



# Optunity Documentation

*Release 1.1.0*

**Marc Claesen, Jaak Simm and Dusan Popovic**

September 30, 2015



<b>1 Example</b>	<b>3</b>
<b>2 Quick setup</b>	<b>5</b>
<b>3 Developer Guide</b>	<b>7</b>
<b>4 Contributors</b>	<b>9</b>
<b>5 Indices and tables</b>	<b>11</b>
5.1 Installing Optunity . . . . .	11
5.2 User Guide . . . . .	13
5.3 Examples . . . . .	38
5.4 Notebooks . . . . .	83
5.5 Using Other Environments . . . . .	116
5.6 Optunity API . . . . .	121
<b>Bibliography</b>	<b>129</b>
<b>Python Module Index</b>	<b>131</b>





---

Optunity is a library containing various optimizers for hyperparameter tuning. Hyperparameter tuning is a recurrent problem in many machine learning tasks, both supervised and unsupervised. This package provides several distinct approaches to solve such problems including some helpful facilities such as cross-validation and a plethora of score functions.

- [Installation](#)
- [User guide](#)
- [Report a problem](#)

#### Obtaining Optunity

- at [PyPI](#) (releases)
- at [GitHub](#) (development)

From an optimization point of view, the tuning problem can be considered as follows: the objective function is non-convex, non-smooth and typically expensive to evaluate. Tuning examples include optimizing regularization or kernel parameters.

The Optunity library is implemented in Python and allows straightforward integration in other machine learning environments. Optunity is currently also supported in [R](#), [MATLAB](#), [GNU Octave](#) and Java through Jython.

If you have any problems, comments or suggestions you can get in touch with us at [gitter](#): Optunity is free software, using a BSD-style license.



---

## Example

---

As a simple example of Optunity's features, the code below demonstrates how to tune an SVM with RBF kernel using Optunity and scikit-learn. This involves optimizing the hyperparameters `gamma` and `C`:

```
import optunity
import optunity.metrics
import sklearn.svm

# score function: twice iterated 10-fold cross-validated accuracy
@optunity.cross_validated(x=data, y=labels, num_folds=10, num_iter=2)
def svm_auc(x_train, y_train, x_test, y_test, logC, logGamma):
    model = sklearn.svm.SVC(C=10 ** logC, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    return optunity.metrics.roc_auc(y_test, decision_values)

# perform tuning
hps, _, _ = optunity.maximize(svm_auc, num_evals=200, logC=[-5, 2], logGamma=[-5, 1])

# train model on the full training set with tuned hyperparameters
optimal_model = sklearn.svm.SVC(C=10 ** hps['logC'], gamma=10 ** hps['logGamma']).fit(data, labels)
```

For more examples, please see our [Examples](#).



### Quick setup

---

Issue the following commands to get started on Linux:

```
git clone https://github.com/claesem/optunity.git  
export PYTHONPATH=$PYTHONPATH:$pwd/optunity/
```

Afterwards, importing optunity should work in Python:

```
python -c 'import optunity'
```

For a proper installation, run the following:

```
python optunity/setup.py install
```

or, if you have pip:

```
pip install optunity
```

Installation may require superuser privileges.



---

**Developer Guide**

---

Opportunity API



## Contributors

---

Optunity is developed at the STADIUS lab of the dept. of electrical engineering at KU Leuven (ESAT). The main contributors to Optunity are:

### **Marc Claesen**

- Python package
- framework design & implementation
- solver implementation
- communication protocol design & implementation
- MATLAB wrapper
- Octave wrapper
- Python, MATLAB and Octave examples

### **Jaak Simm**

- communication protocol design
- R wrapper
- R examples



---

## Indices and tables

---

- genindex
- modindex
- search

## 5.1 Installing Optunity

The source can be obtained from git at <http://git.optunity.net> (recommended), releases can be obtained from <http://releases.optunity.net>. Optunity is compatible with Python 2.7 and 3.x. Note that you must have Python installed before Optunity can be used (all Linux distributions have this, but Windows requires explicit installation).

Obtaining the latest version of Optunity using git can be done as follows:

```
git clone https://github.com/claesem/optunity.git
```

Installation instructions per environment:

- [\*Installing Optunity for Python\*](#)
- [\*Installing Optunity for MATLAB\*](#)
- [\*Installing Optunity for Octave\*](#)
- [\*Installing Optunity for R\*](#)
- [\*Installing Optunity for Java\*](#)

If you encounter any difficulties during installation, please open an issue at <https://github.com/claesem/optunity/issues>.

---

**Note:** Optunity has soft dependencies on [NumPy](#) and [/DEAP2012](#) for the [CMA-ES](#) solver. If these Python libraries are unavailable, the CMA-ES solver will be unavailable.

---

### 5.1.1 Installing Optunity for Python

If you want to use Optunity in another environment, this is not required. Optunity can be installed as a typical Python package, for example:

- Add Optunity's root directory (the one you cloned/extracted) to your PYTHONPATH environment variable. Note that this will only affect your current shell. To make it permanent, add this line to your .bashrc (Linux) or adjust the PYTHONPATH variable manually (Windows).

Extending PYTHONPATH in Linux (not permanent):

```
export PYTHONPATH=$PYTHONPATH:/path/to/optunity/
```

Extending PYTHONPATH in Windows (not permanent):

```
set PYTHONPATH=%PYTHONPATH%;/path/to/optunity
```

- Install using Optunity's setup script:

```
python optunity/setup.py install [--home=/desired/installation/directory/]
```

- Install using pip:

```
pip install optunity [-t /desired/installation/directory/]
```

After these steps, you should be able to import optunity module:

```
python -c 'import optunity'
```

### 5.1.2 Installing Optunity for MATLAB

To install Optunity for MATLAB, you must add `<optunity>/wrappers/matlab/` and its subdirectories to your MATLAB path. You can set the path in *Preferences -> Path* or using the following commands:

```
addpath(genpath('/path/to/optunity/wrappers/matlab/'));  
savepath
```

After these steps, you should be able to run the examples in `<optunity>/wrappers/matlab/optunity_example.m`:

```
optunity_example
```

**Warning:** The MATLAB wrapper requires the entire directory structure of Optunity to remain as is. If you only copy the `<optunity>/wrappers/matlab` subdirectory it will not work.

### 5.1.3 Installing Optunity for Octave

Optunity requires sockets for communication. In Octave, please install the *sockets* package first (available at Octave-Forge) and at <http://octave.sourceforge.net/sockets/>.

To install Optunity for Octave, you must add `<optunity>/wrappers/octave/` and its subdirectories to your Octave path:

```
addpath(genpath('/path/to/optunity/wrappers/octave/'));  
savepath
```

After these steps, you should be able to run the examples in `<optunity>/wrappers/octave/optunity/example/optunity_example.m`:

```
optunity_example
```

**Warning:** The Octave wrapper requires the entire directory structure of Optunity to remain as is. If you only copy the `<optunity>/wrappers/octave` subdirectory it will not work.

## 5.1.4 Installing Optunity for R

First install all dependencies. In R, issue the following command:

```
install.packages(c("rjson", "ROCR", "enrichvs", "plyr"))
```

Subsequently, clone the git repository and then issue the following commands:

```
cd optunity/wrappers
R CMD build R/
R CMD INSTALL optunity_<version number>.tar.gz
```

## 5.1.5 Installing Optunity for Java

Optunity is available for Java through Jython (v2.7+). To use Optunity via Jython the Python package must be installed first (see above).

## 5.2 User Guide

Optunity provides a variety of solvers for hyperparameter tuning problems. A tuning problem is specified by an objective function that provides a score for some tuple of hyperparameters. Specifying the objective function must be done by the user. The software offers a diverse set of solvers to optimize the objective function. A solver determines a good set of hyperparameters.

### Jump to

- [Solver overview](#)
- [Cross-validation](#)
- [Quality metrics](#)
- [Domain constraints](#)
- [Structured search spaces](#)

Optunity consists of a set of core functions that are offered in each environment, which we will now discuss briefly. Clicking on a function will take you to its Python API documentation. If you are using a different environment, you can still get the general idea on the Python pages. To dive into code details straight away, please consult the [Optunity API](#).

If you want to learn by example, please consult our [Examples](#) which use various features of Optunity to cover a wide variety of tuning tasks. In case of confusion, we provide a list of basic [Terminology](#).

In the rest of this section we will discuss the main API functions. We will start with very simple functions that offer basic functionality which should meet the needs of most use cases. Subsequently we will introduce the expert interface functions which have more bells and whistles that can be configured.

