skrobot

Release 1.0.8

Medoid AI

Nov 29, 2020

CONTENTS

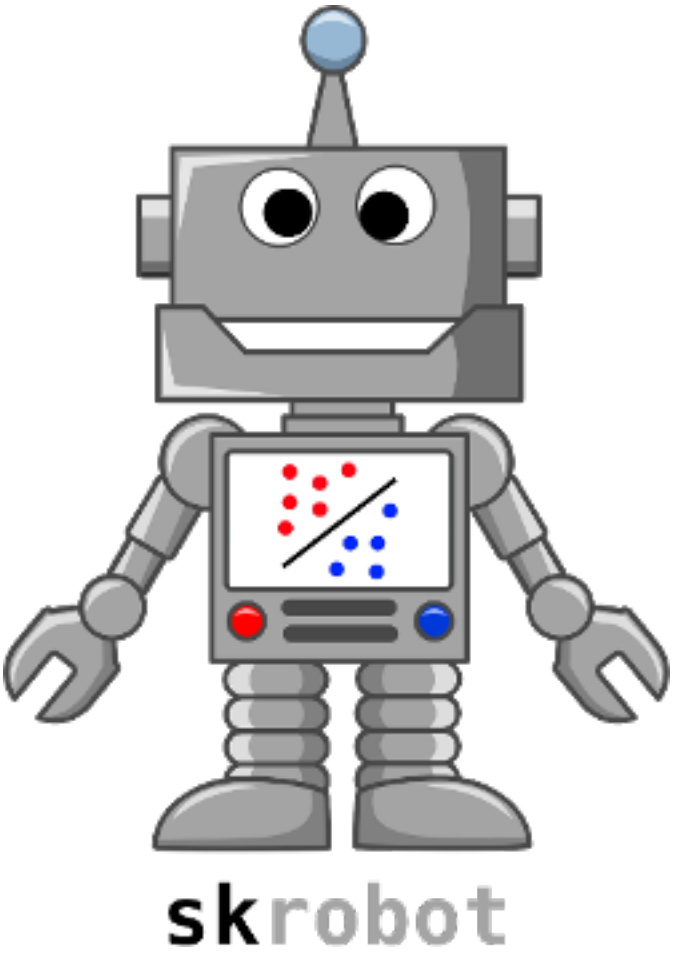
skrobot

CHAPTER

ONE

API REFERENCE

1.1 skrobot package

1.1.1 Subpackages

skrobot.core package

Submodules

skrobot.core.experiment module

class skrobot.core.experiment.**Experiment**(*experiments_repository*)
Bases: object

The *Experiment* class can be used to build, track and run an experiment.

It can run *BaseTask* tasks in the context of an experiment.

When building an experiment and/or running tasks, various metadata as well as task-related files are stored for tracking experiments.

Lastly, an experiment can be configured to send notifications when running a task, which can be useful for teams who need to get notified for the progress of the experiment.

__init__(*experiments_repository*)
This is the constructor method and can be used to create a new object instance of *Experiment* class.

Parameters **experiments_repository** (*str*) – The root directory path under which a unique directory is created for the experiment.

set_notifier(*notifier:* skrobot.notification.base_notifier.BaseNotifier)
Optional method.

Set the experiment's notifier.

Parameters **notifier** (*BaseNotifier*) – The experiment's notifier.

Returns The object instance itself.

Return type *Experiment*

set_source_code_file_path(*source_code_file_path*)
Optional method.

Set the experiment's source code file path.

Parameters **source_code_file_path** (*str*) – The experiment's source code file path.

Returns The object instance itself.

Return type *Experiment*

set_experimenter(*experimenter*)
Optional method.

Set the experimenter's name.

By default the experimenter's name is *anonymous*. However, if you want to override it you can pass a new name.

Parameters **experimenter** (*str*) – The experimenter's name.

Returns The object instance itself.

Return type *Experiment*

build()
Build the *Experiment*.

When an experiment is built, it creates a unique directory under which it stores various experiment-related metadata and files for tracking reasons.

Specifically, under the experiment's directory an *experiment.log* JSON file is created, which contains a unique auto-generated experiment ID, the current date & time, and the experimenter's name.

Also, the experiment's directory name contains the experimenter's name as well as current date & time.

Lastly, in case *set_source_code_file_path()* is used, the experiment's source code file is copied also under the experiment's directory.

Returns The object instance itself.

Return type *Experiment*

run(*task*)
Run a *BaseTask* task.

When running a task, its recorded parameters (e.g., *train_task.params*) and any other task-related generated files are stored under experiment's directory for tracking reasons.

The task's recorded parameters are in JSON format.

Also, in case *set_notifier()* is used to set a notifier, a notification is sent for the success or failure (including the error message) of the task's execution.

Lastly, in case an exception occurs, a text file (e.g., *train_task.errors*) is generated under experiment's directory containing the error message.

Parameters **task** (*BaseTask*) – The task to run.

Returns The task's result.

Return type Depends on the `task` parameter.

skrobot.core.task_runner module

class skrobot.core.task_runner.**TaskRunner**(*output_directory_path*)
Bases: object

The *TaskRunner* class is a simplified version (in functionality) of the *Experiment* class.

It leaves out all the "experiment" stuff and is focused mostly in the execution and tracking of *BaseTask* tasks.

__init__(*output_directory_path*)
This is the constructor method and can be used to create a new object instance of *TaskRunner* class.

Parameters **output_directory_path** (*str*) – The output directory path under which task-related generated files are stored.

run(*task*)
Run a *BaseTask* task.

When running a task, its recorded parameters (e.g., *train_task.params*) and any other task-related generated files are stored under output directory for tracking reasons.

The task's recorded parameters are in JSON format.

Lastly, in case an exception occurs, a text file (e.g., *train_task.errors*) is generated under output directory containing the error message.

Parameters **task** (*BaseTask*) – The task to run.

Returns The task's result.

Return type Depends on the task parameter.

skrobot.feature_selection package

Submodules

skrobot.feature_selection.column_selector module

class skrobot.feature_selection.column_selector.**ColumnSelector**(*cols*, *drop_axis=False*)
Bases: sklearn.base.BaseEstimator

The *ColumnSelector* class is an implementation of a column selector for scikit-learn pipelines.

It can be used for manual feature selection to select specific columns from an input data set.

It can select columns either by integer indices or by names.

__init__(*cols*, *drop_axis=False*)
This is the constructor method and can be used to create a new object instance of *ColumnSelector* class.

Parameters

- **cols** (*list*) – A non-empty list specifying the columns to be selected. For example, [1, 4, 5] to select the 2nd, 5th, and 6th columns, and ['A','C','D'] to select the columns A, C and D.
- **drop_axis** (*bool, optional*) – Can be used to reshape the output data set from (n_samples, 1) to (n_samples) by dropping the last axis. It defaults to False.

fit_transform(*X*, *y=None*)
Returns a slice of the input data set.

Parameters

- **X** (*{NumPy array, pandas DataFrame, SciPy sparse matrix}*) – Input vectors of shape (n_samples, n_features), where n_samples is the number of samples and n_features is the number of features.
- **y** (*None*) – Ignored.

Returns Subset of the input data set of shape (n_samples, k_features), where n_samples is the number of samples and k_features <= n_features.

Return type {NumPy array, SciPy sparse matrix}

transform(*X*, *y=None*)
Returns a slice of the input data set.

Parameters

- **X** (*{NumPy array, pandas DataFrame, SciPy sparse matrix}*) – Input vectors of shape (n_samples, n_features), where n_samples is the number of samples and n_features is the number of features.
- **y** (*None*) – Ignored.

Returns Subset of the input data set of shape (n_samples, k_features), where n_samples is the number of samples and k_features <= n_features.

Return type {NumPy array, SciPy sparse matrix}

fit(*X*, *y=None*)
This is a mock method and does nothing.

Parameters

- **X** (*None*) – Ignored.
- **y** (*None*) – Ignored.

Returns The object instance itself.

Return type *ColumnSelector*

skrobot.notification package

Submodules

skrobot.notification.base_notifier module

class skrobot.notification.base_notifier.**BaseNotifier**
Bases: abc.ABC

The *BaseNotifier* is an abstract base class for implementing notifiers.

A notifier can be used to send notifications.

abstract notify(*message*)
An abstract method for sending the notification.

Parameters **message** (*str*) – The notification's message.

skrobot.tasks package

Submodules

skrobot.tasks.base_cross_validation_task module

class skrobot.tasks.base_cross_validation_task.**BaseCrossValidationTask**(*type_name*, *args*)

Bases: *skrobot.tasks.base_task.BaseTask*

The *BaseCrossValidationTask* is an abstract base class for implementing tasks that use cross-validation functionality.

It can support both stratified k-fold cross-validation as well as cross-validation with user-defined folds.

By default, stratified k-fold cross-validation is used with the default parameters of *stratified_folds()* method.

__init__(*type_name*, *args*)

This is the constructor method and can be used from child *BaseCrossValidationTask* implementations.

Parameters

- **type_name** (*str*) – The task's type name. A common practice is to pass the name of the task's class.
- **args** (*dict*) – The task's parameters. A common practice is to pass the parameters at the time of task's object creation. It is a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

custom_folds(*folds_file_path*, *fold_column='fold'*)

Optional method.

Use cross-validation with user-defined custom folds.

Parameters

- **folds_file_path** (*str*) – The path to the file containing the user-defined folds for the samples. The file needs to be formatted with the same separation delimiter (comma for CSV, tab for TSV, etc.) as the one used in the input data set files provided to the task. The file must contain two data columns and the first row must be the header. The first column is for the sample IDs and needs to be the same as the one used in the input data set files provided to the task. The second column is for the fold IDs (e.g., 1 through 5, A through D, etc.).
- **fold_column** (*str, optional*) – The column name for the fold IDs. It defaults to 'fold'.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

stratified_folds(*total_folds=3*, *shuffle=False*)

Optional method.

Use stratified k-fold cross-validation.

The folds are made by preserving the percentage of samples for each class.

Parameters

- **total_folds** (*int, optional*) – Number of folds. Must be at least 2. It defaults to 3.
- **shuffle** (*bool, optional*) – Whether to shuffle each class's samples before splitting into batches. Note that the samples within each split will not be shuffled. It defaults to False.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

abstract run(*output_directory*)
An abstract method for running the task.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

skrobot.tasks.base_task module

class skrobot.tasks.base_task.**BaseTask**(*type_name*, *args*)
Bases: abc.ABC

The *BaseTask* is an abstract base class for implementing tasks.

A task is a configurable and reproducible piece of code built on top of scikit-learn that can be used in machine learning pipelines.

__init__(*type_name*, *args*)
This is the constructor method and can be used from child *BaseTask* implementations.

Parameters

- **type_name** (*str*) – The task's type name. A common practice is to pass the name of the task's class.
- **args** (*dict*) – The task's parameters. A common practice is to pass the parameters at the time of task's object creation. It is a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

abstract run(*output_directory*)
An abstract method for running the task.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

skrobot.tasks.evaluation_cross_validation_task module

class skrobot.tasks.evaluation_cross_validation_task.**EvaluationCrossValidationTask**(*estimator*, *train_data* *test_data_s* *es-* *ti-* *ma-* *tor_params* *field_delim* *fea-* *ture_colum* *id_column=* *la-* *bel_colum* *ran-* *dom_seed=* *thresh-* *old_selecti* *met-* *ric_greate* *thresh-* *old_tuning* *1.0,* *0.01)*, *ex-* *port_classi* *ex-* *port_confu* *ex-* *port_roc_c* *ex-* *port_pr_cu* *ex-* *port_false* *ex-* *port_false* *ex-* *port_also* *fs-* *core_beta=*

Bases: *skrobot.tasks.base_cross_validation_task.BaseCrossValidationTask*

The *EvaluationCrossValidationTask* class can be used to evaluate a scikit-learn estimator/pipeline

on some data.

The following evaluation results can be generated on-demand for hold-out test data set as well as train/validation cross-validation folds:

- PR / ROC Curves
- Confusion Matrixes
- Classification Reports
- Performance Metrics
- False Positives
- False Negatives

It can support both stratified k-fold cross-validation as well as cross-validation with user-defined folds.

By default, stratified k-fold cross-validation is used with the default parameters of `stratified_folds()` method.

