The aggregation method for the importance scores of the resampling iterations. See details.

The parameter `feature_fraction`, `feature_number` and `subset_sizes` are mutually exclusive.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchRFE`

Methods

Public methods:

- `FSelectorBatchRFE$new()`
- `FSelectorBatchRFE$clone()`

`FSelectorBatchRFE$new()`**:** Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchRFE$new()
```

`FSelectorBatchRFE$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchRFE$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Source

Guyon I, Weston J, Barnhill S, Vapnik V (2002). "Gene Selection for Cancer Classification using Support Vector Machines." *Machine Learning*, **46**(1), 389–422. ISSN 1573-0565. doi:10.1023/A:1012487302797.

See Also

Other FSelector: FSelector, mlr_fselectors, mlr_fselectors_design_points, mlr_fselectors_exhaustive_search, mlr_fselectors_genetic_search, mlr_fselectors_random_search, mlr_fselectors_rfecv, mlr_fselectors_sequential, mlr_fselectors_shadow_variable_search

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("rfe"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  store_models = TRUE
)

# best performing feature subset
instance$result

# all evaluated feature subsets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

`mlr_fselectors_rfecv` *Feature Selection with Recursive Feature Elimination with Cross Validation*

Description

Feature selection using the Recursive Feature Elimination with Cross-Validation (RFE-CV) algorithm. See FSelectorBatchRFE for a description of the base algorithm. RFE-CV runs a recursive feature elimination in each iteration of a cross-validation to determine the optimal number of features. Then a recursive feature elimination is run again on the complete dataset with the optimal number of features as the final feature set size. The performance of the optimal feature set is calculated on the complete data set and should not be reported as the performance of the final model. Only works with mlr3::Learners that can calculate importance scores (see the section on optional extractors in mlr3::Learner).

Details

The resampling strategy is changed during the feature selection. The resampling strategy passed to the instance (`resampling`) is used to determine the optimal number of features. Usually, a cross-validation strategy is used and a recursive feature elimination is run in each iteration of the cross-validation. Internally, mlr3::ResamplingCustom is used to emulate this part of the algorithm. In the final recursive feature elimination run the resampling strategy is changed to mlr3::ResamplingInsample i.e. the complete data set is used for training and testing.

The feature selection terminates itself when the optimal number of features is reached. It is not necessary to set a termination criterion.

Archive

The ArchiveBatchFSelect holds the following additional columns:

- `"iteration"` (`integer(1)`)
  The resampling iteration in which the feature subset was evaluated.
- `"importance"` (`numeric()`)
  The importance score vector of the feature subset.

Resources

The gallery features a collection of case studies and demos about optimization.

- Utilize the built-in feature importance of models with Recursive Feature Elimination.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("rfecv")
```

Control Parameters

`n_features integer(1)`
: The number of features to select. By default half of the features are selected.

`feature_fraction double(1)`
: Fraction of features to retain in each iteration. The default 0.5 retrains half of the features.

`feature_number integer(1)`
: Number of features to remove in each iteration.

`subset_sizes integer()`
: Vector of number of features to retain in each iteration. Must be sorted in decreasing order.

`recursive logical(1)`
: If `TRUE` (default), the feature importance is calculated in each iteration.

The parameter `feature_fraction`, `feature_number` and `subset_sizes` are mutually exclusive.

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchRFECV`

Methods

Public methods:

- `FSelectorBatchRFECV$new()`
- `FSelectorBatchRFECV$clone()`

`FSelectorBatchRFECV$new()`**:** Creates a new instance of this R6 class.

Usage:

```
FSelectorBatchRFECV$new()
```

`FSelectorBatchRFECV$clone()`**:** The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchRFECV$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

See Also

Other FSelector: `FSelector`, `mlr_fselectors`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_genetic_search`, `mlr_fselectors_random_search`, `mlr_fselectors_rfe`, `mlr_fselectors_sequent`, `mlr_fselectors_shadow_variable_search`

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
```

```
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("rfecv"),
  task = task,
  learner = learner,
  resampling = rsmp("cv", folds = 3),
  measure = msr("classif.ce"),
  store_models = TRUE
)

# best performing feature subset
instance$result

# all evaluated feature subsets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

`mlr_fselectors_sequential`

Feature Selection with Sequential Search

Description

Feature selection using Sequential Search Algorithm.