A variety of solvers is available, discussed in more detail [Solver overview](#). Optunity additionally provides [Cross-validation](#) and several [Quality metrics](#).

### Simple interface

For beginning users, we offer a set of functions with simple arguments. These functions should be enough for most of your needs. In case these functions are insufficient, please refer to the expert functions listed below or to submodules.

- **`optunity.maximize()`: maximizes the objective function** Adheres to a prespecified upper bound on the number of function evaluations. The solution will be within given box constraints. Optunity determines the best solver and its configuration for you.
- **`optunity.minimize()`: minimizes the objective function** Adheres to a prespecified upper bound on the number of function evaluations. The solution will be within given box constraints. Optunity determines the best solver and its configuration for you.
- **`optunity.maximize_structured()`: maximizes the objective function with a structured search space** This function extends the functionality of `optunity.maximize` by allowing you to specify a structured search space, that is a search space where the existence of some hyperparameters are contingent upon some discrete choices. Adheres to a prespecified upper bound on the number of function evaluations. The solution will be within given box constraints. Optunity determines the best solver and its configuration for you.
- **`optunity.minimize_structured()`: minimizes the objective function with a structured search space** This function extends the functionality of `optunity.minimize` by allowing you to specify a structured search space, that is a search space where the existence of some hyperparameters are contingent upon some discrete choices. Adheres to a prespecified upper bound on the number of function evaluations. The solution will be within given box constraints. Optunity determines the best solver and its configuration for you.
- **`optunity.manual()`: prints a basic manual (general or solver specific)** Prints a basic manual of Optunity and a list of all registered solvers. If a solver name is specified, its manual will be shown.

---

**Note:** You may alternatively consult the solver API documentation at [Solvers](#) for more details.

---

- **`optunity.cross_validated()`: decorator to perform k-fold cross-validation** Wrap a function with cross-validation functionality. The way cross-validation is performed is highly configurable, including support for strata and clusters. More details are available [here](#).

### Expert interface

The following functions are recommended for more advanced use of Optunity. This part of the API allows you to fully configure every detail about the provided solvers. In addition to more control in configuration, you can also perform parallel function evaluations via this interface (turned off by default due to problems in IPython).

- **`optunity.suggest_solver()`: suggests a solver and its configuration** Suggests a solver and configuration for a given tuning problem, based on the permitted number of function evaluations and box constraints.
- **`optunity.make_solver()`: constructs one of Optunity's registered solvers.** See the solver-specific manuals for more information per solver.
- **`optunity.optimize()`: optimizes an objective function with given solver** Some solvers are capable of vector evaluations. By default, the optimization is done through sequential function evaluations but this can be parallelized by specifying an appropriate pmap argument (cfr. `optunity.parallel.pmap()`).

More details about our solvers can be found [here](#). To learn about adding constraints to the objective function, please consult [this page](#).

#### 5.2.1 Solver overview

We currently recommend using [Particle Swarm Optimization](#) (our default). Based on our experience this is the most reliable solver across different learning algorithms. If you consistently get great results with a solver/algorithm combination, we are happy to hear about your experiences.

You can specify which solver you want to use in `optunity.maximize()` and `optunity.minimize()`, but only limited configuration is possible. If you want to specify detailed settings for a solver, you can use the expert interface, specifically `optunity.make_solver()` in combination with `optunity.optimize()`.

The following solvers are available in Optunity:

## Grid Search

This solver is implemented in `optunity.solvers.GridSearch`. It is available in `optunity.make_solver()` as ‘grid search’.

Grid search is an undirected search method that consists of testing a predefined set of values per hyperparameter. A search grid is constructed that is the Cartesian product of these sets of values.

Grid search can be used to tune a limited number of hyperparameters, as the size of the search grid (number of evaluations) increases exponentially in terms of the number of hyperparameters. We generally recommend the use of directed solvers over grid search, as they waste less time exploring uninteresting regions.

When you use choose this solver in `optunity.maximize()` or `optunity.minimize()`, an equal number of values is determined per hyperparameter (spread uniformly). The number of test values per hyperparameter is determined based on the number of permitted evaluations. If you want to specify the grid points manually, this is possible via `optunity.make_solver()` and `optunity.optimize()`.

## Example

Assume we have two hyperparameters  $x$  and  $y$ . For  $x$  we want to test the following values  $\{0, 1, 2, 3\}$  and for  $y$  we will try  $\{-10, -5, 0, 5, 10\}$ . The search grid consists of all possible pairs ( $4 \times 5 = 20$  in this case), e.g.:

x	y
0	-10
0	-5
0	0
0	5
0	10
1	-10
1	-5
1	0
1	5
1	10
2	-10
2	-5
2	0
2	5
2	10
3	-10
3	-5
3	0
3	5
3	10

## Random Search

This solver is implemented in `opportunity.solvers.RandomSearch`. It is available in `opportunity.make_solver()` as ‘random search’.

This strategy consists of testing a predefined number of randomly sampled hyperparameter tuples. Sampling is done uniform at random within specified box constraints.

This solver implements the search strategy described in [RAND].

## Particle Swarm Optimization

This solver is implemented in `opportunity.solvers.ParticleSwarm`. It is available in `opportunity.make_solver()` as ‘particle swarm’.

Particle swarm optimization (PSO) is a heuristic optimization technique. It simulates a set of particles (candidate solutions) that are moving around in the search-space [PSO2010], [PSO2002].

In the context of hyperparameter search, the position of a particle represents a set of hyperparameters and its movement is influenced by the goodness of the objective function value.

PSO is an iterative algorithm:

1. **Initialization:** a set of particles is initialized with random positions and initial velocities. The initialization step is essentially equivalent to Random Search.
2. **Iteration:** every particle’s position is updated based on its velocity, the particle’s historically best position and the entire swarm’s historical optimum.

```

1 for k ← 1 to ngen do
2   for i ← 1 to npart do
3     |   vi ← vi + φ1 · (p̄i - x̄i) + φ2 · (p̄g - x̄i)
4     |   x̄i ← x̄i + v̄i
5     |   if f(x̄i) is better than f(p̄i) then p̄i ← x̄i
6   end
7   p̄g ← best({p̄i | i = 1 : npart

```

Fig. 5.1: Particle swarm iterations:

- $\vec{x}_i$  is a particle’s position,
- $\vec{v}_i$  its velocity,
- $\vec{p}_i$  its historically best position,
- $\vec{p}_g$  is the swarm’s optimum,
- $\phi_1$  and  $\phi_2$  are vectors of uniformly sampled values in  $(0, \phi_1)$  and  $(0, \phi_2)$ , respectively.

PSO has 5 parameters that can be configured (see `opportunity.solvers.ParticleSwarm`):

- `num_particles`: the number of particles to use
- `num_generations`: the number of generations (iterations)
- `phi1`: the impact of each particle’s historical best on its movement
- `phi2`: the impact of the swarm’s optimum on the movement of each particle
- `max_speed`: an upper bound for  $\vec{v}_i$

The number of function evaluations that will be performed is  $num\_particles * num\_generations$ . A high number of particles focuses on global, undirected search (just like Random Search), whereas a high number of generations leads to more localized search since all particles will have time to converge.

Bibliographic references:

### Nelder-Mead simplex

This solver is implemented in `optunity.solvers.NelderMead`. It is available in `optunity.make_solver()` as ‘nelder-mead’.

This is a heuristic, nonlinear optimization method based on the concept of a simplex, originally introduced by Nelder and Mead [[NELDERMEAD](#)]. We implemented the version as described on [Wikipedia](#).

This method requires an initial starting point  $x_0$ . It is a good local search method, but will get stuck in bad regions when a poor starting point is specified.

### CMA-ES

This solver is implemented in `optunity.solvers.CMA_ES`. It is available in `optunity.make_solver()` as ‘cma-es’.

CMA-ES stands for Covariance Matrix Adaptation Evolutionary Strategy. This is an evolutionary strategy for continuous function optimization. It can dynamically adapt its search resolution per hyperparameter, allowing for efficient searches on different scales. More information is available in [[HANSEN2001](#)].

Opportunity’s implementation of this solver is done using the DEAP toolbox [[DEAP2012](#)]. This, in turn, requires NumPy. Both dependencies must be met to use this solver.

Bibliographic references:

### Tree-structured Parzen Estimator

This solver is implemented in `optunity.solvers.TPE`. It is available in `optunity.make_solver()` as ‘TPE’.

The Tree-structured Parzen Estimator (TPE) is a sequential model-based optimization (SMBO) approach. SMBO methods sequentially construct models to approximate the performance of hyperparameters based on historical measurements, and then subsequently choose new hyperparameters to test based on this model.