__init__(*estimator*, *train_data_set_file_path*, *test_data_set_file_path=None*, *estimator_params=None*, *field_delimiter=','*, *feature_columns='all'*, *id_column='id'*, *label_column='label'*, *random_seed=42*, *threshold_selection_by='f1'*, *metric_greater_is_better=True*, *threshold_tuning_range=(0.01, 1.0, 0.01)*, *export_classification_reports=False*, *export_confusion_matrixes=False*, *export_roc_curves=False*, *export_pr_curves=False*, *export_false_positives_reports=False*, *export_false_negatives_reports=False*, *export_also_for_train_folds=False*, *fscore_beta=1*)

This is the constructor method and can be used to create a new object instance of `EvaluationCrossValidationTask` class.

Parameters

- **estimator** (*scikit-learn {estimator, pipeline}*) – It can be either an estimator (e.g., LogisticRegression) or a pipeline ending with an estimator. The estimator needs to be able to predict probabilities through a `predict_proba` method.
- **train_data_set_file_path** (*str*) – The file path of the input train data set. It can be either a URL or a disk file path.
- **test_data_set_file_path** (*str, optional*) – The file path of the input test data set. It can be either a URL or a disk file path. It defaults to None.
- **estimator_params** (*dict, optional*) – The parameters to override in the provided estimator/pipeline. It defaults to None.
- **field_delimiter** (*str, optional*) – The separation delimiter (comma for CSV, tab for TSV, etc.) used in the input train/test data set files. It defaults to ‘,’.
- **feature_columns** (*{str, list}, optional*) – Either ‘all’ to use from the input train/test data set files all the columns or a list of column names to select specific columns. It defaults to ‘all’.
- **id_column** (*str, optional*) – The name of the column in the input train/test data set files containing the sample IDs. It defaults to ‘id’.
- **label_column** (*str, optional*) – The name of the column in the input train/test data set files containing the ground truth labels. It defaults to ‘label’.
- **random_seed** (*int, optional*) – The random seed used in the random number generator. It can be used to reproduce the output. It defaults to 42.

- **threshold_selection_by** (*{str, float}, optional*) – The evaluation results will be generated either for a specific provided threshold value (e.g., 0.49) or for the best threshold found from threshold tuning, based on a specific provided metric (e.g., 'f1', 'f0.55'). It defaults to 'f1'.
- **metric_greater_is_better** (*bool, optional*) – This flag will control the direction of searching of the best threshold and it depends on the provided metric in `threshold_selection_by`. True, means that greater metric values is better and False means the opposite. It defaults to True.
- **threshold_tuning_range** (*tuple, optional*) – A range in form (start_value, stop_value, step_size) for generating a sequence of threshold values in threshold tuning. It generates the sequence by incrementing the start value using the step size until it reaches the stop value. It defaults to (0.01, 1.0, 0.01).
- **export_classification_reports** (*bool, optional*) – If this task will export classification reports. It defaults to False.
- **export_confusion_matrixes** (*bool, optional*) – If this task will export confusion matrixes. It defaults to False.
- **export_roc_curves** (*bool, optional*) – If this task will export ROC curves. It defaults to False.
- **export_pr_curves** (*bool, optional*) – If this task will export PR curves. It defaults to False.
- **export_false_positives_reports** (*bool, optional*) – If this task will export false positives reports. It defaults to False.
- **export_false_negatives_reports** (*bool, optional*) – If this task will export false negatives reports. It defaults to False.
- **export_also_for_train_folds** (*bool, optional*) – If this task will export the evaluation results also for the train folds of cross-validation. It defaults to False.
- **fscore_beta** (*float, optional*) – The beta parameter in F-measure. It determines the weight of recall in the score. *beta < 1* lends more weight to precision, while *beta > 1* favors recall (*beta -> 0* considers only precision, *beta -> +inf* only recall). It defaults to 1.

run (*output_directory*)
Run the task.

All of the evaluation results are stored as files under the output directory path.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

Returns The task's result. Specifically, the threshold used along with its related performance metrics and summary metrics from all cross-validation splits as well as hold-out test data set.

Return type dict

custom_folds (*folds_file_path*, *fold_column='fold'*)
Optional method.

Use cross-validation with user-defined custom folds.

Parameters

- **folds_file_path** (*str*) – The path to the file containing the user-defined folds for the samples. The file needs to be formatted with the same separation delimiter (comma for

CSV, tab for TSV, etc.) as the one used in the input data set files provided to the task. The file must contain two data columns and the first row must be the header. The first column is for the sample IDs and needs to be the same as the one used in the input data set files provided to the task. The second column is for the fold IDs (e.g., 1 through 5, A through D, etc.).

- **fold_column** (*str, optional*) – The column name for the fold IDs. It defaults to 'fold'.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

stratified_folds(*total_folds=3*, *shuffle=False*)
Optional method.

Use stratified k-fold cross-validation.

The folds are made by preserving the percentage of samples for each class.

Parameters

- **total_folds** (*int, optional*) – Number of folds. Must be at least 2. It defaults to 3.
- **shuffle** (*bool, optional*) – Whether to shuffle each class's samples before splitting into batches. Note that the samples within each split will not be shuffled. It defaults to False.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

skrobot.tasks.feature_selection_cross_validation_task module

class skrobot.tasks.feature_selection_cross_validation_task.**FeatureSelectionCrossValidatio**

Bases: *skrobot.tasks.base_cross_validation_task.BaseCrossValidationTask*

The *FeatureSelectionCrossValidationTask* class can be used to perform feature selection with Recursive Feature Elimination using a scikit-learn estimator on some data.

A scikit-learn preprocessor can be used on the input train data set before feature selection runs.

It can support both stratified k-fold cross-validation as well as cross-validation with user-defined folds.

By default, stratified k-fold cross-validation is used with the default parameters of *stratified_folds()* method.

__init__(*estimator*, *train_data_set_file_path*, *estimator_params=None*, *field_delimiter=','*, *preprocessor=None*, *preprocessor_params=None*, *min_features_to_select=1*, *scoring='f1'*, *feature_columns='all'*, *id_column='id'*, *label_column='label'*, *random_seed=42*, *verbose=3*, *n_jobs=1*)

This is the constructor method and can be used to create a new object instance of *FeatureSelectionCrossValidationTask* class.

Parameters

- **estimator** (*scikit-learn estimator*) – An estimator (e.g., LogisticRegression). It needs to provide feature importances through either a `coef_` or a `feature_importances_` attribute.
- **train_data_set_file_path** (*str*) – The file path of the input train data set. It can be either a URL or a disk file path.

- **estimator_params** (*dict, optional*) – The parameters to override in the provided estimator. It defaults to None.
- **field_delimiter** (*str, optional*) – The separation delimiter (comma for CSV, tab for TSV, etc.) used in the input train data set file. It defaults to ‘,’.
- **preprocessor** (*scikit-learn preprocessor, optional*) – The preprocessor you want to run on the input train data set before feature selection. You can set for example a scikit-learn ColumnTransformer, OneHotEncoder, etc. It defaults to None.
- **preprocessor_params** (*dict, optional*) – The parameters to override in the provided preprocessor. It defaults to None.
- **min_features_to_select** (*int, optional*) – The minimum number of features to be selected. This number of features will always be scored. It defaults to 1.
- **scoring** (*{str, callable}, optional*) – A single scikit-learn scorer string (e.g., ‘f1’) or a callable that is built with scikit-learn `make_scorer`. Note that when using custom scorers, each scorer should return a single value. It defaults to ‘f1’.
- **feature_columns** (*{str, list}, optional*) – Either ‘all’ to use from the input train data set file all the columns or a list of column names to select specific columns. It defaults to ‘all’.
- **id_column** (*str, optional*) – The name of the column in the input train data set file containing the sample IDs. It defaults to ‘id’.
- **label_column** (*str, optional*) – The name of the column in the input train data set file containing the ground truth labels. It defaults to ‘label’.
- **random_seed** (*int, optional*) – The random seed used in the random number generator. It can be used to reproduce the output. It defaults to 42.
- **verbose** (*int, optional*) – Controls the verbosity of output. The higher, the more messages. It defaults to 3.
- **n_jobs** (*int, optional*) – Number of jobs to run in parallel. -1 means using all processors. It defaults to 1.

run (*output_directory*)
Run the task.

The selected features are returned as a result and also stored in a *features_selected.txt* text file under the output directory path.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

Returns The task’s result. Specifically, the selected features, which can be either column names from the input train data set or column indexes from the preprocessed data set, depending on whether a `preprocessor` was used or not.

Return type list

custom_folds (*folds_file_path*, *fold_column='fold'*)
Optional method.

Use cross-validation with user-defined custom folds.

Parameters

- **folds_file_path** (*str*) – The path to the file containing the user-defined folds for the samples. The file needs to be formatted with the same separation delimiter (comma for CSV, tab for TSV, etc.) as the one used in the input data set files provided to the task. The

file must contain two data columns and the first row must be the header. The first column is for the sample IDs and needs to be the same as the one used in the input data set files provided to the task. The second column is for the fold IDs (e.g., 1 through 5, A through D, etc.).

- **fold_column** (*str, optional*) – The column name for the fold IDs. It defaults to 'fold'.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

stratified_folds(*total_folds=3*, *shuffle=False*)
Optional method.

Use stratified k-fold cross-validation.

The folds are made by preserving the percentage of samples for each class.

Parameters

- **total_folds** (*int, optional*) – Number of folds. Must be at least 2. It defaults to 3.
- **shuffle** (*bool, optional*) – Whether to shuffle each class's samples before splitting into batches. Note that the samples within each split will not be shuffled. It defaults to False.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

skrobot.tasks.hyperparameters_search_cross_validation_task module

class skrobot.tasks.hyperparameters_search_cross_validation_task.**HyperParametersSearchCross**

Bases: *skrobot.tasks.base_cross_validation_task.BaseCrossValidationTask*

The *HyperParametersSearchCrossValidationTask* class can be used to search the best hyperparameters of a scikit-learn estimator/pipeline on some data.

Cross-Validation

It can support both stratified k-fold cross-validation as well as cross-validation with user-defined folds.

By default, stratified k-fold cross-validation is used with the default parameters of *stratified_folds()* method.

Search

It can support both grid search as well as random search.

By default, grid search is used.

__init__(*estimator*, *search_params*, *train_data_set_file_path*, *estimator_params=None*, *field_delimiter=','*, *scorers=['roc_auc', 'average_precision', 'f1', 'precision', 'recall', 'accuracy']*, *feature_columns='all'*, *id_column='id'*, *label_column='label'*, *objective_score='f1'*, *random_seed=42*, *verbose=3*, *n_jobs=1*, *return_train_score=True*)

This is the constructor method and can be used to create a new object instance of *HyperParametersSearchCrossValidationTask* class.

Parameters

- **estimator** (*scikit-learn {estimator, pipeline}*) – It can be either an estimator (e.g., LogisticRegression) or a pipeline ending with an estimator.
- **search_params** (*{dict, list of dictionaries}*) – Dictionary with hyperparameters names as keys and lists of hyperparameter settings to try as values, or a list of such dictionaries, in which case the grids spanned by each dictionary in the list are explored. This enables searching over any sequence of hyperparameter settings.
- **train_data_set_file_path** (*str*) – The file path of the input train data set. It can be either a URL or a disk file path.
- **estimator_params** (*dict, optional*) – The parameters to override in the provided estimator/pipeline. It defaults to None.
- **field_delimiter** (*str, optional*) – The separation delimiter (comma for CSV, tab for TSV, etc.) used in the input train data set file. It defaults to ‘,’.
- **scorers** (*{list, dict}, optional*) – Multiple metrics to evaluate the predictions on the hold-out data. Either give a list of (unique) strings or a dict with names as keys and callables as values. The callables should be scorers built using scikit-learn `make_scorer`. Note that when using custom scorers, each scorer should return a single value. It defaults to [‘roc_auc’, ‘average_precision’, ‘f1’, ‘precision’, ‘recall’, ‘accuracy’].
- **feature_columns** (*{str, list}, optional*) – Either ‘all’ to use from the input train data set file all the columns or a list of column names to select specific columns. It defaults to ‘all’.
- **id_column** (*str, optional*) – The name of the column in the input train data set file containing the sample IDs. It defaults to ‘id’.
- **label_column** (*str, optional*) – The name of the column in the input train data set file containing the ground truth labels. It defaults to ‘label’.
- **objective_score** (*str, optional*) – The scorer that would be used to find the best hyperparameters for refitting the best estimator/pipeline at the end. It defaults to ‘f1’.
- **random_seed** (*int, optional*) – The random seed used in the random number generator. It can be used to reproduce the output. It defaults to 42.
- **verbose** (*int, optional*) – Controls the verbosity of output. The higher, the more messages. It defaults to 3.
- **n_jobs** (*int, optional*) – Number of jobs to run in parallel. -1 means using all processors. It defaults to 1.
- **return_train_score** (*bool, optional*) – If False, training scores will not be computed and returned. Computing training scores is used to get insights on how different parameter settings impact the overfitting/underfitting trade-off. It defaults to True.

grid_search()

Optional method.