Details

Sequential forward selection (`strategy = fsf`) extends the feature set in each iteration with the feature that increases the model's performance the most. Sequential backward selection (`strategy = fsb`) follows the same idea but starts with all features and removes features from the set.

The feature selection terminates itself when `min_features` or `max_features` is reached. It is not necessary to set a termination criterion.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("sequential")
```

Control Parameters

`min_features integer(1)`
: Minimum number of features. By default, 1.

`max_features integer(1)`
: Maximum number of features. By default, number of features in mlr3::Task.

`strategy character(1)`
: Search method `sfs` (forward search) or `sbs` (backward search).

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchSequential`

Methods

Public methods:

- `FSelectorBatchSequential$new()`
- `FSelectorBatchSequential$optimization_path()`
- `FSelectorBatchSequential$clone()`

`FSelectorBatchSequential$new()`: Creates a new instance of this R6 class.‘

Usage:

```
FSelectorBatchSequential$new()
```

`FSelectorBatchSequential$optimization_path()`: Returns the optimization path.

Usage:

```
FSelectorBatchSequential$optimization_path(inst, include_uhash = FALSE)
```

Arguments:

`inst` (FSelectInstanceBatchSingleCrit)
: Instance optimized with FSelectorBatchSequential.

`include_uhash (logical(1))`
: Include uhash column?

Returns: `data.table::data.table()`

`FSelectorBatchSequential$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchSequential$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

See Also

Other FSelector: `FSelector`, `mlr_fselectors`, `mlr_fselectors_design_points`, `mlr_fselectors_exhaustive_search`, `mlr_fselectors_genetic_search`, `mlr_fselectors_random_search`, `mlr_fselectors_rfe`, `mlr_fselectors_rfecv`, `mlr_fselectors_shadow_variable_search`

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("sequential"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
  term_evals = 10
)

# best performing feature set
instance$result

# all evaluated feature sets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

---

mlr_fselectors_shadow_variable_search

Feature Selection with Shadow Variable Search

Description

Feature selection using the Shadow Variable Search Algorithm. Shadow variable search creates for each feature a permuted copy and stops when one of them is selected.

Details

The feature selection terminates itself when the first shadow variable is selected. It is not necessary to set a termination criterion.

Resources

The gallery features a collection of case studies and demos about optimization.

- Run a feature selection with Shadow Variable Search.

Dictionary

This FSelector can be instantiated with the associated sugar function `fs()`:

```
fs("shadow_variable_search")
```

Super classes

`FSelector` -> `FSelectorBatch` -> `FSelectorBatchShadowVariableSearch`

Methods

Public methods:

- `FSelectorBatchShadowVariableSearch$new()`
- `FSelectorBatchShadowVariableSearch$optimization_path()`
- `FSelectorBatchShadowVariableSearch$clone()`

`FSelectorBatchShadowVariableSearch$new()`: Creates a new instance of this R6 class.‘

Usage:

```
FSelectorBatchShadowVariableSearch$new()
```

`FSelectorBatchShadowVariableSearch$optimization_path()`: Returns the optimization path.

Usage:

```
FSelectorBatchShadowVariableSearch$optimization_path(inst)
```

Arguments:

`inst` (FSelectInstanceBatchSingleCrit)
Instance optimized with FSelectorBatchShadowVariableSearch.

Returns: data.table::data.table

`FSelectorBatchShadowVariableSearch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
FSelectorBatchShadowVariableSearch$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Source

Thomas J, Hepp T, Mayr A, Bischl B (2017). “Probing for Sparse and Fast Variable Selection with Model-Based Boosting.” *Computational and Mathematical Methods in Medicine*, **2017**, 1–8. doi:10.1155/2017/1421409.

Wu Y, Boos DD, Stefanski LA (2007). “Controlling Variable Selection by the Addition of Pseudovariables.” *Journal of the American Statistical Association*, **102**(477), 235–243. doi:10.1198/016214506000000843.

See Also

Other FSelector: FSelector, mlr_fselectors, mlr_fselectors_design_points, mlr_fselectors_exhaustive_search mlr_fselectors_genetic_search, mlr_fselectors_random_search, mlr_fselectors_rfe, mlr_fselectors_rfecv, mlr_fselectors_sequential

Examples