The TPE approach models  $P(x|y)$  and  $P(y)$  where  $x$  represents hyperparameters and  $y$  the associated quality score.  $P(x|y)$  is modeled by transforming the generative process of hyperparameters, replacing the distributions of the configuration prior with non-parametric densities. In this solver, Opportunity only supports uniform priors within given box constraints. For more exotic search spaces, please refer to [[Hyperopt](#)]. This optimization approach is described in detail in [[TPE2011](#)] and [[TPE2013](#)].

Opportunity provides the TPE solver as is implemented in [[Hyperopt](#)]. This solver is only available if Hyperopt is installed, which in turn requires NumPy. Both dependencies must be met to use this solver.

### Sobol sequences

A Sobol sequence is a low discrepancy quasi-random sequence. Sobol sequences were designed to cover the unit hypercube with lower discrepancy than completely random sampling (e.g. [Random Search](#)). Opportunity supports Sobol sequences in up to 40 dimensions (e.g. 40 hyperparameters).

The figures below show the differences between a Sobol sequence and sampling uniformly at random. These figures can be recreated using the code in `bin/examples/python/sobol_vs_random.py`.

The mathematical details on Sobol sequences are available in the following papers: [[SOBOL](#)], [[SOBOL2](#)], [[ANTONOV](#)], [[BRATLEY](#)], [[FOX](#)].

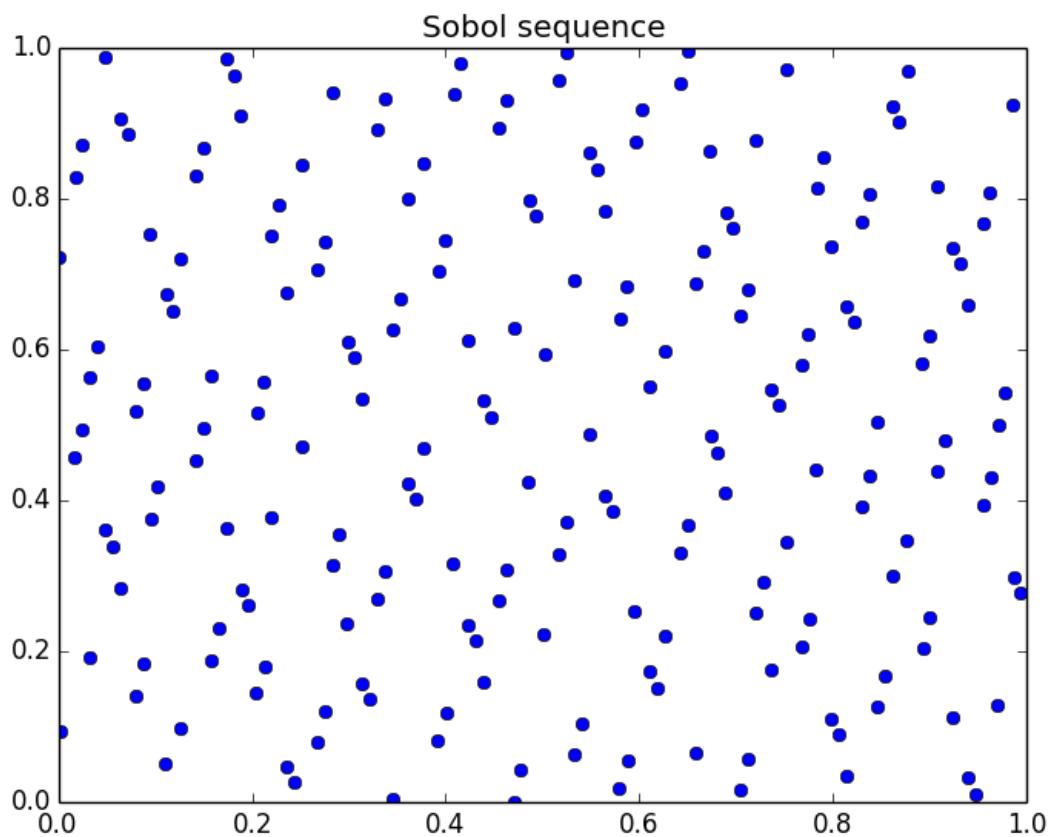


Fig. 5.2: 200 points sampled in 2D with a Sobol sequence.

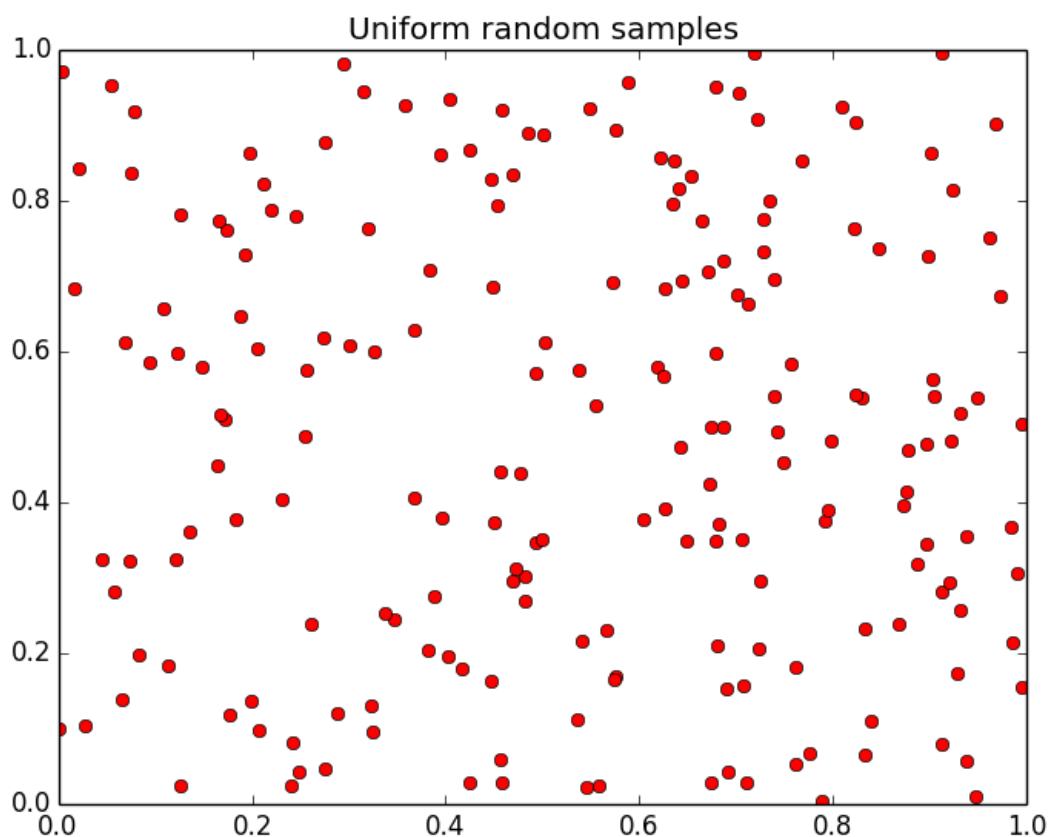


Fig. 5.3: 200 points sampled in 2D uniformly at random.

Opportunity's default solver is Particle Swarm Optimization.

Grid Search, Random Search and Sobol sequences are completely undirected algorithms and consequently not very efficient. Of these three, Sobol sequences is most efficient as uses a low-discrepancy quasirandom sequence.

Nelder-Mead simplex works well for objective functions that are smooth, unimodal and not too noisy (it is good for local search when you have a good idea about optimal regions for your hyperparameters).

For general searches, Particle Swarm Optimization and CMA-ES are most robust. Finally, the Tree-structured Parzen Estimator solver exposes Hyperopt's TPE solver in Opportunity's familiar API.

### 5.2.2 Cross-validation

Opportunity offers a simple interface to k-fold cross-validation. This is a statistical approach to measure a model's generalization performance. In the context of hyperparameter search, cross-validation is used to estimate the performance of a hyperparameter tuple. The cross-validation routines we offer are optional and can be replaced by comparable routines from other packages or some other method to estimate generalization performance.

The fold generation procedure in Opportunity allows for iterated cross-validation and is aware of both strata (data instances that must be spread across folds) and clusters (sets of instances that must assigned to a single fold). Please refer to `opportunity.cross_validated()` for implementation and API details.