Use the grid search method when searching the best hyperparameters.

Returns The object instance itself.

Return type *HyperParametersSearchCrossValidationTask*

random_search(*n_iters=200*)
Optional method.

Use the random search method when searching the best hyperparameters.

Parameters **n_iters** (*int, optional*) – Number of hyperparameter settings that are sampled. n_iters trades off runtime vs quality of the solution. It defaults to 200.

Returns The object instance itself.

Return type *HyperParametersSearchCrossValidationTask*

run(*output_directory*)
Run the task.

The search results (search_results) are stored also in a *search_results.html* file as a static HTML table under the output directory path.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

Returns The task's result. Specifically, **1)** best_estimator: The estimator/pipeline that was chosen by the search, i.e. estimator/pipeline which gave best score on the hold-out data. **2)** best_params: The hyperparameters setting that gave the best results on the hold-out data. **3)** best_score: Mean cross-validated score of the best_estimator. **4)** search_results: Metrics measured for each of the hyperparameters setting in the search. **5)** best_index: The index (of the search_results) which corresponds to the best candidate hyperparameters setting.

Return type dict

custom_folds(*folds_file_path*, *fold_column='fold'*)
Optional method.

Use cross-validation with user-defined custom folds.

Parameters

- **folds_file_path** (*str*) – The path to the file containing the user-defined folds for the samples. The file needs to be formatted with the same separation delimiter (comma for CSV, tab for TSV, etc.) as the one used in the input data set files provided to the task. The file must contain two data columns and the first row must be the header. The first column is for the sample IDs and needs to be the same as the one used in the input data set files provided to the task. The second column is for the fold IDs (e.g., 1 through 5, A through D, etc.).
- **fold_column** (*str, optional*) – The column name for the fold IDs. It defaults to 'fold'.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

stratified_folds(*total_folds=3*, *shuffle=False*)
Optional method.

Use stratified k-fold cross-validation.

The folds are made by preserving the percentage of samples for each class.

Parameters

- **total_folds** (*int, optional*) – Number of folds. Must be at least 2. It defaults to 3.
- **shuffle** (*bool, optional*) – Whether to shuffle each class's samples before splitting into batches. Note that the samples within each split will not be shuffled. It defaults to False.

Returns The object instance itself.

Return type *BaseCrossValidationTask*

skrobot.tasks.prediction_task module

class skrobot.tasks.prediction_task.**PredictionTask**(*estimator*, *data_set_file_path*, *field_delimiter=','*, *feature_columns='all'*, *id_column='id'*, *prediction_column='prediction'*, *threshold=0.5*)
Bases: *skrobot.tasks.base_task.BaseTask*

The *PredictionTask* class can be used to predict new data using a scikit-learn estimator/pipeline.

__init__(*estimator*, *data_set_file_path*, *field_delimiter=','*, *feature_columns='all'*, *id_column='id'*, *prediction_column='prediction'*, *threshold=0.5*)
This is the constructor method and can be used to create a new object instance of *PredictionTask* class.

Parameters

- **estimator** (*scikit-learn {estimator, pipeline}*) – It can be either an estimator (e.g., LogisticRegression) or a pipeline ending with an estimator. The estimator needs to be able to predict probabilities through a `predict_proba` method.
- **data_set_file_path** (*str*) – The file path of the input data set. It can be either a URL or a disk file path.
- **field_delimiter** (*str, optional*) – The separation delimiter (comma for CSV, tab for TSV, etc.) used in the input data set file. It defaults to ','.
- **feature_columns** (*{str, list}, optional*) – Either 'all' to use from the input data set file all the columns or a list of column names to select specific columns. It defaults to 'all'.
- **id_column** (*str, optional*) – The name of the column in the input data set file containing the sample IDs. It defaults to 'id'.

- **prediction_column** (*str, optional*) – The name of the column for the predicted binary class labels. It defaults to 'prediction'.
- **threshold** (*float, optional*) – The threshold to use for converting the predicted probability into a binary class label. It defaults to 0.5.

run(*output_directory*)
Run the task.

The predictions are returned as a result and also stored in a *predictions.csv* CSV file under the output directory path.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

Returns The task's result. Specifically, the predictions for the input data set, containing the sample IDs, the predicted binary class labels, and the predicted probabilities for the positive class.

Return type pandas DataFrame

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

skrobot.tasks.train_task module

class skrobot.tasks.train_task.**TrainTask**(*estimator*, *train_data_set_file_path*, *estimator_params=None*, *field_delimiter=','*, *feature_columns='all'*, *id_column='id'*, *label_column='label'*, *random_seed=42*)
Bases: *skrobot.tasks.base_task.BaseTask*

The *TrainTask* class can be used to fit a scikit-learn estimator/pipeline on train data.

__init__(*estimator*, *train_data_set_file_path*, *estimator_params=None*, *field_delimiter=','*, *feature_columns='all'*, *id_column='id'*, *label_column='label'*, *random_seed=42*)
This is the constructor method and can be used to create a new object instance of *TrainTask* class.

Parameters

- **estimator** (*scikit-learn {estimator, pipeline}*) – It can be either an estimator (e.g., LogisticRegression) or a pipeline ending with an estimator.
- **train_data_set_file_path** (*str*) – The file path of the input train data set. It can be either a URL or a disk file path.
- **estimator_params** (*dict, optional*) – The parameters to override in the provided estimator/pipeline. It defaults to None.
- **field_delimiter** (*str, optional*) – The separation delimiter (comma for CSV, tab for TSV, etc.) used in the input train data set file. It defaults to ','.

- **feature_columns** (*{str, list}, optional*) – Either 'all' to use from the input train data set file all the columns or a list of column names to select specific columns. It defaults to 'all'.
- **id_column** (*str, optional*) – The name of the column in the input train data set file containing the sample IDs. It defaults to 'id'.
- **label_column** (*str, optional*) – The name of the column in the input train data set file containing the ground truth labels. It defaults to 'label'.
- **random_seed** (*int, optional*) – The random seed used in the random number generator. It can be used to reproduce the output. It defaults to 42.

run(*output_directory*)
Run the task.

The fitted estimator/pipeline is returned as a result and also stored in a *trained_model.pkl* pickle file under the output directory path.

Parameters **output_directory** (*str*) – The output directory path under which task-related generated files are stored.

Returns The task's result. Specifically, the fitted estimator/pipeline.

Return type dict

get_configuration()
Get the task's parameters.

Returns The task's parameters as a dictionary of key-value pairs, where the key is the parameter name and the value is the parameter value.

Return type dict

get_type()
Get the task's type name.

Returns The task's type name.

Return type str

CHAPTER

TWO

WHAT IS IT ABOUT?

skrobot is a Python module for designing, running and tracking Machine Learning experiments / tasks. It is built on top of scikit-learn framework.

CHAPTER

THREE

WHY DOES IT EXISTS?

It can help Data Scientists and Machine Learning Engineers:

- to keep track of modelling experiments / tasks
- to automate the repetitive (and boring) stuff when designing modelling pipelines
- to spend more time on the things that truly matter when solving a problem

CHAPTER

FOUR

HOW DO I INSTALL IT?

```
$ pip install skrobot
```

CHAPTER

FIVE

WHICH ARE THE COMPONENTS?

NOTE : Currently, skrobot can be used only for binary classification problems.

5.1 For the module's users

Component	What is this?
Train Task	This task can be used to fit a scikit-learn estimator on some data.
Prediction Task	This task can be used to predict new data using a scikit-learn estimator.
Evaluation Cross Validation Task	This task can be used to evaluate a scikit-learn estimator on some data.
Feature Selection Cross Validation Task	This task can be used to perform feature selection with Recursive Feature Elimination using a scikit-learn estimator on some data.
Hyperparameters Search Cross Validation Task	This task can be used to search the best hyperparameters of a scikit-learn estimator on some data.
Experiment	This is used to build, track and run an experiment. It can run tasks in the context of an experiment.
Task Runner	This is a simplified version (in functionality) of the Experiment component. It leaves out all the "experiment" stuff and is focused mostly in the execution and tracking of tasks.

5.2 For the module's developers

Component	What is this?
Base Task	All tasks inherit from this component. A task is a configurable and reproducible piece of code built on top of scikit-learn that can be used in machine learning pipelines.
Base Cross Validation Task	All tasks that use cross validation functionality inherit from this component.
Base Notifier	All notifiers inherit from this component. A notifier can be used to send success / failure notifications for tasks execution.

CHAPTER

SIX

HOW DO I USE IT?

The following examples use many of skrobot's components to built a machine learning modelling pipeline. Please try them and we would love to have your feedback! Furthermore, many examples can be found in the project's repository.

6.1 Example on Titanic Dataset

The following example has generated the following results.

```
from sklearn.compose import ColumnTransformer
from sklearn.pipeline import Pipeline
from sklearn.impute import SimpleImputer
from sklearn.preprocessing import StandardScaler, OneHotEncoder
from sklearn.linear_model import LogisticRegression

from skrobot.core import Experiment
from skrobot.tasks import TrainTask
from skrobot.tasks import PredictionTask
from skrobot.tasks import FeatureSelectionCrossValidationTask
from skrobot.tasks import EvaluationCrossValidationTask
from skrobot.tasks import HyperParametersSearchCrossValidationTask
from skrobot.feature_selection import ColumnSelector
from skrobot.notification import BaseNotifier

######### Initialization Code

train_data_set_file_path = 'https://bit.ly/titanic-data-train'

test_data_set_file_path = 'https://bit.ly/titanic-data-test'

new_data_set_file_path = 'https://bit.ly/titanic-data-new'

random_seed = 42

id_column = 'PassengerId'

label_column = 'Survived'

numerical_features = ['Age', 'Fare', 'SibSp', 'Parch']

categorical_features = ['Embarked', 'Sex', 'Pclass']

numeric_transformer = Pipeline(steps=[
    ('imputer', SimpleImputer()),
```

(continues on next page)

(continued from previous page)

```
    ('scaler', StandardScaler())])

categorical_transformer = Pipeline(steps=[
    ('imputer', SimpleImputer(strategy='most_frequent')),
    ('encoder', OneHotEncoder(handle_unknown='ignore'))])

preprocessor = ColumnTransformer(transformers=[
    ('numerical_transformer', numeric_transformer, numerical_features),
    ('categorical_transformer', categorical_transformer, categorical_features)])

classifier = LogisticRegression(solver='liblinear', random_state=random_seed)

search_params = {
    "classifier__C" : [ 1.e-01, 1.e+00, 1.e+01 ],
    "classifier__penalty" : [ "l1", "l2" ],
    "preprocessor__numerical_transformer__imputer__strategy" : [ "mean", "median" ]
}

######### skrobot Code

# Define a Notifier (This is optional and you can implement any notifier you want, e.
↪g. for Slack / Trello / Discord)
class ConsoleNotifier(BaseNotifier):
    def notify (self, message):
        print(message)

# Build an Experiment
experiment = Experiment('experiments-output').set_source_code_file_path(__file__).set_
↪experimenter('echatzikyriakidis').set_notifier(ConsoleNotifier()).build()

# Run Feature Selection Task
features_columns = experiment.run(FeatureSelectionCrossValidationTask␣
↪(estimator=classifier,
                                                                    train_data_set_
↪file_path=train_data_set_file_path,
                                                                    ␣
↪preprocessor=preprocessor,
                                                                    min_features_
↪to_select=4,
                                                                    id_column=id_
↪column,
                                                                    label_
↪column=label_column,
                                                                    random_
↪seed=random_seed).stratified_folds(total_folds=5, shuffle=True))

pipe = Pipeline(steps=[('preprocessor', preprocessor),
                       ('selector', ColumnSelector(cols=features_columns)),
                       ('classifier', classifier)])

# Run Hyperparameters Search Task
hyperparameters_search_results = experiment.
↪run(HyperParametersSearchCrossValidationTask (estimator=pipe,
↪    search_params=search_params,
↪    train_data_set_file_path=train_data_set_file_path,
```