```
# Feature Selection


# retrieve task and load learner
task = tsk("penguins")
learner = lrn("classif.rpart")

# run feature selection on the Palmer Penguins data set
instance = fselect(
  fselector = fs("shadow_variable_search"),
  task = task,
  learner = learner,
  resampling = rsmp("holdout"),
  measure = msr("classif.ce"),
)

# best performing feature subset
instance$result

# all evaluated feature subsets
as.data.table(instance$archive)

# subset the task and fit the final model
task$select(instance$result_feature_set)
learner$train(task)
```

ObjectiveFSelect *Class for Feature Selection Objective*

Description

Stores the objective function that estimates the performance of feature subsets. This class is usually constructed internally by the FSelectInstanceBatchSingleCrit or FSelectInstanceBatchMultiCrit.

Super class

bbotk::Objective -> ObjectiveFSelect

Public fields

`task` (mlr3::Task).

`learner` (mlr3::Learner).

`resampling` (mlr3::Resampling).

`measures` (list of mlr3::Measure).

`store_models` (`logical(1)`).

`store_benchmark_result` (`logical(1)`).

`callbacks` (List of CallbackBatchFSelects).

Methods

Public methods:

- `ObjectiveFSelect$new()`
- `ObjectiveFSelect$clone()`

`ObjectiveFSelect$new()`: Creates a new instance of this R6 class.

Usage:

```
ObjectiveFSelect$new(
  task,
  learner,
  resampling,
  measures,
  check_values = TRUE,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  callbacks = NULL
)
```

Arguments:

`task` (mlr3::Task)
: Task to operate on.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (list of mlr3::Measure)
: Measures to optimize. If `NULL`, **mlr3**'s default measure is used.

`check_values` (`logical(1)`)
: Check the parameters before the evaluation and the results for validity?

`store_benchmark_result` (`logical(1)`)
: Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ObjectiveFSelect$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ObjectiveFSelect$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`ObjectiveFSelectAsync` *Class for Feature Selection Objective*

Description

Stores the objective function that estimates the performance of feature subsets. This class is usually constructed internally by the FSelectInstanceAsyncSingleCrit or FSelectInstanceAsyncMultiCrit.

Super classes

`bbotk::Objective` -> `ObjectiveFSelect` -> `ObjectiveFSelectAsync`

Methods

Public methods:

- `ObjectiveFSelectAsync$clone()`

`ObjectiveFSelectAsync$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ObjectiveFSelectAsync$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

`ObjectiveFSelectBatch` *Class for Feature Selection Objective*

Description

Stores the objective function that estimates the performance of feature subsets. This class is usually constructed internally by the FSelectInstanceBatchSingleCrit or FSelectInstanceBatchMultiCrit.

Super classes

`bbotk::Objective` -> `ObjectiveFSelect` -> `ObjectiveFSelectBatch`

Public fields

`archive` (ArchiveBatchFSelect).

Methods

Public methods:

- `ObjectiveFSelectBatch$new()`
- `ObjectiveFSelectBatch$clone()`

`ObjectiveFSelectBatch$new()`: Creates a new instance of this R6 class.

Usage:

```
ObjectiveFSelectBatch$new(
  task,
  learner,
  resampling,
  measures,
  check_values = TRUE,
  store_benchmark_result = TRUE,
  store_models = FALSE,
  archive = NULL,
  callbacks = NULL
)
```

Arguments:

`task` (mlr3::Task)
: Task to operate on.

`learner` (mlr3::Learner)
: Learner to optimize the feature subset for.

`resampling` (mlr3::Resampling)
: Resampling that is used to evaluated the performance of the feature subsets. Uninstantiated resamplings are instantiated during construction so that all feature subsets are evaluated on the same data splits. Already instantiated resamplings are kept unchanged.

`measures` (list of mlr3::Measure)
: Measures to optimize. If `NULL`, **mlr3**'s default measure is used.

`check_values` (`logical(1)`)
Check the parameters before the evaluation and the results for validity?

`store_benchmark_result` (`logical(1)`)
Store benchmark result in archive?

`store_models` (`logical(1)`). Store models in benchmark result?

`archive` (ArchiveBatchFSelect)
Reference to the archive of FSelectInstanceBatchSingleCrit | FSelectInstanceBatchMultiCrit. If `NULL` (default), benchmark result and models cannot be stored.

`callbacks` (list of CallbackBatchFSelect)
List of callbacks.

`ObjectiveFSelectBatch$clone()`: The objects of this class are cloneable with this method.

Usage:

```
ObjectiveFSelectBatch$clone(deep = FALSE)
```

Arguments:

`deep` Whether to make a deep clone.

Index

∗ **Dictionary**
mlr_fselectors, 79
∗ **FSelectorAsync**
mlr_fselectors_async_design_points, 80
mlr_fselectors_async_exhaustive_search, 81
mlr_fselectors_async_random_search, 82
∗ **FSelector**
FSelector, 63
mlr_fselectors, 79
mlr_fselectors_design_points, 83
mlr_fselectors_exhaustive_search, 85
mlr_fselectors_genetic_search, 87
mlr_fselectors_random_search, 88
mlr_fselectors_rfe, 90
mlr_fselectors_rfecv, 93
mlr_fselectors_sequential, 95
mlr_fselectors_shadow_variable_search, 97

Archive, *59*
ArchiveAsyncFSelect, *4*, 4, *5*, *8, 9*, *48*, *51*, *54*, *66*, *76*
ArchiveAsyncFSelectFrozen, *8*, 8, *76*
ArchiveBatchFSelect, *10*, 10, *11*, *48*, *69*, *91*, *93*, *102, 103*
assert_async_fselect_callback, 14
assert_async_fselect_callbacks *(*assert_async_fselect_callback*)*, 14
auto_fselector, 20, *35*
auto_fselector(), *15*, *20*
AutoFSelector, *15*, 15, *16*, *20–22*, *28*, *30*, *32*, *44, 45*

bbotk::Archive, *5*, *8*, *11*
bbotk::ArchiveAsync, *5*, *8*
bbotk::ArchiveAsyncFrozen, *8*
bbotk::ArchiveBatch, *11*
bbotk::CallbackAsync, *22*
bbotk::CallbackBatch, *23*
bbotk::Codomain, *12*
bbotk::ContextAsync, *30*
bbotk::ContextBatch, *28*, *31*
bbotk::EvalInstance, *52*, *54*, *57*, *60*
bbotk::Objective, *99*, *101, 102*
bbotk::OptimInstance, *52*, *54*, *57*, *60*
bbotk::OptimInstanceAsync, *52*, *54*
bbotk::OptimInstanceAsyncMultiCrit, *52*
bbotk::OptimInstanceAsyncSingleCrit, *54*
bbotk::OptimInstanceBatch, *57*, *60*
bbotk::OptimInstanceBatchMultiCrit, *57*
bbotk::OptimInstanceBatchSingleCrit, *60*
bbotk::Terminator, *15*, *17*, *21*, *35*, *47*, *49*, *53*, *55*, *58, 59*, *61*, *67*, *70*, *72*, *74*
bbotk::TerminatorCombo, *49*

c(), *39*
callback_async_fselect, 24
callback_async_fselect(), *22*, *30*
callback_batch_fselect, 28
callback_batch_fselect(), *23*, *31*
CallbackAsyncFSelect, *14, 15*, *22*, 22, *25*, *30*, *76*
CallbackBatchFSelect, *17*, *21*, *23*, 23, *28*, *31*, *35*, *49*, *53*, *55*, *58*, *61*, *70*, *72*, *74*, *76*, *100, 101*, *103*
clbk(), *22, 23*, *25*, *28*
ContextAsyncFSelect, *25*, *28*, 30
ContextBatchFSelect, *28*, *31*, 31

data.table::data.table, *16*, *32*, *37, 38*, *40*, *42*, *80*, *82, 83*, *98*
data.table::data.table(), *5*, *7, 8*, *10, 11*, *13, 14*, *37*, *44–47*, *66*, *69*, *79*, *96*

dictionary, *22, 23*, *25*, *28*, *47*

embedded_ensemble_fselect, 32
embedded_ensemble_fselect(), *36*
ensemble_fs_result, 36
ensemble_fselect, 34
ensemble_fselect(), *36*
EnsembleFSResult, *32–37*, *39, 40*
EnsembleFSResult *(*ensemble_fs_result*)*, 36
extract_inner_fselect_archives, 44
extract_inner_fselect_results, 45

faggregate, 46
fastVoteR::rank_candidates(), *40*
fs, 47
fs(), *79–83*, *85*, *87, 88*, *91*, *93*, *95*, *98*
fselect, 48
fselect(), *51*, *54*, *56*, *59*
fselect_nested, 69
FSelectInstanceAsyncMultiCrit, *51*, 51, *66*, *73, 74*, *81*, *101*
FSelectInstanceAsyncSingleCrit, *54*, 54, *66*, *73, 74*, *81*, *101*
FSelectInstanceBatchMultiCrit, *13*, *48–50*, *53*, *56*, 56, *58*, *62*, *67*, *69*, *71, 72*, *99*, *102, 103*
FSelectInstanceBatchSingleCrit, *16, 17*, *21*, *48–50*, *59*, 59, *67*, *69–72*, *96*, *98, 99*, *102, 103*
FSelector, *15–17*, *20, 21*, *34, 35*, *47, 48*, *56*, *58, 59*, *62*, 63, *66*, *68*, *70*, *79–89*, *91–96*, *98, 99*
FSelectorAsync, *51*, *53, 54*, *56*, *66*, 66, *80–82*
FSelectorAsyncDesignPoints *(*mlr_fselectors_async_design_points*)*, 80
FSelectorAsyncExhaustiveSearch *(*mlr_fselectors_async_exhaustive_search*)*, 81
FSelectorAsyncFromOptimizerAsync, *80*
FSelectorAsyncRandomSearch *(*mlr_fselectors_async_random_search*)*, 82