We will build examples step by step. The basic setup is a `train` and `predict` function along with some data to construct folds over:

```
from __future__ import print_function
import opportunity as opt

def train(x, y, filler=''):
    print(filler + 'Training data:')
    for instance, label in zip(x, y):
        print(filler + str(instance) + ' ' + str(label))

def predict(x, filler=''):
    print(filler + 'Testing data:')
    for instance in x:
        print(filler + str(instance))

data = list(range(9))
labels = [0] * 9
```

The recommended way to perform cross-validation is using the `opportunity.cross_validation.cross_validated()` function decorator. To use it, you must specify an objective function. This function should contain the logic that is placed in the inner loop in cross-validation (e.g. train a model, predict test set, compute score), with the following signature: `f(x_train, y_train, x_test, y_test, hyperpar_1, hyperpar_2, ...)` (argument names are important):

- `x_train`: training data
- `y_train`: training labels (optional)
- `x_test`: test data
- `y_test`: test labels (optional)
- `hyperparameter names`: the hyperparameters that must be optimized

The `cross_validated` decorator takes care of generating folds, partitioning the data, iterating over folds and aggregating partial results. After decoration, the arguments `x_train`, `y_train`, `x_test` and `y_test` will be bound (e.g. the decorated

function does not take these as arguments). The decorated function will have hyperparameters as (keyword) arguments and returns a cross-validation result.

A simple code example:

```
@opt.cross_validated(x=data, y=labels, num_folds=3)
def cved(x_train, y_train, x_test, y_test):
    train(x_train, y_train)
    predict(x_test)
    return 0.0

cved()
```

## Reusing folds

If you want to compare different aspects of the learning approach (learning algorithms, score function, ...), it is a good idea to use the same cross-validation folds. This is very easy by using the *cross\_validated* decorator without syntactic sugar. Lets say we want to compare an SVM with RBF kernel and polynomial kernel with the same cross-validation configuration:

```
import sklearn.svm as svm

def svm_rbf(x_train, y_train, x_test, y_test, C, gamma):
    model = svm.SVC(kernel='rbf', C=C, gamma=gamma).fit(x_train, y_train)
    y_pred = model.predict(x_test)
    return opt.score_functions.accuracy(y_test, y_pred)

def svm_poly(x_train, y_train, x_test, y_test, C, d):
    model = svm.SVC(kernel='poly', C=C, degree=d).fit(x_train, y_train)
    y_pred = model.predict(x_test)
    return opt.score_functions.accuracy(y_test, y_pred)

cv_decorator = opt.cross_validated(x=data, y=labels, num_folds=3)

svm_rbf_cv = cv_decorator(svm_rbf)
svm_poly_cv = cv_decorator(svm_poly)
```

In this example, the function *svm\_rbf\_cv* takes keyword arguments *C* and *gamma* while *svm\_poly\_cv* takes *C* and *d*. Both perform cross-validation on the same data, using the same folds.

## Nested cross-validation

Nested cross-validation is a commonly used approach to estimate the generalization performance of a modeling process which includes model selection internally. A good summary is provided [here](#).

Nested cv consists of two cross-validation procedures wrapped around eachother. The inner cv is used for model selection, the outer cv estimates generalization performance.

This can be done in a straightforward manner using Opportunity:

```
@opt.cross_validated(x=data, y=labels, num_folds=3)
def nested_cv(x_train, y_train, x_test, y_test):

    @opt.cross_validated(x=x_train, y=y_train, num_folds=3)
    def inner_cv(x_train, y_train, x_test, y_test):
        train(x_train, y_train, ...)
        predict(x_test, ...)
```

```
    return 0.0

    inner_cv()
    predict(x_test)
    return 0.0

nested_cv()
```

The inner `opportunity.cross_validated()` decorator has access to the train and test folds generated by the outer procedure (`x_train` and `x_test`). For notational simplicity we assume a problem without labels here.

---

**Note:** The inner folds are regenerated in every iteration (since we are redefining `inner_cv` each time). The inner folds will therefore be different each time. The outer folds remain static, unless `regenerate_folds=True` is passed.

---

A complete example of nested cross-validation is available in [Basic: nested cross-validation](#).

## Aggregators

Opportunity's cross-validation implementation allows you to specify an *aggregator*. This is the function that will be used to compute the cross-validation result based on the results of individual folds. The default function is `mean`. You can specify any function to compute another measure if desired (for instance `min`, `max`, ...).

### Computing multiple performance measures during cross-validation

Sometimes it is desired to compute multiple performance measures using cross-validation. This is particularly useful for nested cross-validation. This is possible in Opportunity by letting the wrapped function return multiple scores and using the `opportunity.cross_validation.list_mean()` aggregator:

```
@opportunity.cross_validated(x=data, y=labels, num_folds=3,
                               aggregator=opportunity.cross_validation.list_mean)
def f(x_train, y_train, x_test, y_test):
    model = train(x_train, y_train)
    predictions = model.predict(x_test)
    score_1 = score_1(y_test, predictions)
    score_2 = score_2(y_test, predictions)
    return score_1, score_2
```

For even more flexibility, you can use `opportunity.cross_validation.identity()` as aggregator, which will return a list of return values for every cross-validation fold.

### 5.2.3 Quality metrics

Quality metrics (score/loss functions) are used to quantify the performance of a given model. Score functions are typically maximized (e.g. accuracy, concordance, ...) whereas loss functions should be minimized (e.g. mean squared error, error rate, ...). Opportunity provides common score/loss functions for your convenience.

We use the following calling convention:

- `y` (iterable): the true labels/function values
- `yhat` (iterable): the predicted labels/function values
- we assume `y` and `yhat` are of the same length (though we do not assert this).

- potential parameters of the score function must be added by keyword

All functions listed here are available in `opportunity.metrics`.

## Score functions

Score functions are typically maximized (e.g. `opportunity.maximize()`).

### Classification

Score	Associated Opportunity function
accuracy	<code>accuracy()</code>
area under ROC curve	<code>roc_auc()</code>
area under PR curve	<code>pr_auc()</code>
$F_\beta$	<code>fbeta()</code>
precision/PPV	<code>precision()</code>
recall/sensitivity	<code>recall()</code>
specificity/NPV	<code>npv()</code>
PU score	<code>pu_score()</code>

### Regression

Score	Associated Opportunity function
R squared	<code>r_squared()</code>

## Loss functions

Loss functions are typically minimized (e.g. `opportunity.minimize()`).

### Classification

Score	Associated Opportunity function
Brier score	<code>brier()</code>
error rate	<code>error_rate()</code>
log loss	<code>logloss()</code>

### Regression

Score	Associated Opportunity function
mean squared error	<code>mse()</code>
absolute error	<code>absolute_error()</code>

## 5.2.4 Domain constraints

Opportunity supports domain constraints on the objective function. Domain constraints are used to enforce solvers to remain within a prespecified search space. Most solvers that Opportunity provides are implicitly unconstrained (cfr.

Solver overview), though hyperparameters are usually constrained in some way (ex: regularization coefficients must be positive).

A set of simple constraints and facilities to use them are provided in [Domain constraints](#). Specifically, the following constraints are provided:

- `lb_{oc}`: assigns a lower bound (open or closed)
- `ub_{oc}`: assigns an upper bound (open or closed)
- `range_{oc}{oc}`: creates a range constraint (e.g.  $A < x < B$ )

All of the above constraints apply to a specific hyperparameter. Multidimensional constraints are possible, but you would need to implement them yourself (see [Implementing custom constraints](#)).

Note that the functions `opportunity.maximize()` and `opportunity.minimize()` wrap explicit box constraints around the objective function prior to starting the solving process. The expert function `opportunity.optimize()` does not do this for you, which allows more flexibility at the price of verbosity.

Constraint violations in Opportunity raise a `ConstraintViolation` exception by default. The usual way we handle these exceptions is by returning a certain (typically bad) default function value (using the `opportunity.constraints.violations_defaulted()` decorator). This will cause solvers to stop searching in the infeasible region.

To add a series of constraints, we recommend using the `opportunity.wrap_constraints()` function. This function takes care of assigning default values on constraint violations if desired.

### Implementing custom constraints

Constraints are implemented as a binary functions, which yield false in case of a constraint violation. You can design your own constraint according to this principle. For instance, assume we have a binary function with two arguments  $x$  and  $y$ :

```
def f(x, y):  
    return x + y
```

Opportunity provides all univariate constraints you need, but lets say we want to constrain the domain to be the unit circle in  $x,y$ -space. We can do this using the following constraint:

```
constraint = lambda x, y: (x ** 2 + y ** 2) <= 1.0
```

To constrain  $f$ , we use `opportunity.wrap_constraints()`:

```
fc = opportunity.wrap_constraints(f, custom=[constraint])
```

The constrained function  $fc(x, y)$  will yield  $x + y$  if the arguments are within the unit circle, or raise a `ConstraintViolation` exception otherwise.