(continues on next page)

(continued from previous page)

```
↪    id_column=id_column,
↪    label_column=label_column,
↪    random_seed=random_seed).random_search(n_iters=100).stratified_folds(total_
↪folds=5, shuffle=True))

# Run Evaluation Task
evaluation_results = experiment.run(EvaluationCrossValidationTask(estimator=pipe,
                                                                  estimator_
↪params=hyperparameters_search_results['best_params'],
                                                                  train_data_set_file_
↪path=train_data_set_file_path,
                                                                  test_data_set_file_
↪path=test_data_set_file_path,
                                                                  id_column=id_column,
                                                                  label_column=label_
↪column,
                                                                  random_seed=random_
↪seed,
                                                                  export_
↪classification_reports=True,
                                                                  export_confusion_
↪matrixes=True,
                                                                  export_pr_
↪curves=True,
                                                                  export_roc_
↪curves=True,
                                                                  export_false_
↪positives_reports=True,
                                                                  export_false_
↪negatives_reports=True,
                                                                  export_also_for_
↪train_folds=True).stratified_folds(total_folds=5, shuffle=True))

# Run Train Task
train_results = experiment.run(TrainTask(estimator=pipe,
                                         estimator_params=hyperparameters_search_
↪results['best_params'],
                                         train_data_set_file_path=train_data_set_file_
↪path,
                                         id_column=id_column,
                                         label_column=label_column,
                                         random_seed=random_seed))

# Run Prediction Task
predictions = experiment.run(PredictionTask(estimator=train_results['estimator'],
                                            data_set_file_path=new_data_set_file_path,
                                            id_column=id_column,
                                            prediction_column=label_column,
                                            threshold=evaluation_results['threshold
↪']))

# Print in-memory results
print(features_columns)
```

(continues on next page)

(continued from previous page)

```
print(hyperparameters_search_results['best_params'])
print(hyperparameters_search_results['best_index'])
print(hyperparameters_search_results['best_estimator'])
print(hyperparameters_search_results['best_score'])
print(hyperparameters_search_results['search_results'])

print(evaluation_results['threshold'])
print(evaluation_results['cv_threshold_metrics'])
print(evaluation_results['cv_splits_threshold_metrics'])
print(evaluation_results['cv_splits_threshold_metrics_summary'])
print(evaluation_results['test_threshold_metrics'])

print(train_results['estimator'])

print(predictions)
```

6.2 Example on SMS Spam Collection Dataset

The following example has generated the following results.

```
from sklearn.pipeline import Pipeline
from sklearn.feature_extraction.text import CountVectorizer, TfidfTransformer
from sklearn.feature_selection import SelectPercentile, chi2
from sklearn.linear_model import SGDClassifier

from skrobot.core import Experiment
from skrobot.tasks import TrainTask
from skrobot.tasks import PredictionTask
from skrobot.tasks import EvaluationCrossValidationTask
from skrobot.tasks import HyperParametersSearchCrossValidationTask
from skrobot.feature_selection import ColumnSelector

######### Initialization Code

train_data_set_file_path = 'https://bit.ly/sms-spam-ham-data-train'

test_data_set_file_path = 'https://bit.ly/sms-spam-ham-data-test'

new_data_set_file_path = 'https://bit.ly/sms-spam-ham-data-new'

field_delimiter = '\t'

random_seed = 42

pipe = Pipeline(steps=[
    ('column_selection', ColumnSelector(cols=['message'], drop_axis=True)),
    ('vectorizer', CountVectorizer()),
    ('tfidf', TfidfTransformer()),
    ('feature_selection', SelectPercentile(chi2)),
    ('classifier', SGDClassifier(loss='log'))])

search_params = {
    'classifier__max_iter': [ 20, 50, 80 ],
    'classifier__alpha': [ 0.00001, 0.000001 ],
```

(continues on next page)

(continued from previous page)

```
    'classifier__penalty': [ 'l2', 'elasticnet' ],
    "vectorizer__stop_words" : [ "english", None ],
    "vectorizer__ngram_range" : [ (1, 1), (1, 2) ],
    "vectorizer__max_df": [ 0.5, 0.75, 1.0 ],
    "tfidf__use_idf" : [ True, False ],
    "tfidf__norm" : [ 'l1', 'l2' ],
    "feature_selection__percentile" : [ 70, 60, 50 ]
}

######### skrobot Code

# Build an Experiment
experiment = Experiment('experiments-output').set_source_code_file_path(__file__).set_
↪experimenter('echatzikyriakidis').build()

# Run Hyperparameters Search Task
hyperparameters_search_results = experiment.
↪run(HyperParametersSearchCrossValidationTask (estimator=pipe,
↪     search_params=search_params,
↪     train_data_set_file_path=train_data_set_file_path,
↪     field_delimiter=field_delimiter,
↪     random_seed=random_seed).random_search().stratified_folds(total_folds=5,␣
↪shuffle=True))

# Run Evaluation Task
evaluation_results = experiment.run(EvaluationCrossValidationTask(estimator=pipe,
                                                                  estimator_
↪params=hyperparameters_search_results['best_params'],
                                                                  train_data_set_file_
↪path=train_data_set_file_path,
                                                                  test_data_set_file_
↪path=test_data_set_file_path,
                                                                  field_
↪delimiter=field_delimiter,
                                                                  random_seed=random_
↪seed,
                                                                  export_
↪classification_reports=True,
                                                                  export_confusion_
↪matrixes=True,
                                                                  export_pr_
↪curves=True,
                                                                  export_roc_
↪curves=True,
                                                                  export_false_
↪positives_reports=True,
                                                                  export_false_
↪negatives_reports=True,
                                                                  export_also_for_
↪train_folds=True).stratified_folds(total_folds=5, shuffle=True))

# Run Train Task
train_results = experiment.run(TrainTask(estimator=pipe,
```

(continues on next page)

(continued from previous page)

```
                                            estimator_params=hyperparameters_search_
↪results['best_params'],
                                            train_data_set_file_path=train_data_set_file_
↪path,
                                            field_delimiter=field_delimiter,
                                            random_seed=random_seed))

# Run Prediction Task
predictions = experiment.run(PredictionTask(estimator=train_results['estimator'],
                                            data_set_file_path=new_data_set_file_path,
                                            field_delimiter=field_delimiter,
                                            threshold=evaluation_results['threshold
↪']))

# Print in-memory results
print(hyperparameters_search_results['best_params'])
print(hyperparameters_search_results['best_index'])
print(hyperparameters_search_results['best_estimator'])
print(hyperparameters_search_results['best_score'])
print(hyperparameters_search_results['search_results'])

print(evaluation_results['threshold'])
print(evaluation_results['cv_threshold_metrics'])
print(evaluation_results['cv_splits_threshold_metrics'])
print(evaluation_results['cv_splits_threshold_metrics_summary'])
print(evaluation_results['test_threshold_metrics'])

print(train_results['estimator'])

print(predictions)
```

CHAPTER

SEVEN

SAMPLE OF GENERATED RESULTS?