FSelectorBatch, *67*, 67, *83*, *85*, *87*, *89*, *91*, *94*, *96*, *98*
FSelectorBatchDesignPoints *(*mlr_fselectors_design_points*)*, 83
FSelectorBatchExhaustiveSearch *(*mlr_fselectors_exhaustive_search*)*, 85
FSelectorBatchFromOptimizerBatch, *83*
FSelectorBatchGeneticSearch *(*mlr_fselectors_genetic_search*)*, 87
FSelectorBatchRandomSearch *(*mlr_fselectors_random_search*)*, 88
FSelectorBatchRFE, *93*
FSelectorBatchRFE *(*mlr_fselectors_rfe*)*, 90
FSelectorBatchRFECV *(*mlr_fselectors_rfecv*)*, 93
FSelectorBatchSequential, *96*
FSelectorBatchSequential *(*mlr_fselectors_sequential*)*, 95
FSelectorBatchShadowVariableSearch, *98*
FSelectorBatchShadowVariableSearch *(*mlr_fselectors_shadow_variable_search*)*, 97
FSelectors, *47*
fsi, 71
fsi(), *56*, *59*
fsi_async, 73
fsi_async(), *51*, *54*
fss *(*fs*)*, 47
fss(), *79*

genalg::rbga.bin(), *87*

measure, *39*
mlr3::benchmark(), *15*, *22*, *59*, *67*
mlr3::BenchmarkResult, *5*, *8*, *10, 11*, *31, 32*, *37, 38*, *44–47*, *76*
mlr3::callback_resample(), *26, 27*
mlr3::ContextResample, *28*
mlr3::Learner, *6*, *9*, *13*, *15–18*, *20, 21*, *33*, *35*, *48, 49*, *52*, *55*, *57*, *61*, *70, 71*, *74*, *90*, *93*, *100*, *102*
mlr3::Measure, *5*, *11, 12*, *15*, *17*, *20, 21*, *33*, *35*, *37, 38*, *47*, *49*, *53*, *55*, *58*, *61*, *70*, *72*, *74*, *100*, *102*
mlr3::Prediction, *6*, *9*, *13*
mlr3::resample(), *15*, *22*
mlr3::ResampleResult, *4, 5*, *7*, *10, 11*, *13*, *44–47*, *50, 51*, *54*, *57*, *60*, *70*

mlr3::Resampling, *15*, *17*, *20–22*, *33*, *35*, *49*, *52*, *55*, *58*, *61*, *70*, *72*, *74*, *100*, *102*
mlr3::ResamplingBootstrap, *33*, *35*
mlr3::ResamplingCustom, *93*
mlr3::ResamplingInsample, *93*
mlr3::ResamplingSubsampling, *33*, *35*
mlr3::Task, *6*, *12*, *33*, *34*, *49*, *52*, *55*, *57*, *61*, *70*, *71*, *74*, *81*, *82*, *85*, *89*, *96*, *100*, *102*
mlr3fselect *(*mlr3fselect-package*)*, 3
mlr3fselect-package, 3
mlr3fselect.async_freeze_archive, *8*, 76
mlr3fselect.backup, 76
mlr3fselect.internal_tuning, 77
mlr3fselect.one_se_rule, 77
mlr3fselect.svm_rfe, 78
mlr3learners::LearnerClassifSVM, *78*
mlr3misc::Callback, *22*, *23*
mlr3misc::Context, *30*, *31*
mlr3misc::Dictionary, *47*, *79*
mlr3misc::dictionary_sugar_get(), *47*
mlr_callbacks, *22*, *23*, *25*, *28*
mlr_fselectors, *47*, *65*, 79, *84*, *86*, *87*, *89*, *92*, *94*, *96*, *99*
mlr_fselectors_async_design_points, 80, *82*, *83*
mlr_fselectors_async_exhaustive_search, *80*, 81, *83*
mlr_fselectors_async_random_search, *80*, *82*, 82
mlr_fselectors_design_points, *65*, *79*, 83, *86*, *87*, *89*, *92*, *94*, *96*, *99*
mlr_fselectors_exhaustive_search, *65*, *79*, *84*, 85, *87*, *89*, *92*, *94*, *96*, *99*
mlr_fselectors_genetic_search, *65*, *79*, *84*, *86*, 87, *89*, *92*, *94*, *96*, *99*
mlr_fselectors_random_search, *65*, *79*, *84*, *86*, *87*, 88, *92*, *94*, *96*, *99*
mlr_fselectors_rfe, *65*, *79*, *84*, *86*, *87*, *89*, 90, *94*, *96*, *99*
mlr_fselectors_rfecv, *65*, *79*, *84*, *86*, *87*, *89*, *92*, 93, *96*, *99*
mlr_fselectors_sequential, *65*, *79*, *84*, *86*, *87*, *89*, *92*, *94*, 95, *99*
mlr_fselectors_shadow_variable_search, *65*, *79*, *84*, *86*, *87*, *89*, *92*, *94*, *96*, 97
mlr_reflections$fselect_properties, *64*, *68*
mlr_terminators, *47*

ObjectiveFSelect, 99, *101*, *102*
ObjectiveFSelectAsync, 101
ObjectiveFSelectBatch, *59*, 102

paradox::ParamSet, *6*, *12*, *64*, *68*

R6, *5*, *9*, *12*, *16*, *38*, *52*, *55*, *57*, *61*, *64*, *68*, *80–82*, *84*, *85*, *87*, *89*, *91*, *94*, *96*, *98*, *100*, *102*
R6::R6Class, *48*, *79*
requireNamespace(), *64*, *68*
rush::Rush, *4*

stabm::listStabilityMeasures(), *41*

Terminators, *47*, *49*