### 5.2.5 Structured search spaces

Some hyperparameter optimization problems have a hierarchical nature, comprising discrete choices and depending on the choice additional hyperparameters may exist. A common example is optimizing a kernel, without choosing a certain family of kernel functions in advance (e.g. polynomial, RBF, linear, ...).

Opportunity provides the functions `opportunity.maximize_structured()` and `opportunity.minimize_structured()` for such structured search spaces. Structured search spaces can be specified as nested dictionaries, which generalize the standard way of specifying box constraints:

- hyperparameters within box constraints: specified as dictionary entries, where `key=parameter name` and `value=box constraints (list)`.
- discrete choices:** specified as a dictionary, where each entry represents a choice, that is `key=option name` and `value` has two possible values:
  - a new dictionary of conditional hyperparameters, following the same rules
  - `None`, to indicate a choice which doesn't imply further hyperparameterization

Structured search spaces can be nested to form any graph-like search space. It's worth noting that the addition of discrete choices naturally generalizes Opportunity's search space definition in `opportunity.minimize()` and `opportunity.maximize()`, since box constraints are specified as keyword arguments there, so Python's `kwargs` to these functions effectively follows the exact same structure, e.g.:

```
_ = opportunity.minimize(fun, num_evals=1, A=[0, 1], B=[-1, 2])
# kwargs = {A: [0, 1], B: [-1, 2]}
```

### Example: SVM kernel hyperparameterization

Suppose we want to optimize the kernel, choosing from the following options:

- linear kernel  $\kappa_{\text{linear}}(u, v) = u^T v$ : no additional hyperparameters
- polynomial kernel  $\kappa_{\text{poly}}(u, v) = (u^T v + \text{coef0})^{\text{degree}}$ : 2 hyperparameters (degree and coef0)
- RBF kernel  $\kappa_{\text{RBF}}(u, v) = \exp(-\gamma * |u - v|^2)$  1 hyperparameter (gamma)

When we put this all together, the SVM kernel search space can be defined as follows (Python syntax):

```
search = {'kernel': {'linear': None,
                     'rbf': {'gamma': [gamma_min, gamma_max]},
                     'poly': {'degree': [degree_min, degree_max],
                              'coef0': [coef0_min, coef0_max]}}
          }
```

The structure of the search space directly matches the hyperparameterization of every kernel function. We use `None` in the linear kernel as there are no additional hyperparameters. The hyperparameters of the RBF and polynomial kernel follow Opportunity's default syntax based on dictionnaries.

Opportunity also supports nested choices, for example an outer choice for the learning algorithm (e.g. SVM, naive Bayes,...) and an inner choice for the SVM kernel function. This is illustrated in the following notebook: [sklearn: automated learning method selection and tuning](#).

## 5.2.6 Terminology

To avoid confusion, here is some basic terminology we will use consistently throughout the documentation:

**hyperparameters** User-specified parameters for a given machine learning approach. These will serve as optimization variables in our context.

Example: kernel parameters.

**score** Some measure which quantifies the quality of a certain modeling approach (model type + hyperparameters).

Example: cross-validated accuracy of a classifier.

**objective function** The function that must be optimized. The arguments to this function are hyperparameters and the result is some score measure of a model constructed using these hyperparameters. Opportunity can minimize and maximize, depending on the requirements.

Example: accuracy of an SVM classifier as a function of kernel and regularization parameters.

**box constraints** Every hyperparameter of the tuning problem must be within a prespecified interval. The optimal solution will be within the hyperrectangle (box) specified by the ranges.

**solver** A strategy to optimize hyperparameters, such as *grid search*.

**train-predict-score (TPS) chain** A sequence of code which trains a model, uses it to predict an independent test set and then computes some score measure based on the predictions. TPS chains must be specified by the user as they depend entirely on the method that is being tuned and the evaluation criteria.

## Installing Opportunity

The source can be obtained from git at <http://git.opportunity.net> (recommended), releases can be obtained from <http://releases.opportunity.net>. Opportunity is compatible with Python 2.7 and 3.x. Note that you must have Python installed before Opportunity can be used (all Linux distributions have this, but Windows requires explicit installation).

Obtaining the latest version of Opportunity using git can be done as follows:

```
git clone https://github.com/claesem/opportunity.git
```

Installation instructions per environment:

- [Installing Opportunity for Python](#)
- [Installing Opportunity for MATLAB](#)
- [Installing Opportunity for Octave](#)
- [Installing Opportunity for R](#)
- [Installing Opportunity for Java](#)

If you encounter any difficulties during installation, please open an issue at <https://github.com/claesem/opportunity/issues>.

---

**Note:** Opportunity has soft dependencies on NumPy and [DEAP2012] for the CMA-ES solver. If these Python libraries are unavailable, the CMA-ES solver will be unavailable.

---

## Installing Opportunity for Python

If you want to use Opportunity in another environment, this is not required. Opportunity can be installed as a typical Python package, for example:

- Add Opportunity's root directory (the one you cloned/extracted) to your PYTHONPATH environment variable. Note that this will only affect your current shell. To make it permanent, add this line to your .bashrc (Linux) or adjust the PYTHONPATH variable manually (Windows).

Extending PYTHONPATH in Linux (not permanent):

```
export PYTHONPATH=$PYTHONPATH:/path/to/opportunity/
```

Extending PYTHONPATH in Windows (not permanent):

```
set PYTHONPATH=%PYTHONPATH%;/path/to/opportunity
```

- Install using Opportunity' setup script:

```
python opportunity/setup.py install [--home=/desired/installation/directory/]
```

- Install using pip:

```
pip install optunity [-t /desired/installation/directory/]
```

After these steps, you should be able to import optunity module:

```
python -c 'import optunity'
```

### Installing Optunity for MATLAB

To install Optunity for MATLAB, you must add `<optunity>/wrappers/matlab/` and its subdirectories to your MATLAB path. You can set the path in *Preferences -> Path* or using the following commands:

```
addpath(genpath('/path/to/optunity/wrappers/matlab/'));  
savepath
```

After these steps, you should be able to run the examples in `<optunity>/wrappers/matlab/optunity_example.m`:

```
optunity_example
```

**Warning:** The MATLAB wrapper requires the entire directory structure of Optunity to remain as is. If you only copy the `<optunity>/wrappers/matlab` subdirectory it will not work.

### Installing Optunity for Octave

Optunity requires sockets for communication. In Octave, please install the *sockets* package first (available at Octave-Forge) and at <http://octave.sourceforge.net/sockets/>.

To install Optunity for Octave, you must add `<optunity>/wrappers/octave/` and its subdirectories to your Octave path:

```
addpath(genpath('/path/to/optunity/wrappers/octave/'));  
savepath
```

After these steps, you should be able to run the examples in `<optunity>/wrappers/octave/optunity/example/optunity_example.m`:

```
optunity_example
```

**Warning:** The Octave wrapper requires the entire directory structure of Optunity to remain as is. If you only copy the `<optunity>/wrappers/octave` subdirectory it will not work.

### Installing Optunity for R

First install all dependencies. In R, issue the following command:

```
install.packages(c("rjson", "ROCR", "enrichvs", "plyr"))
```

Subsequently, clone the git repository and then issue the following commands:

```
cd optunity/wrappers  
R CMD build R/  
R CMD INSTALL optunity_<version number>.tar.gz
```

## Installing Opportunity for Java

Opportunity is available for Java through Jython (v2.7+). To use Opportunity via Jython the Python package must be installed first (see above).

## Solver overview

We currently recommend using **Particle Swarm Optimization** (our default). Based on our experience this is the most reliable solver across different learning algorithms. If you consistently get great results with a solver/algorithm combination, we are happy to hear about your experiences.

You can specify which solver you want to use in `opportunity.maximize()` and `opportunity.minimize()`, but only limited configuration is possible. If you want to specify detailed settings for a solver, you can use the expert interface, specifically `opportunity.make_solver()` in combination with `opportunity.optimize()`.

The following solvers are available in Opportunity:

### Grid Search

This solver is implemented in `opportunity.solvers.GridSearch`. It is available in `opportunity.make_solver()` as ‘grid search’.

Grid search is an undirected search method that consists of testing a predefined set of values per hyperparameter. A search grid is constructed that is the Cartesian product of these sets of values.