7.1 Classification Reports

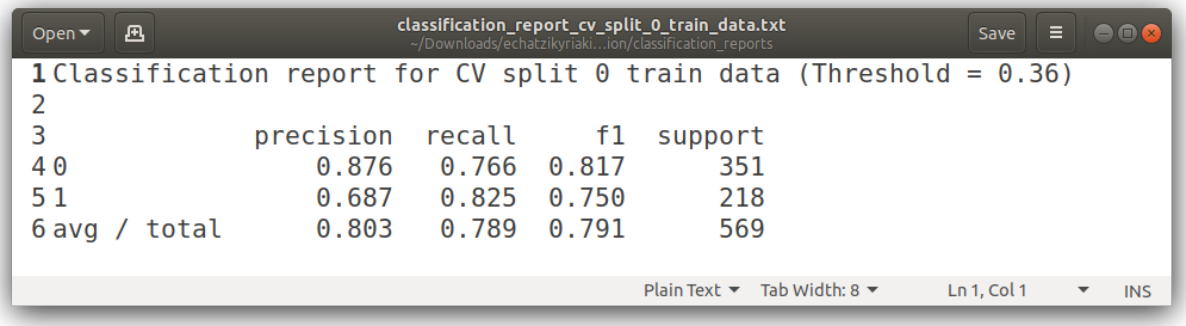

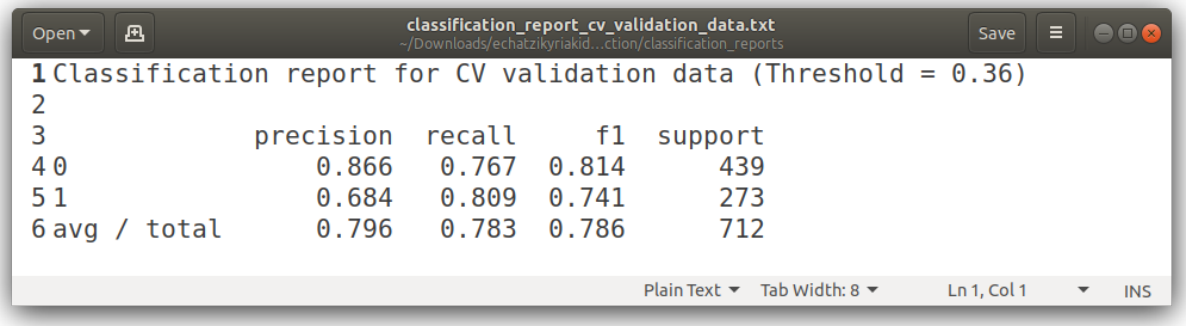

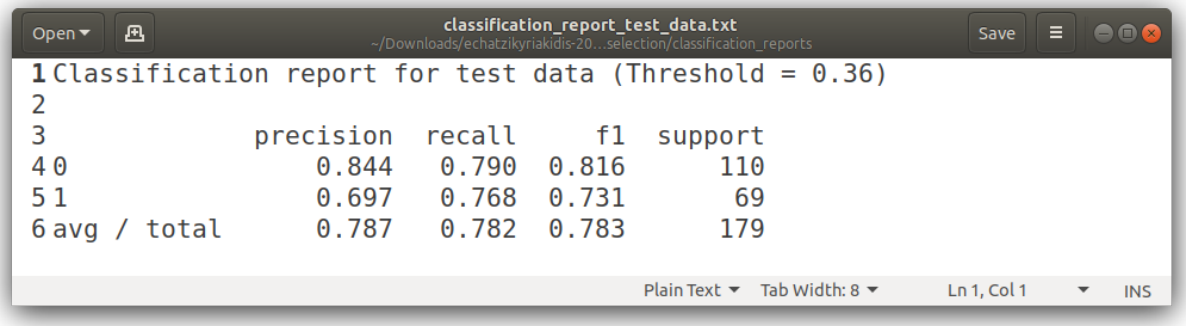

7.2 Confusion Matrixes

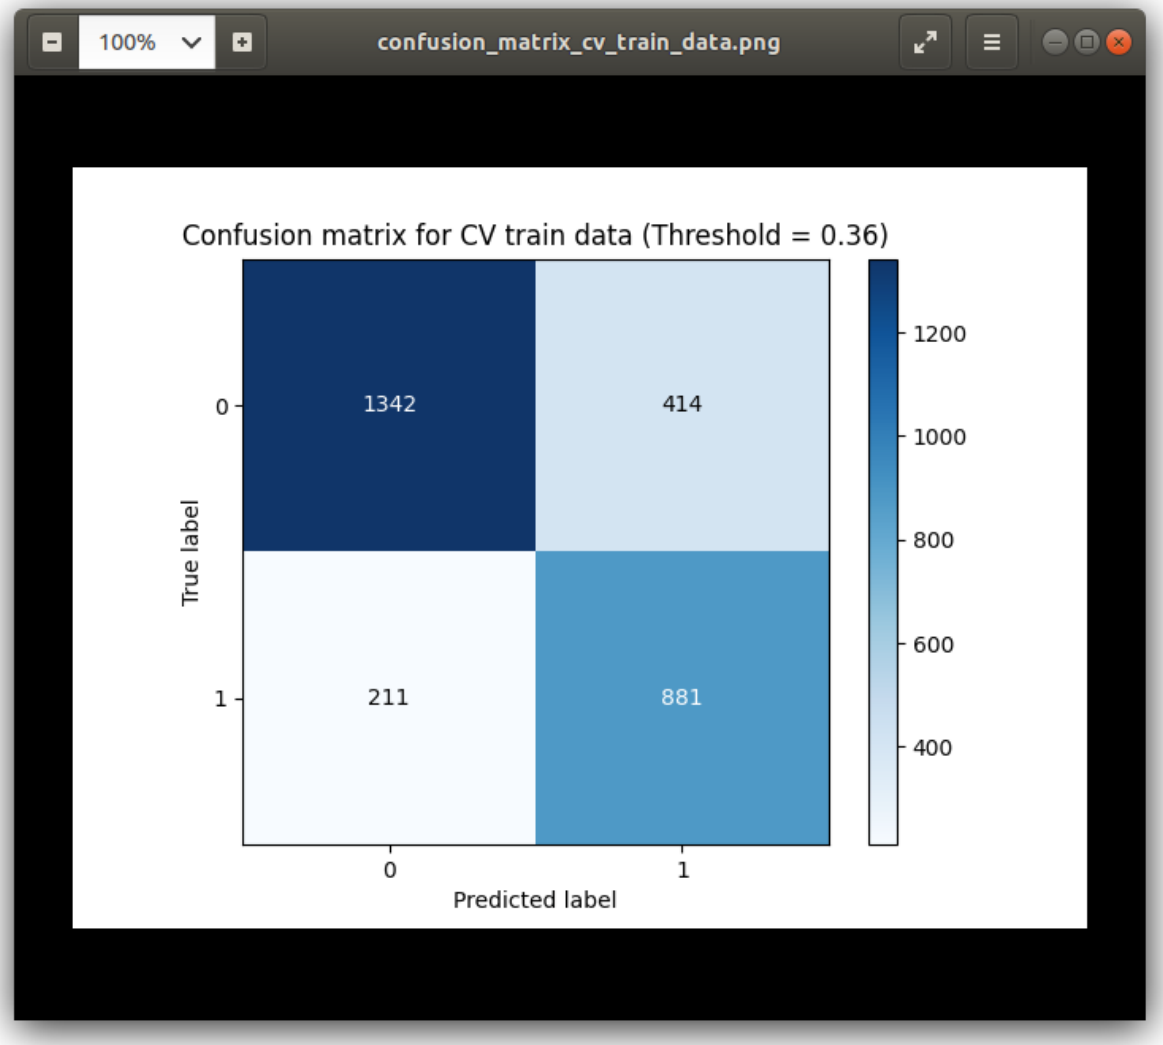
100%
confusion_matrix_cv_train_data.png
Confusion matrix for CV train data (Threshold = 0.36)
1342
414
211
881
0
1
True label
0
1
Predicted label
1200
1000
800
600
400


100%
confusion_matrix_test_data.png
Confusion matrix for test data (Threshold = 0.36)
87
23
16
53
0
1
True label
0
1
Predicted label
80
70
60
50
40
30
20

7.3 False Negatives

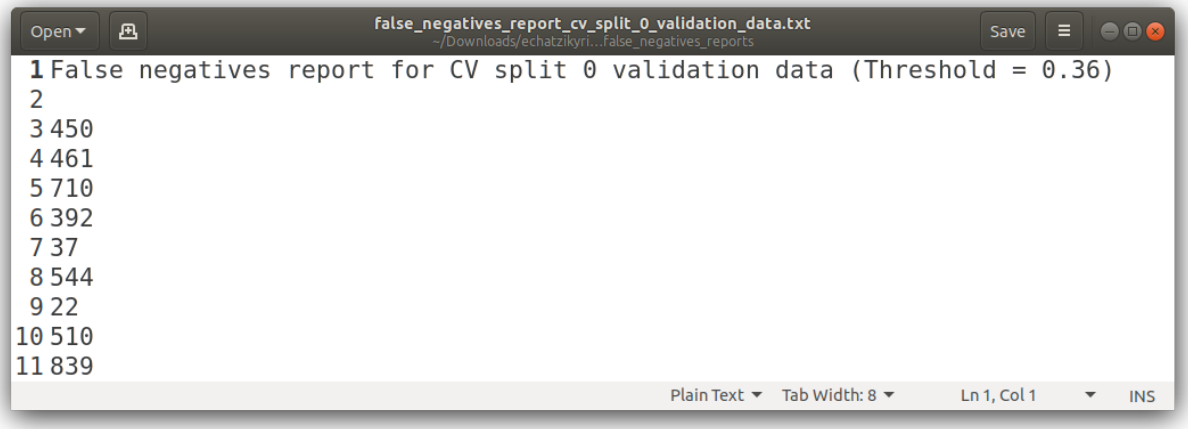

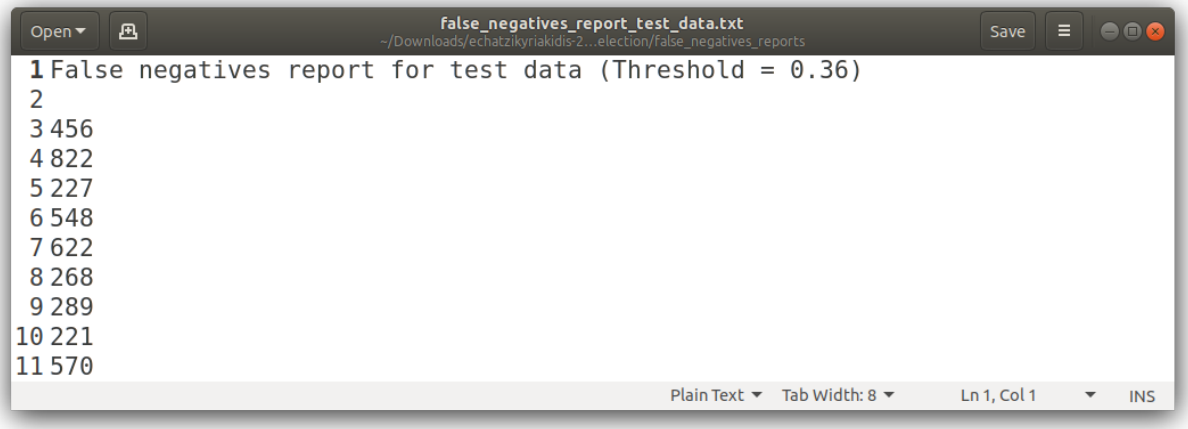

7.4 False Positives

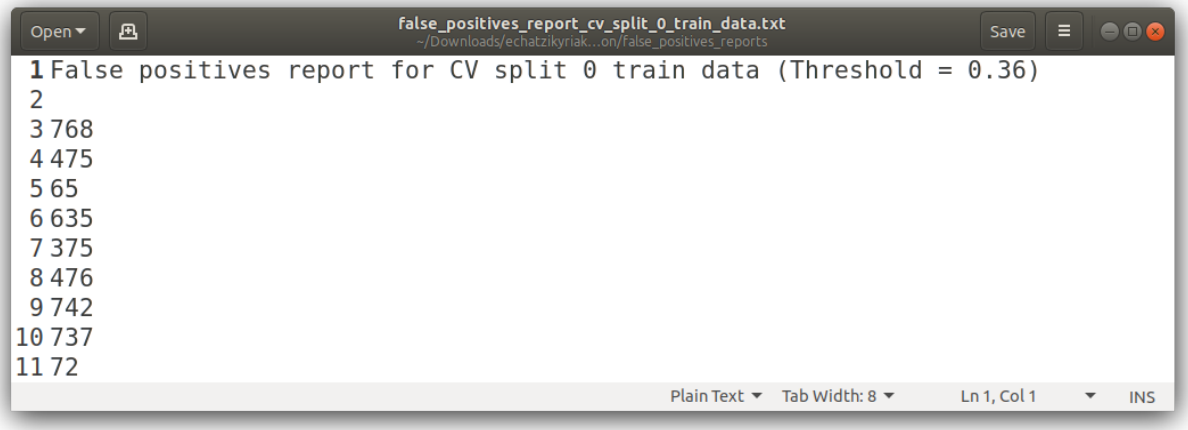

Open | false_positives_report_test_data.txt ~/Downloads/echatzikyriakidis-2...election/false_positives_reports | Save

```
1 False positives report for test data (Threshold = 0.36)
2
3 358
4 246
5 333
6 528
7 594
8 542
9 410
10 618
11 453
```