Grid search can be used to tune a limited number of hyperparameters, as the size of the search grid (number of evaluations) increases exponentially in terms of the number of hyperparameters. We generally recommend the use of directed solvers over grid search, as they waste less time exploring uninteresting regions.

When you use choose this solver in `opportunity.maximize()` or `opportunity.minimize()`, an equal number of values is determined per hyperparameter (spread uniformly). The number of test values per hyperparameter is determined based on the number of permitted evaluations. If you want to specify the grid points manually, this is possible via `opportunity.make_solver()` and `opportunity.optimize()`.

**Example** Assume we have two hyperparameters  $x$  and  $y$ . For  $x$  we want to test the following values  $[0, 1, 2, 3]$  and for  $y$  we will try  $[-10, -5, 0, 5, 10]$ . The search grid consists of all possible pairs ( $4 \times 5 = 20$  in this case), e.g.:

x	y
0	-10
0	-5
0	0
0	5
0	10
1	-10
1	-5
1	0
1	5
1	10
2	-10
2	-5
2	0
2	5
2	10
3	-10
3	-5
3	0
3	5
3	10

## Random Search

This solver is implemented in `opportunity.solvers.RandomSearch`. It is available in `opportunity.make_solver()` as ‘random search’.

This strategy consists of testing a predefined number of randomly sampled hyperparameter tuples. Sampling is done uniform at random within specified box constraints.

This solver implements the search strategy described in [RAND].

## Particle Swarm Optimization

This solver is implemented in `opportunity.solvers.ParticleSwarm`. It is available in `opportunity.make_solver()` as ‘particle swarm’.

Particle swarm optimization (PSO) is a heuristic optimization technique. It simulates a set of particles (candidate solutions) that are moving around in the search-space [PSO2010], [PSO2002].

In the context of hyperparameter search, the position of a particle represents a set of hyperparameters and its movement is influenced by the goodness of the objective function value.

PSO is an iterative algorithm:

1. **Initialization:** a set of particles is initialized with random positions and initial velocities. The initialization step is essentially equivalent to Random Search.
2. **Iteration:** every particle’s position is updated based on its velocity, the particle’s historically best position and the entire swarm’s historical optimum.

PSO has 5 parameters that can be configured (see `opportunity.solvers.ParticleSwarm`):

- `num_particles`: the number of particles to use
- `num_generations`: the number of generations (iterations)

```

1 for k ← 1 to  $n_{gen}$  do
2   for i ← 1 to  $n_{part}$  do
3      $\vec{v}_i \leftarrow \vec{v}_i + \phi_1 \cdot (\vec{p}_i - \vec{x}_i) + \phi_2 \cdot (\vec{p}_g - \vec{x}_i)$ 
4      $\vec{x}_i \leftarrow \vec{x}_i + \vec{v}_i$ 
5     if  $f(\vec{x}_i)$  is better than  $f(\vec{p}_i)$  then  $\vec{p}_i \leftarrow \vec{x}_i$ 
6   end
7    $\vec{p}_g \leftarrow \text{best}(\{\vec{p}_i \mid i = 1 : n_{part}\})$ 
8 end

```

Fig. 5.4: Particle swarm iterations:

- $\vec{x}_i$  is a particle's position,
- $\vec{v}_i$  its velocity,
- $\vec{p}_i$  its historically best position,
- $\vec{p}_g$  is the swarm's optimum,
- $\phi_1$  and  $\phi_2$  are vectors of uniformly sampled values in  $(0, \phi_1)$  and  $(0, \phi_2)$ , respectively.

- *phi1*: the impact of each particle's historical best on its movement
- *phi2*: the impact of the swarm's optimum on the movement of each particle
- *max\_speed*: an upper bound for  $\vec{v}_i$

The number of function evaluations that will be performed is  $num\_particles * num\_generations$ . A high number of particles focuses on global, undirected search (just like Random Search), whereas a high number of generations leads to more localized search since all particles will have time to converge.

Bibliographic references:

### Nelder-Mead simplex

This solver is implemented in `optunity.solvers.NelderMead`. It is available in `optunity.make_solver()` as ‘nelder-mead’.

This is a heuristic, nonlinear optimization method based on the concept of a simplex, originally introduced by Nelder and Mead [[NELDERMEAD](#)]. We implemented the version as described on [Wikipedia](#).

This method requires an initial starting point  $x_0$ . It is a good local search method, but will get stuck in bad regions when a poor starting point is specified.

### CMA-ES

This solver is implemented in `optunity.solvers.CMA_ES`. It is available in `optunity.make_solver()` as ‘cma-es’.

CMA-ES stands for Covariance Matrix Adaptation Evolutionary Strategy. This is an evolutionary strategy for continuous function optimization. It can dynamically adapt its search resolution per hyperparameter, allowing for efficient searches on different scales. More information is available in [[HANSEN2001](#)].

Optunity’s implementation of this solver is done using the DEAP toolbox [[DEAP2012](#)]. This, in turn, requires NumPy. Both dependencies must be met to use this solver.

Bibliographic references:

## Tree-structured Parzen Estimator

This solver is implemented in `opportunity.solvers.TPE`. It is available in `opportunity.make_solver()` as ‘TPE’.

The Tree-structured Parzen Estimator (TPE) is a sequential model-based optimization (SMBO) approach. SMBO methods sequentially construct models to approximate the performance of hyperparameters based on historical measurements, and then subsequently choose new hyperparameters to test based on this model.

The TPE approach models  $P(x|y)$  and  $P(y)$  where  $x$  represents hyperparameters and  $y$  the associated quality score.  $P(x|y)$  is modeled by transforming the generative process of hyperparameters, replacing the distributions of the configuration prior with non-parametric densities. In this solver, Opportunity only supports uniform priors within given box constraints. For more exotic search spaces, please refer to [\[Hyperopt\]](#). This optimization approach is described in detail in [\[TPE2011\]](#) and [\[TPE2013\]](#).

Opportunity provides the TPE solver as is implemented in [\[Hyperopt\]](#). This solver is only available if Hyperopt is installed, which in turn requires NumPy. Both dependencies must be met to use this solver.

## Sobol sequences

A Sobol sequence is a low discrepancy quasi-random sequence. Sobol sequences were designed to cover the unit hypercube with lower discrepancy than completely random sampling (e.g. [Random Search](#)). Opportunity supports Sobol sequences in up to 40 dimensions (e.g. 40 hyperparameters).

The figures below show the differences between a Sobol sequence and sampling uniformly at random. These figures can be recreated using the code in `bin/examples/python/sobol_vs_random.py`.

The mathematical details on Sobol sequences are available in the following papers: [\[SOBOL\]](#), [\[SOBOL2\]](#), [\[ANTONOV\]](#), [\[BRATLEY\]](#), [\[FOX\]](#).

Opportunity’s default solver is Particle Swarm Optimization.

[Grid Search](#), [Random Search](#) and [Sobol sequences](#) are completely undirected algorithms and consequently not very efficient. Of these three, [Sobol sequences](#) is most efficient as uses a low-discrepancy quasirandom sequence.

[Nelder-Mead simplex](#) works well for objective functions that are smooth, unimodal and not too noisy (it is good for local search when you have a good idea about optimal regions for your hyperparameters).

For general searches, [Particle Swarm Optimization](#) and [CMA-ES](#) are most robust. Finally, the [Tree-structured Parzen Estimator](#) solver exposes Hyperopt’s TPE solver in Opportunity’s familiar API.

## Cross-validation

Opportunity offers a simple interface to k-fold [cross-validation](#). This is a statistical approach to measure a model’s generalization performance. In the context of hyperparameter search, cross-validation is used to estimate the performance of a hyperparameter tuple. The cross-validation routines we offer are optional and can be replaced by comparable routines from other packages or some other method to estimate generalization performance.

The fold generation procedure in Opportunity allows for iterated cross-validation and is aware of both strata (data instances that must be spread across folds) and clusters (sets of instances that must be assigned to a single fold). Please refer to `opportunity.cross_validated()` for implementation and API details.

We will build examples step by step. The basic setup is a `train` and `predict` function along with some data to construct folds over:

```
from __future__ import print_function
import opportunity as opt
```

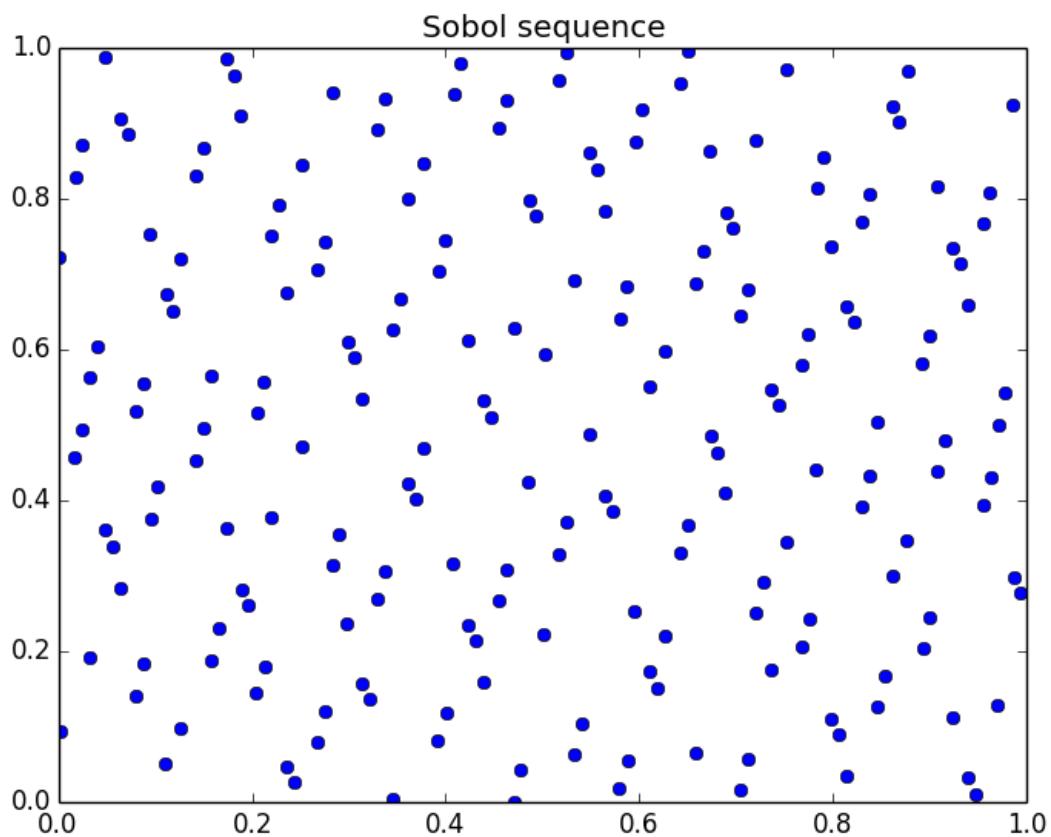


Fig. 5.5: 200 points sampled in 2D with a Sobol sequence.

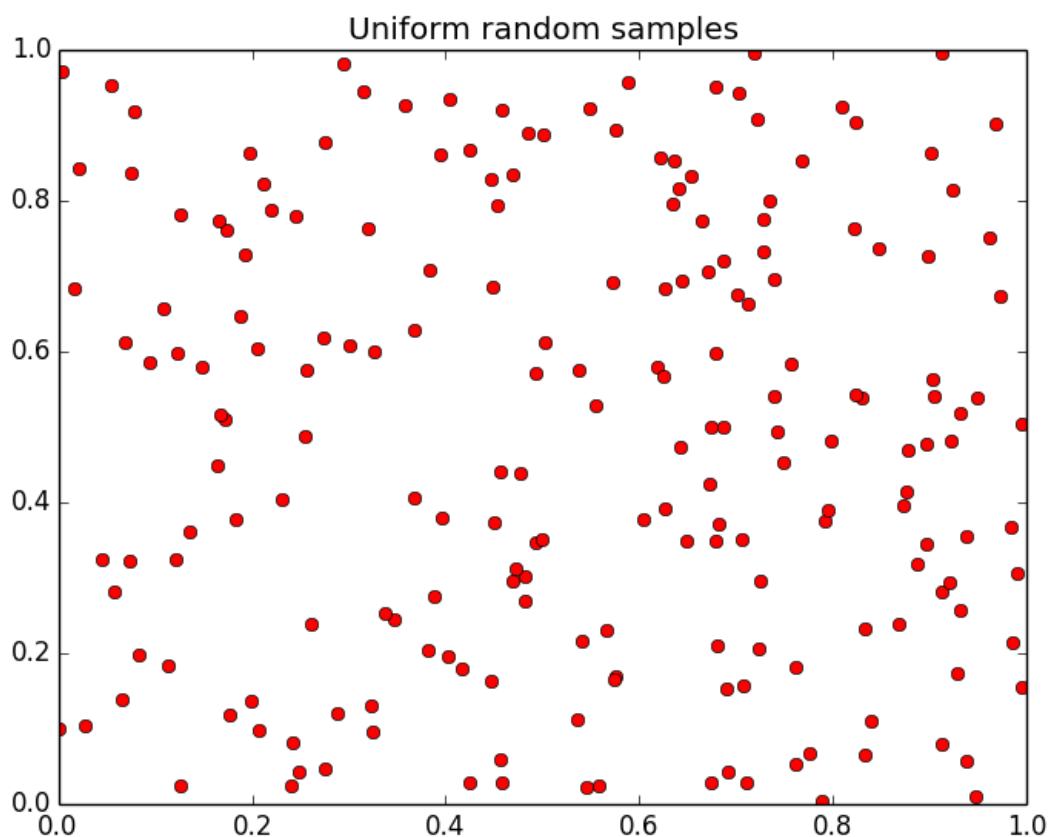


Fig. 5.6: 200 points sampled in 2D uniformly at random.

```

def train(x, y, filler=''):
    print(filler + 'Training data:')
    for instance, label in zip(x, y):
        print(filler + str(instance) + ' ' + str(label))

def predict(x, filler=''):
    print(filler + 'Testing data:')
    for instance in x:
        print(filler + str(instance))

data = list(range(9))
labels = [0] * 9

```

The recommended way to perform cross-validation is using the `opportunity.cross_validation.cross_validated()` function decorator. To use it, you must specify an objective function. This function should contain the logic that is placed in the inner loop in cross-validation (e.g. train a model, predict test set, compute score), with the following signature: `f(x_train, y_train, x_test, y_test, hyperpar_1, hyperpar_2, ...)` (argument names are important):

- `x_train`: training data
- `y_train`: training labels (optional)
- `x_test`: test data
- `y_test`: test labels (optional)
- *hyperparameter names*: the hyperparameters that must be optimized

The `cross_validated` decorator takes care of generating folds, partitioning the data, iterating over folds and aggregating partial results. After decoration, the arguments `x_train`, `y_train`, `x_test` and `y_test` will be bound (e.g. the decorated function does not take these as arguments). The decorated function will have hyperparameters as (keyword) arguments and returns a cross-validation result.

A simple code example:

```

@opt.cross_validated(x=data, y=labels, num_folds=3)
def cved(x_train, y_train, x_test, y_test):
    train(x_train, y_train)
    predict(x_test)
    return 0.0

cved()

```

### Reusing folds

If you want to compare different aspects of the learning approach (learning algorithms, score function, ...), it is a good idea to use the same cross-validation folds. This is very easy by using the `cross_validated` decorator without syntactic sugar. Lets say we want to compare an SVM with RBF kernel and polynomial kernel with the same cross-validation configuration:

```

import sklearn.svm as svm

def svm_rbf(x_train, y_train, x_test, y_test, C, gamma):
    model = svm.SVC(kernel='rbf', C=C, gamma=gamma).fit(x_train, y_train)
    y_pred = model.predict(x_test)
    return opt.score_functions.accuracy(y_test, y_pred)

def svm_poly(x_train, y_train, x_test, y_test, C, d):

```

```

model = svm.SVC(kernel='poly', C=C, degree=d).fit(x_train, y_train)
y_pred = model.predict(x_test)
return opt.score_functions.accuracy(y_test, y_pred)

cv_decorator = opt.cross_validated(x=data, y=labels, num_folds=3)

svm_rbf_cv = cv_decorator(svm_rbf)
svm_poly_cv = cv_decorator(svm_poly)

```

In this example, the function `svm_rbf_cv` takes keyword arguments `C` and `gamma` while `svm_poly_cv` takes `C` and `d`. Both perform cross-validation on the same data, using the same folds.

## Nested cross-validation

Nested cross-validation is a commonly used approach to estimate the generalization performance of a modeling process which includes model selection internally. A good summary is provided [here](#).

Nested cv consists of two cross-validation procedures wrapped around each other. The inner cv is used for model selection, the outer cv estimates generalization performance.