Plain Text | Tab Width: 8 | Ln 1, Col 1 | INS

7.5 PR Curves

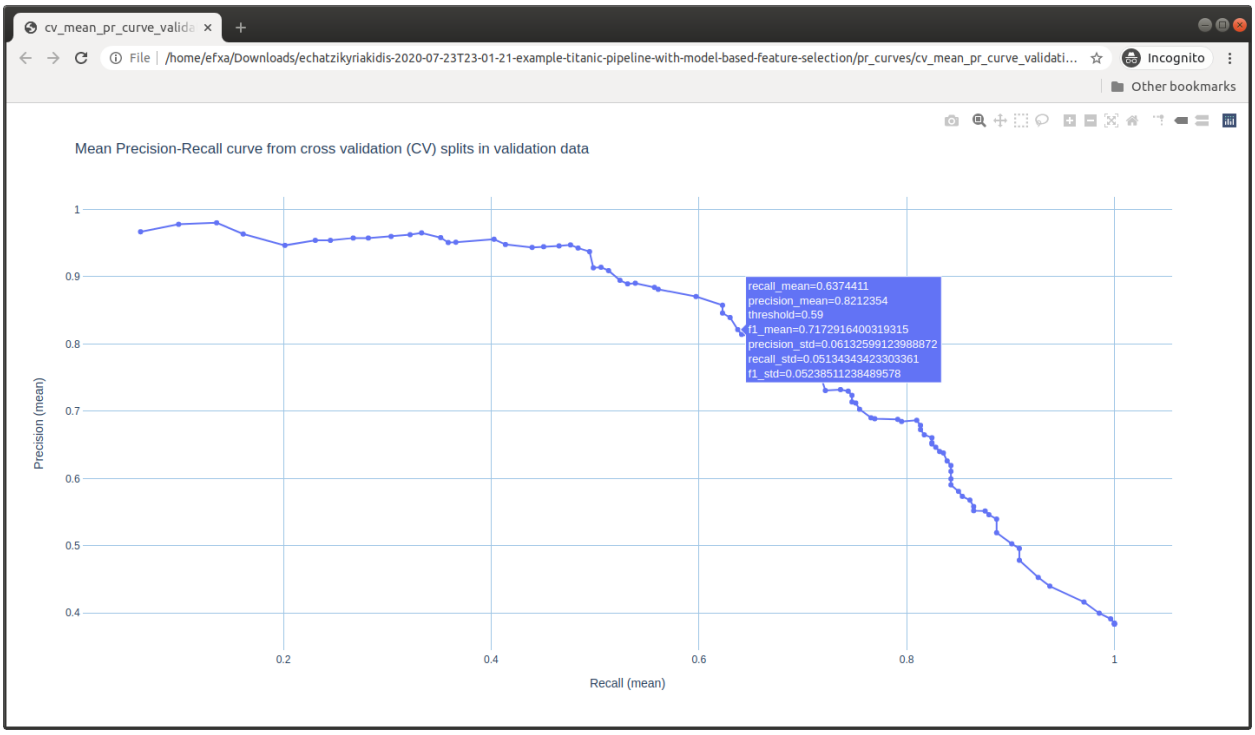

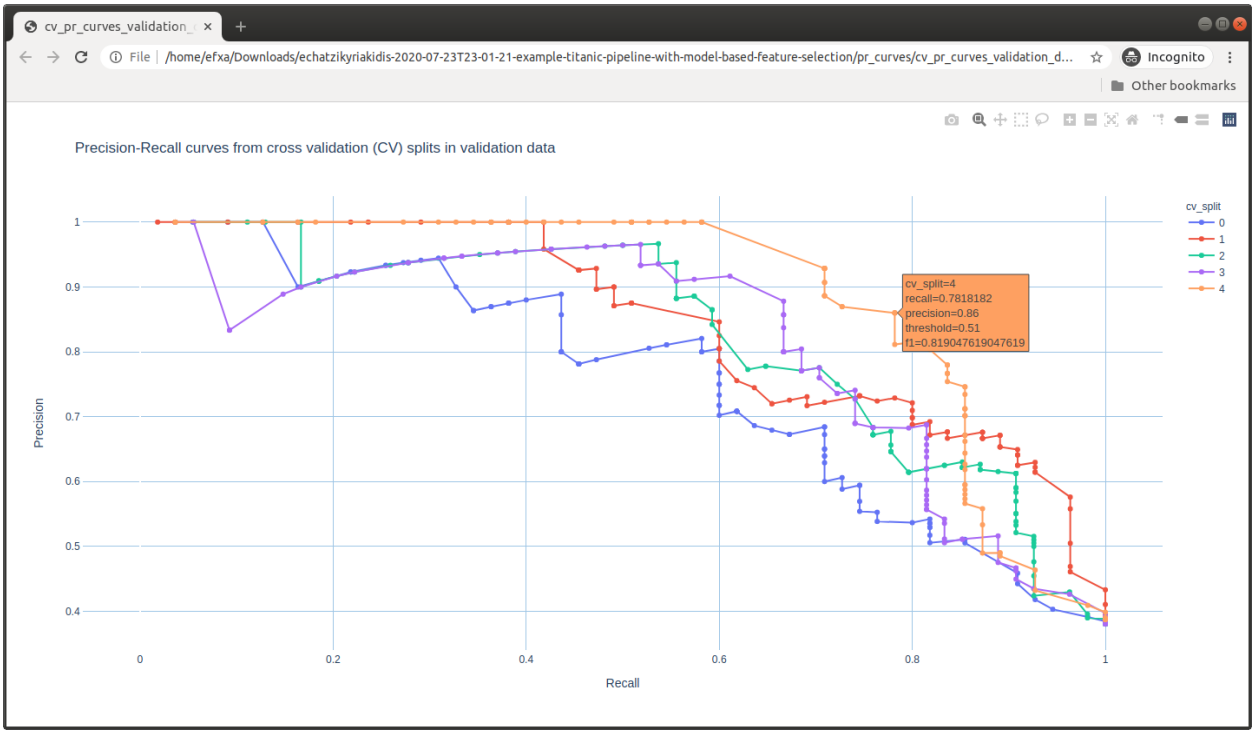

7.6 ROC Curves

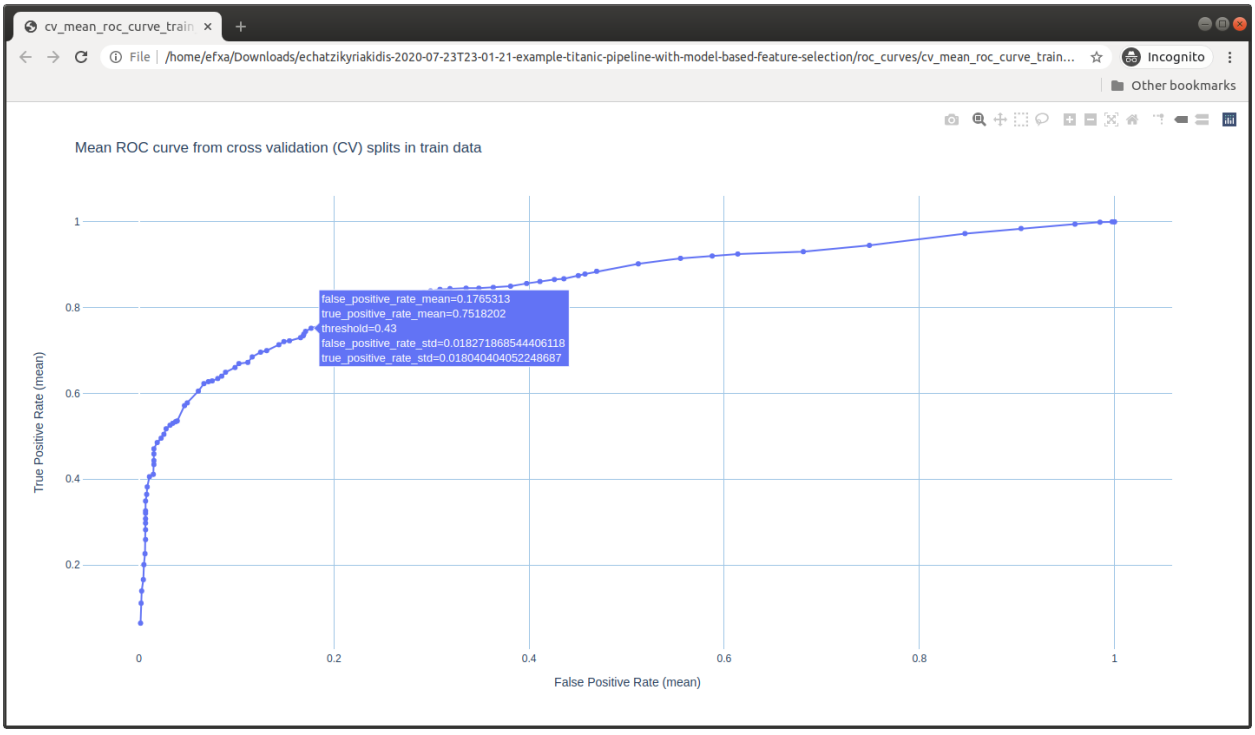

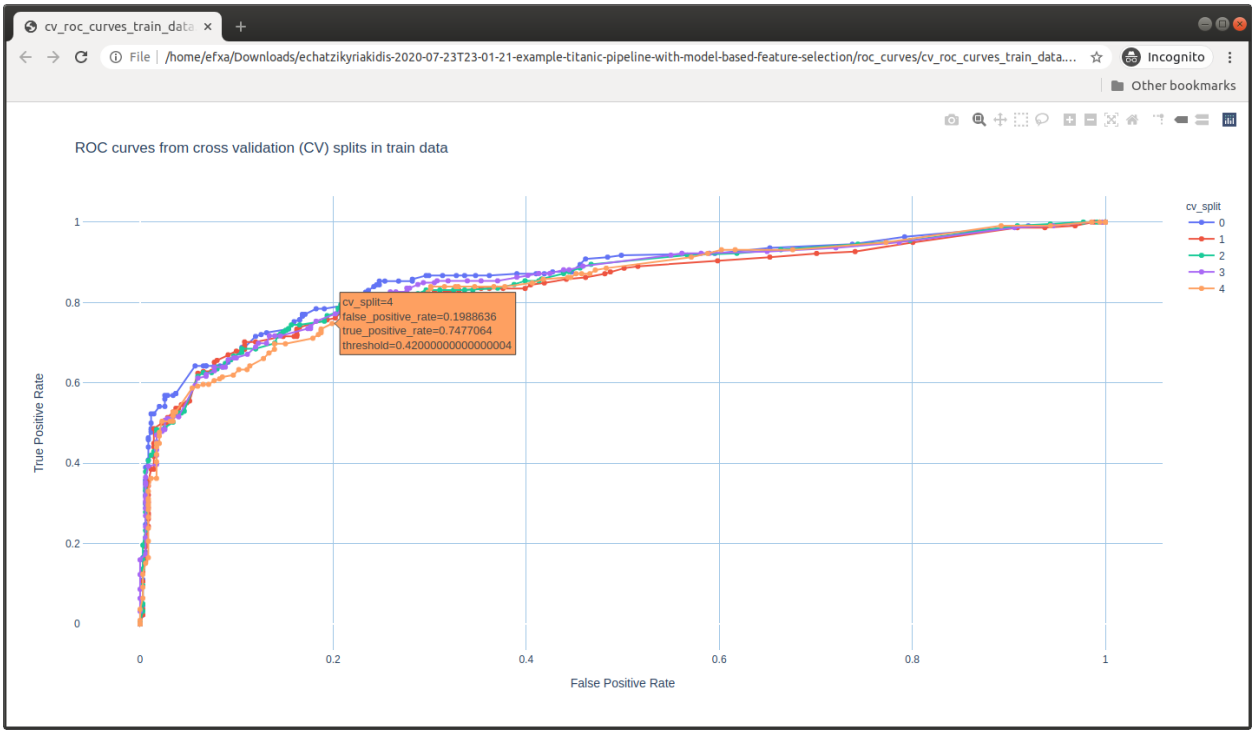

7.7 Performance Metrics

On train / validation CV folds:

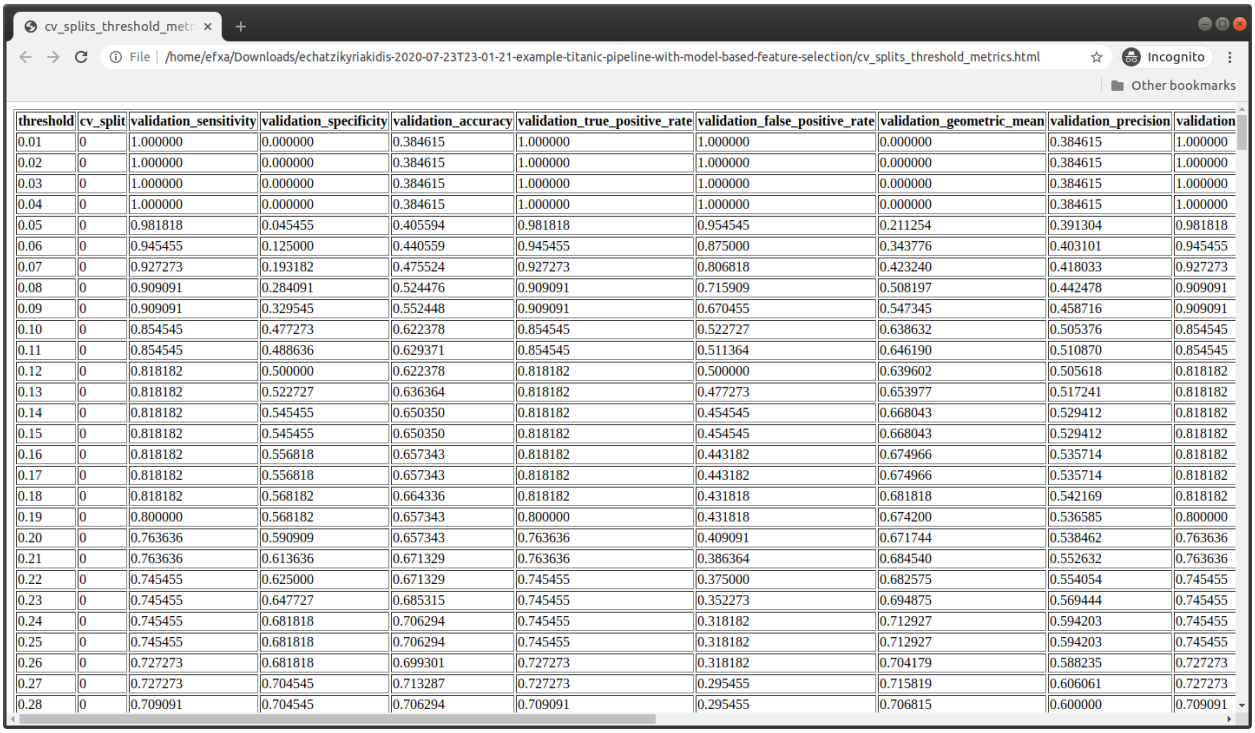

threshold	cv_split	validation_sensitivity	validation_specificity	validation_accuracy	validation_true_positive_rate	validation_false_positive_rate	validation_geometric_mean	validation_precision	validation
0.01	0	1.000000	0.000000	0.384615	1.000000	1.000000	0.000000	0.384615	1.000000
0.02	0	1.000000	0.000000	0.384615	1.000000	1.000000	0.000000	0.384615	1.000000
0.03	0	1.000000	0.000000	0.384615	1.000000	1.000000	0.000000	0.384615	1.000000
0.04	0	1.000000	0.000000	0.384615	1.000000	1.000000	0.000000	0.384615	1.000000
0.05	0	0.981818	0.045455	0.405594	0.981818	0.954545	0.211254	0.391304	0.981818
0.06	0	0.945455	0.125000	0.440559	0.945455	0.875000	0.343776	0.403101	0.945455
0.07	0	0.927273	0.193182	0.475524	0.927273	0.806818	0.423240	0.418033	0.927273
0.08	0	0.909091	0.284091	0.524476	0.909091	0.715909	0.508197	0.442478	0.909091
0.09	0	0.909091	0.329545	0.552448	0.909091	0.670455	0.547345	0.458716	0.909091
0.10	0	0.854545	0.477273	0.622378	0.854545	0.522727	0.638632	0.505376	0.854545
0.11	0	0.854545	0.488636	0.629371	0.854545	0.511364	0.646190	0.510870	0.854545
0.12	0	0.818182	0.500000	0.622378	0.818182	0.500000	0.639602	0.505618	0.818182
0.13	0	0.818182	0.522727	0.636364	0.818182	0.477273	0.653977	0.517241	0.818182
0.14	0	0.818182	0.545455	0.650350	0.818182	0.454545	0.668043	0.529412	0.818182
0.15	0	0.818182	0.545455	0.650350	0.818182	0.454545	0.668043	0.529412	0.818182
0.16	0	0.818182	0.556818	0.657343	0.818182	0.443182	0.674966	0.535714	0.818182
0.17	0	0.818182	0.556818	0.657343	0.818182	0.443182	0.674966	0.535714	0.818182
0.18	0	0.818182	0.568182	0.664336	0.818182	0.431818	0.681818	0.542169	0.818182
0.19	0	0.800000	0.568182	0.657343	0.800000	0.431818	0.674200	0.536585	0.800000
0.20	0	0.763636	0.590909	0.657343	0.763636	0.409091	0.671744	0.538462	0.763636
0.21	0	0.763636	0.613636	0.671329	0.763636	0.386364	0.684540	0.552632	0.763636
0.22	0	0.745455	0.625000	0.671329	0.745455	0.375000	0.682575	0.554054	0.745455
0.23	0	0.745455	0.647727	0.685315	0.745455	0.352273	0.694875	0.569444	0.745455
0.24	0	0.745455	0.681818	0.706294	0.745455	0.318182	0.712927	0.594203	0.745455
0.25	0	0.745455	0.681818	0.706294	0.745455	0.318182	0.712927	0.594203	0.745455
0.26	0	0.727273	0.681818	0.699301	0.727273	0.318182	0.704179	0.588235	0.727273
0.27	0	0.727273	0.704545	0.713287	0.727273	0.295455	0.715819	0.606061	0.727273
0.28	0	0.709091	0.704545	0.706294	0.709091	0.295455	0.706815	0.600000	0.709091

On hold-out test set:

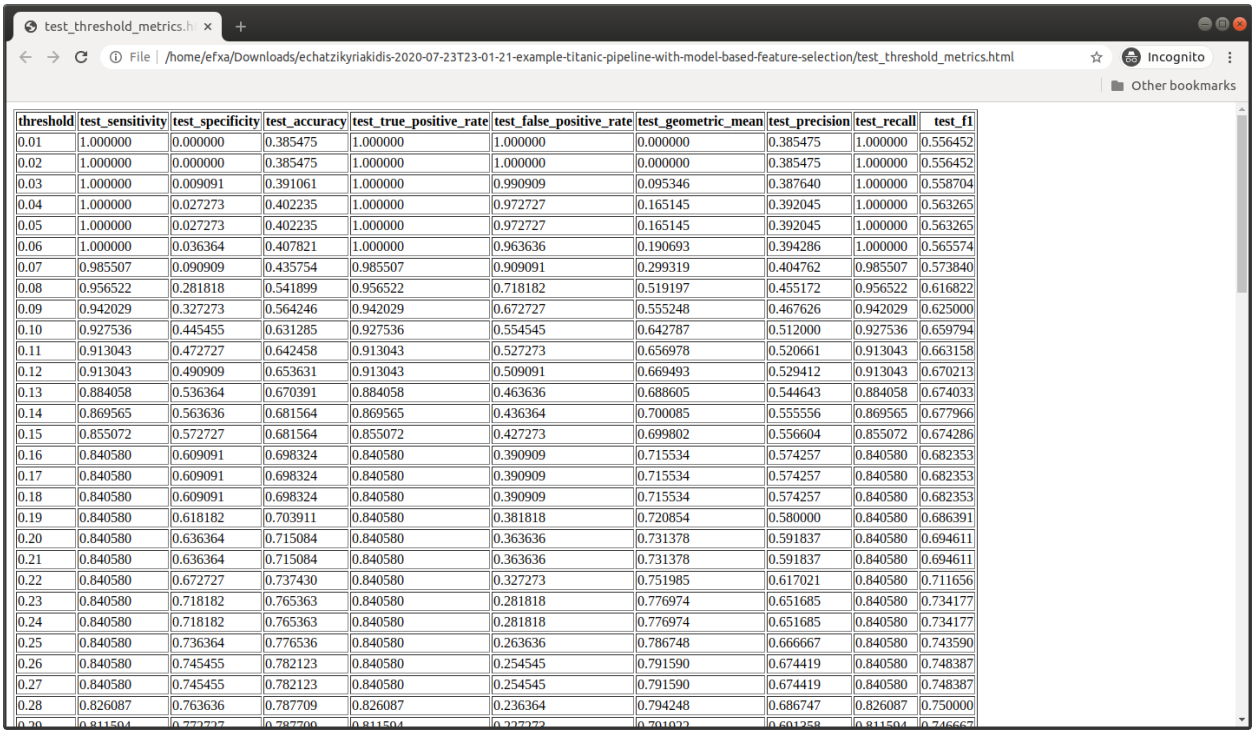

test_threshold_metrics.h ×

File | /home/efxa/Downloads/echatzikyriakidis-2020-07-23T23-01-21-example-titanic-pipeline-with-model-based-feature-selection/test_threshold_metrics.html

Incognito

Other bookmarks

threshold	test_sensitivity	test_specificity	test_accuracy	test_true_positive_rate	test_false_positive_rate	test_geometric_mean	test_precision	test_recall	test_f1
0.01	1.000000	0.000000	0.385475	1.000000	1.000000	0.000000	0.385475	1.000000	0.556452
0.02	1.000000	0.000000	0.385475	1.000000	1.000000	0.000000	0.385475	1.000000	0.556452
0.03	1.000000	0.009091	0.391061	1.000000	0.990909	0.095346	0.387640	1.000000	0.558704
0.04	1.000000	0.027273	0.402235	1.000000	0.972727	0.165145	0.392045	1.000000	0.563265
0.05	1.000000	0.027273	0.402235	1.000000	0.972727	0.165145	0.392045	1.000000	0.563265
0.06	1.000000	0.036364	0.407821	1.000000	0.963636	0.190693	0.394286	1.000000	0.565574
0.07	0.985507	0.090909	0.435754	0.985507	0.909091	0.299319	0.404762	0.985507	0.573840
0.08	0.956522	0.281818	0.541899	0.956522	0.718182	0.519197	0.455172	0.956522	0.616822
0.09	0.942029	0.327273	0.564246	0.942029	0.672727	0.555248	0.467626	0.942029	0.625000
0.10	0.927536	0.445455	0.631285	0.927536	0.554545	0.642787	0.512000	0.927536	0.659794
0.11	0.913043	0.472727	0.642458	0.913043	0.527273	0.656978	0.520661	0.913043	0.663158
0.12	0.913043	0.490909	0.653631	0.913043	0.509091	0.669493	0.529412	0.913043	0.670213
0.13	0.884058	0.536364	0.670391	0.884058	0.463636	0.688605	0.544643	0.884058	0.674033
0.14	0.869565	0.563636	0.681564	0.869565	0.436364	0.700085	0.555556	0.869565	0.677966
0.15	0.855072	0.572727	0.681564	0.855072	0.427273	0.699802	0.556604	0.855072	0.674286
0.16	0.840580	0.609091	0.698324	0.840580	0.390909	0.715534	0.574257	0.840580	0.682353
0.17	0.840580	0.609091	0.698324	0.840580	0.390909	0.715534	0.574257	0.840580	0.682353
0.18	0.840580	0.609091	0.698324	0.840580	0.390909	0.715534	0.574257	0.840580	0.682353
0.19	0.840580	0.618182	0.703911	0.840580	0.381818	0.720854	0.580000	0.840580	0.686391
0.20	0.840580	0.636364	0.715084	0.840580	0.363636	0.731378	0.591837	0.840580	0.694611
0.21	0.840580	0.636364	0.715084	0.840580	0.363636	0.731378	0.591837	0.840580	0.694611
0.22	0.840580	0.672727	0.737430	0.840580	0.327273	0.751985	0.617021	0.840580	0.711656
0.23	0.840580	0.718182	0.765363	0.840580	0.281818	0.776974	0.651685	0.840580	0.734177
0.24	0.840580	0.718182	0.765363	0.840580	0.281818	0.776974	0.651685	0.840580	0.734177
0.25	0.840580	0.736364	0.776536	0.840580	0.263636	0.786748	0.666667	0.840580	0.743590
0.26	0.840580	0.745455	0.782123	0.840580	0.254545	0.791590	0.674419	0.840580	0.748387
0.27	0.840580	0.745455	0.782123	0.840580	0.254545	0.791590	0.674419	0.840580	0.748387
0.28	0.826087	0.763636	0.787709	0.826087	0.236364	0.794248	0.686747	0.826087	0.750000
0.29	0.811594	0.772727	0.787709	0.811594	0.227272	0.791922	0.691358	0.811594	0.746667

7.8 Hyperparameters Search Results

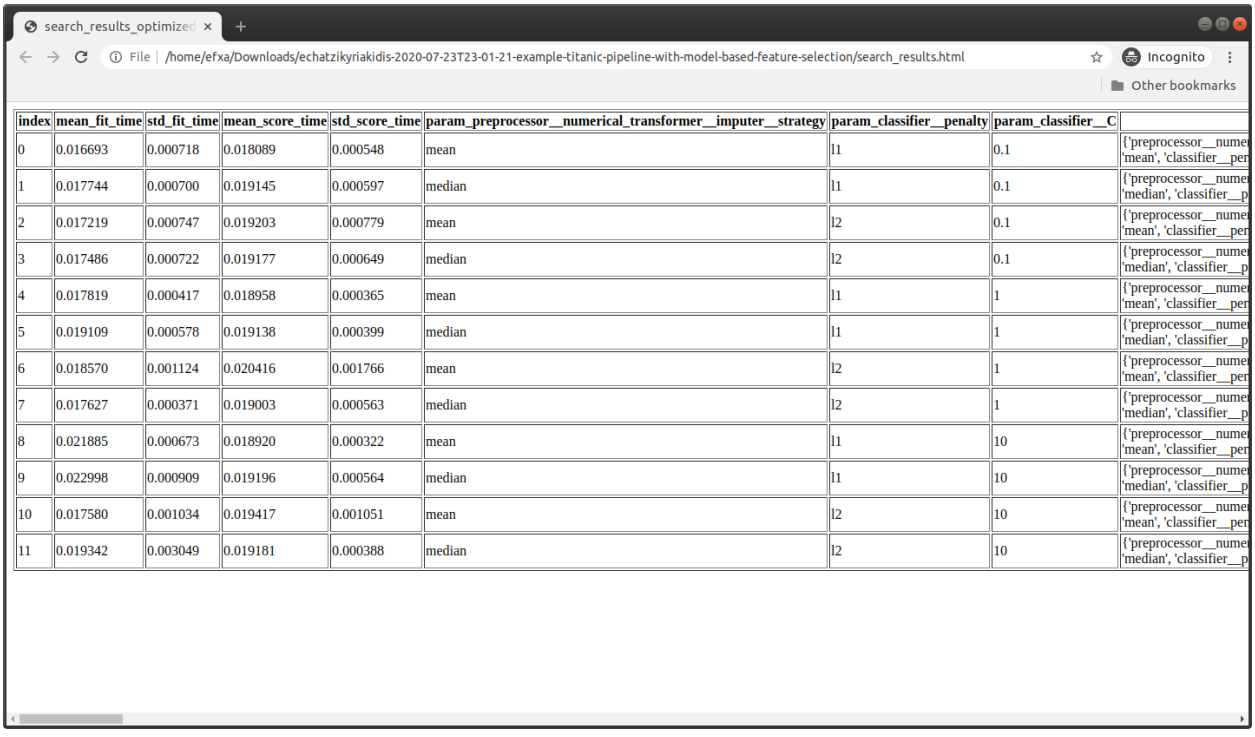

search_results_optimized ×

File | /home/efxa/Downloads/echatzikyriakidis-2020-07-23T23-01-21-example-titanic-pipeline-with-model-based-feature-selection/search_results.html

Incognito

Other bookmarks

index	mean_fit_time	std_fit_time	mean_score_time	std_score_time	param_preprocessor__numerical_transformer__imputer__strategy	param_classifier__penalty	param_classifier__C	
0	0.016693	0.000718	0.018089	0.000548	mean	l1	0.1	{'preprocessor__numer 'mean', 'classifier__per
1	0.017744	0.000700	0.019145	0.000597	median	l1	0.1	{'preprocessor__numer 'median', 'classifier__p
2	0.017219	0.000747	0.019203	0.000779	mean	l2	0.1	{'preprocessor__numer 'mean', 'classifier__per
3	0.017486	0.000722	0.019177	0.000649	median	l2	0.1	{'preprocessor__numer 'median', 'classifier__p
4	0.017819	0.000417	0.018958	0.000365	mean	l1	1	{'preprocessor__numer 'mean', 'classifier__per
5	0.019109	0.000578	0.019138	0.000399	median	l1	1	{'preprocessor__numer 'median', 'classifier__p
6	0.018570	0.001124	0.020416	0.001766	mean	l2	1	{'preprocessor__numer 'mean', 'classifier__per
7	0.017627	0.000371	0.019003	0.000563	median	l2	1	{'preprocessor__numer 'median', 'classifier__p
8	0.021885	0.000673	0.018920	0.000322	mean	l1	10	{'preprocessor__numer 'mean', 'classifier__per
9	0.022998	0.000909	0.019196	0.000564	median	l1	10	{'preprocessor__numer 'median', 'classifier__p
10	0.017580	0.001034	0.019417	0.001051	mean	l2	10	{'preprocessor__numer 'mean', 'classifier__per
11	0.019342	0.003049	0.019181	0.000388	median	l2	10	{'preprocessor__numer 'median', 'classifier__p

7.9 Task Parameters Logging

evaluation_cross_validation_task.params

```
{
 "type": "evaluation_cross_validation_task",
 "estimator": "Pipeline(steps=[('preprocessor',\n                 ColumnTransformer(transformers=[('numerical_transformer',\n
 "train_data_set_file_path": "https://bit.ly/titanic-data-train",
 "test_data_set_file_path": "https://bit.ly/titanic-data-test",
 "estimator_params": {
  "preprocessor__numerical_transformer__imputer__strategy": "mean",
  "classifier__penalty": "l1",
  "classifier__C": 10.0
 },
 "field_delimiter": ",",
 "feature_columns": "all",
 "id_column": "PassengerId",
 "label_column": "Survived",
 "random_seed": 42,
 "threshold_selection_by": "f1",
 "metric_greater_is_better": true,
 "threshold_tuning_range": [
  0.01,
  1.0,
  0.01
 ],
 "export_classification_reports": true,
 "export_confusion_matrixes": true,
 "export_roc_curves": true,
 "export_pr_curves": true,
 "export_false_positives_reports": true,
 "export_false_negatives_reports": true,
 "export_also_for_train_folds": true,
 "fscore_beta": 1,
 "fold_options": {
  "total_folds": 5,
  "shuffle": true
 },
 "fold_method": "stratified"
}
```

Plain Text Tab Width: 8 Ln 1, Col 1 INS

feature_selection_cross_validation_task.params

```
{
 "type": "feature_selection_cross_validation_task",
 "estimator": "LogisticRegression(random_state=42, solver='liblinear')",
 "train_data_set_file_path": "https://bit.ly/titanic-data-train",
 "estimator_params": null,
 "field_delimiter": ",",
 "preprocessor": "ColumnTransformer(transformers=[('numerical_transformer',\n                                                  Pipeline(steps=[('imputer', Si
 "preprocessor_params": null,
 "min_features_to_select": 4,
 "scoring": "f1",
 "feature_columns": "all",
 "id_column": "PassengerId",
 "label_column": "Survived",
 "random_seed": 42,
 "verbose": 3,
 "n_jobs": 1,
 "fold_options": {
  "total_folds": 5,
  "shuffle": true
 },
 "fold_method": "stratified"
}
```

Plain Text Tab Width: 8 Ln 1, Col 1 INS


Open
hyper_parameters_search_cross_validation_task.params
Save
{
"type": "hyper_parameters_search_cross_validation_task",
"estimator": "Pipeline(steps=[('preprocessor',\n ColumnTransformer(transformers=[('numerical_transformer',\n
"search_params": {
"classifier__C": [
0.1,
1.0,
10.0
],
"classifier__penalty": [
"l1",
"l2"
],
"preprocessor__numerical_transformer__imputer__strategy": [
"mean",
"median"
]
},
"train_data_set_file_path": "https://bit.ly/titanic-data-train",
"estimator_params": null,
"field_delimiter": ",",
"scorers": [
"roc_auc",
"average_precision",
"f1",
"precision",
"recall",
"accuracy"
],
"feature_columns": "all",
"id_column": "PassengerId",
"label_column": "Survived",
"objective_score": "f1",
"random_seed": 42,
"verbose": 3,
"n_jobs": 1,
"return_train_score": true,
"fold_options": {
Plain Text
Tab Width: 8
Ln 1, Col 1
INS


Open
train_task.params
Save
{
"type": "train_task",
"estimator": "Pipeline(steps=[('preprocessor',\n ColumnTransformer(transformers=[('numerical_transformer',\n
"train_data_set_file_path": "https://bit.ly/titanic-data-train",
"estimator_params": {
"preprocessor__numerical_transformer__imputer__strategy": "mean",
"classifier__penalty": "l1",
"classifier__C": 10.0
},
"field_delimiter": ",",
"feature_columns": "all",
"id_column": "PassengerId",
"label_column": "Survived",
"random_seed": 42
}
Plain Text
Tab Width: 8
Ln 1, Col 1
INS

prediction_task.params

```
{
 "type": "prediction_task",
 "estimator": "Pipeline(steps=[('preprocessor',\n                 ColumnTransformer(transformers=[('numerical_transformer',\n
 "data_set_file_path": "https://bit.ly/titanic-data-new",
 "field_delimiter": ",",
 "feature_columns": "all",
 "id_column": "PassengerId",
 "prediction_column": "Survived",
 "threshold": 0.36
}
```

7.10 Experiment Logging

experiment.log

```
{
 "datetime": "2020-07-23T23-01-21",
 "experimenter": "echatzikyriakidis",
 "experiment_id": "dc1095b704f546debd615928bcf7a994"
}
```

7.11 Features Selected

The selected column indexes from the transformed features (this is generated when a preprocessor is used):

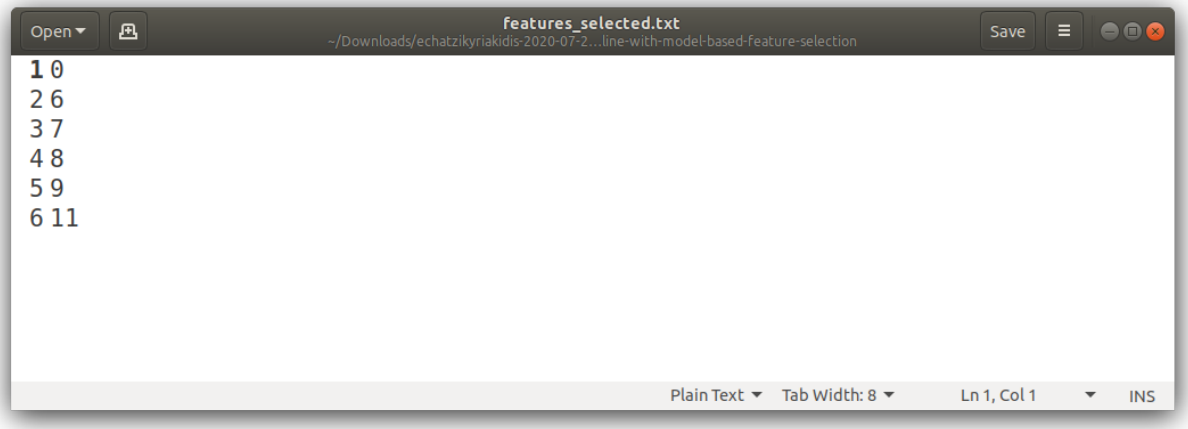

The selected column names from the original features (this is generated when no preprocessor is used):

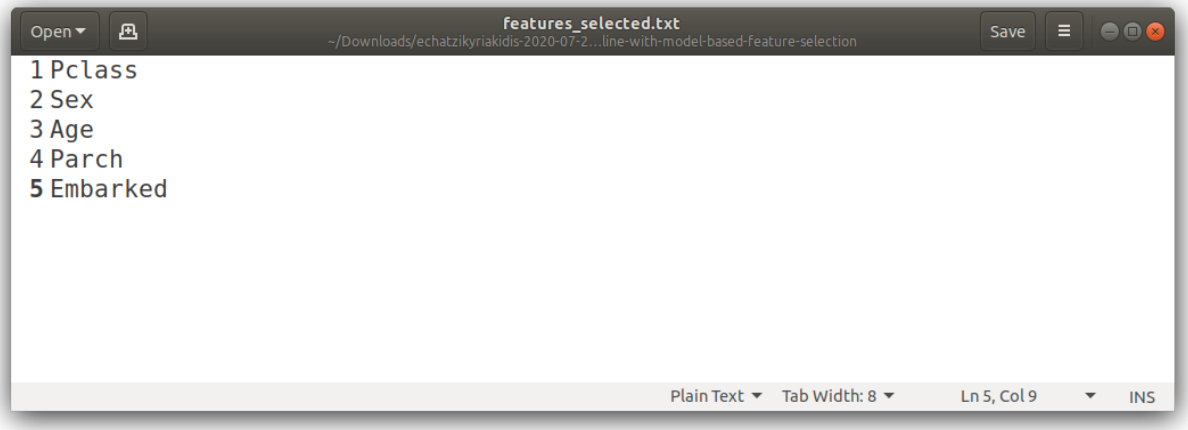

7.12 Experiment Source Code

example_titanic_pipeline_with_model_based_feature_selection.py

```python
from sklearn.compose import ColumnTransformer
from sklearn.pipeline import Pipeline
from sklearn.impute import SimpleImputer
from sklearn.preprocessing import StandardScaler, OneHotEncoder
from sklearn.linear_model import LogisticRegression

from skrobot.core import Experiment
from skrobot.tasks import TrainTask
from skrobot.tasks import PredictionTask
from skrobot.tasks import FeatureSelectionCrossValidationTask
from skrobot.tasks import EvaluationCrossValidationTask
from skrobot.tasks import HyperParametersSearchCrossValidationTask
from skrobot.feature_selection import ColumnSelector
from skrobot.notification import BaseNotifier

######### Initialization Code

train_data_set_file_path = 'https://bit.ly/titanic-data-train'

test_data_set_file_path = 'https://bit.ly/titanic-data-test'

new_data_set_file_path = 'https://bit.ly/titanic-data-new'

random_seed = 42

id_column = 'PassengerId'

label_column = 'Survived'

numerical_features = ['Age', 'Fare', 'SibSp', 'Parch']

categorical_features = ['Embarked', 'Sex', 'Pclass']

numeric_transformer = Pipeline(steps=[
    ('imputer', SimpleImputer()),
    ('scaler', StandardScaler())])

categorical_transformer = Pipeline(steps=[
```

7.13 Predictions

predictions.csv

```
PassengerId,Survived,probability
530,0,0.24575978907773116
760,1,0.8367357017875539
532,0,0.13299699284500366
227,0,0.244909831257034
884,0,0.244909831257034
637,0,0.1333657497174666
718,1,0.7307918212381642
278,0,0.2361056710832674
```

CHAPTER

EIGHT

THE PEOPLE BEHIND IT?

Development:

- Efstathios Chatzikyriakidis

Support, testing and features recommendation:

- Lefteris Kouloubris
- Antonis Markou
- Christina Chrysouli
- Michalis Chaviaras

And last but not least, all the open-source contributors whose work went into RELEASES.

CHAPTER

NINE

CAN I CONTRIBUTE?

Of course, the project is Free Software and you can contribute to it!

CHAPTER

TEN

WHAT LICENSE DO YOU USE?

See our LICENSE for more details.

PYTHON MODULE INDEX

S

INDEX

Symbols

B

C

E

F

G

T