This can be done in a straightforward manner using Opportunity:

```

@opt.cross_validated(x=data, y=labels, num_folds=3)
def nested_cv(x_train, y_train, x_test, y_test):

    @opt.cross_validated(x=x_train, y=y_train, num_folds=3)
    def inner_cv(x_train, y_train, x_test, y_test):
        train(x_train, y_train, ...)
        predict(x_test, ...)
        return 0.0

    inner_cv()
    predict(x_test)
    return 0.0

nested_cv()

```

The inner `opportunity.cross_validated()` decorator has access to the train and test folds generated by the outer procedure (`x_train` and `x_test`). For notational simplicity we assume a problem without labels here.

---

**Note:** The inner folds are regenerated in every iteration (since we are redefining `inner_cv` each time). The inner folds will therefore be different each time. The outer folds remain static, unless `regenerate_folds=True` is passed.

A complete example of nested cross-validation is available in [Basic: nested cross-validation](#).

## Aggregators

Opportunity's cross-validation implementation allows you to specify an *aggregator*. This is the function that will be used to compute the cross-validation result based on the results of individual folds. The default function is `mean`. You can specify any function to compute another measure if desired (for instance `min`, `max`, ...).

**Computing multiple performance measures during cross-validation** Sometimes it is desired to compute multiple performance measures using cross-validation. This is particularly useful for nested cross-

validation. This is possible in Opportunity by letting the wrapped function return multiple scores and using the `opportunity.cross_validation.list_mean()` aggregator:

```
@opportunity.cross_validated(x=data, y=labels, num_folds=3,
                             aggregator=opportunity.cross_validation.list_mean)
def f(x_train, y_train, x_test, y_test):
    model = train(x_train, y_train)
    predictions = model.predict(x_test)
    score_1 = score_1(y_test, predictions)
    score_2 = score_2(y_test, predictions)
    return score_1, score_2
```

For even more flexibility, you can use `opportunity.cross_validation.identity()` as aggregator, which will return a list of return values for every cross-validation fold.

## Domain constraints

Opportunity supports domain constraints on the objective function. Domain constraints are used to enforce solvers to remain within a prespecified search space. Most solvers that Opportunity provides are implicitly unconstrained (cfr. [Solver overview](#)), though hyperparameters are usually constrained in some way (ex: regularization coefficients must be positive).

A set of simple constraints and facilities to use them are provided in [Domain constraints](#). Specifically, the following constraints are provided:

- `lb_{oc}`: assigns a lower bound (open or closed)
- `ub_{oc}`: assigns an upper bound (open or closed)
- `range_{oc}{oc}`: creates a range constraint (e.g.  $A < x < B$ )

All of the above constraints apply to a specific hyperparameter. Multidimensional constraints are possible, but you would need to implement them yourself (see [Implementing custom constraints](#)).

Note that the functions `opportunity.maximize()` and `opportunity.minimize()` wrap explicit box constraints around the objective function prior to starting the solving process. The expert function `opportunity.optimize()` does not do this for you, which allows more flexibility at the price of verbosity.

Constraint violations in Opportunity raise a `ConstraintViolation` exception by default. The usual way we handle these exceptions is by returning a certain (typically bad) default function value (using the `opportunity.constraints.violations_defaulted()` decorator). This will cause solvers to stop searching in the infeasible region.

To add a series of constraints, we recommend using the `opportunity.wrap_constraints()` function. This function takes care of assigning default values on constraint violations if desired.

## Implementing custom constraints

Constraints are implemented as a binary functions, which yield false in case of a constraint violation. You can design your own constraint according to this principle. For instance, assume we have a binary function with two arguments  $x$  and  $y$ :

```
def f(x, y):
    return x + y
```

Opportunity provides all univariate constraints you need, but lets say we want to constrain the domain to be the unit circle in  $x,y$ -space. We can do this using the following constraint:

```
constraint = lambda x, y: (x ** 2 + y ** 2) <= 1.0
```

To constrain  $f$ , we use `opportunity.wrap_constraints()`:

```
fc = opportunity.wrap_constraints(f, custom=[constraint])
```

The constrained function  $fc(x, y)$  will yield  $x + y$  if the arguments are within the unit circle, or raise a *ConstraintViolation* exception otherwise.

## Structured search spaces

Some hyperparameter optimization problems have a hierarchical nature, comprising discrete choices and depending on the choice additional hyperparameters may exist. A common example is optimizing a kernel, without choosing a certain family of kernel functions in advance (e.g. polynomial, RBF, linear, ...).

Opportunity provides the functions `opportunity.maximize_structured()` and `opportunity.minimize_structured()` for such structured search spaces. Structured search spaces can be specified as nested dictionaries, which generalize the standard way of specifying box constraints:

- hyperparameters within box constraints: specified as dictionary entries, where `key=parameter name` and `value=box constraints (list)`.
- **discrete choices: specified as a dictionary, where each entry represents a choice, that is `key=option name` and `value` has two cases**
  - a new dictionary of conditional hyperparameters, following the same rules
  - *None*, to indicate a choice which doesn't imply further hyperparameterization

Structured search spaces can be nested to form any graph-like search space. It's worth noting that the addition of discrete choices naturally generalizes Opportunity's search space definition in `opportunity.minimize()` and `opportunity.maximize()`, since box constraints are specified as keyword arguments there, so Python's `kwargs` to these functions effectively follows the exact same structure, e.g.:

```
_ = opportunity.minimize(fun, num_evals=1, A=[0, 1], B=[-1, 2])
# kwargs = {A: [0, 1], B: [-1, 2]}
```

## Example: SVM kernel hyperparameterization

Suppose we want to optimize the kernel, choosing from the following options:

- linear kernel  $\kappa_{linear}(u, v) = u^T v$ : no additional hyperparameters
- polynomial kernel  $\kappa_{poly}(u, v) = (u^T v + \text{coef0})^{\text{degree}}$ : 2 hyperparameters (degree and coef0)
- RBF kernel  $\kappa_{RBF}(u, v) = \exp(-\gamma * |u - v|^2)$  1 hyperparameter (gamma)

When we put this all together, the SVM kernel search space can be defined as follows (Python syntax):

```
search = {'kernel': {'linear': None,
                     'rbf': {'gamma': [gamma_min, gamma_max]},
                     'poly': {'degree': [degree_min, degree_max],
                             'coef0': [coef0_min, coef0_max]}}
          }
```

The structure of the search space directly matches the hyperparameterization of every kernel function. We use *None* in the linear kernel as there are no additional hyperparameters. The hyperparameters of the RBF and polynomial kernel follow Opportunity's default syntax based on dictionnaries.

Opportunity also supports nested choices, for example an outer choice for the learning algorithm (e.g. SVM, naive Bayes, ...) and an inner choice for the SVM kernel function. This is illustrated in the following notebook: [sklearn: automated learning method selection and tuning](#).

## 5.3 Examples

Here you can find a variety of examples to illustrate how Opportunity can be used in tandem with other machine learning software. We have split the examples per language and library.

To contribute examples, please send us a pull request on [Github](#).

### 5.3.1 Python

The following examples are available as IPython notebooks in the *OPPORTUNITY/notebooks* folder:

#### Basic: cross-validation

This notebook explores the main elements of Opportunity's cross-validation facilities, including:

- standard cross-validation
- using strata and clusters while constructing folds
- using different aggregators

We recommend perusing the related documentation for more details.

Nested cross-validation is available as a separate notebook.

```
import opportunity
import opportunity.cross_validation
```

We start by generating some toy data containing 6 instances which we will partition into folds.

```
data = list(range(6))
labels = [True] * 3 + [False] * 3
```

#### Standard cross-validation

Each function to be decorated with cross-validation functionality must accept the following arguments: - `x_train`: training data - `x_test`: test data - `y_train`: training labels (required only when `y` is specified in the cross-validation decorator) - `y_test`: test labels (required only when `y` is specified in the cross-validation decorator)

These arguments will be set implicitly by the cross-validation decorator to match the right folds. Any remaining arguments to the decorated function remain as free parameters that must be set later on.

Lets start with the basics and look at Opportunity's cross-validation in action. We use an objective function that simply prints out the train and test data in every split to see what's going on.

```
def f(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t test labels:\t" + str(y_test))
    return 0.0
```

We start with 2 folds, which leads to equally sized train and test partitions.

```
f_2folds = opportunity.cross_validated(x=data, y=labels, num_folds=2) (f)
print("using 2 folds")
f_2folds()
```

```
using 2 folds

train data: [4, 2, 0]      train labels: [False, True, True]
test data: [3, 1, 5]       test labels: [False, True, False]

train data: [3, 1, 5]      train labels: [False, True, False]
test data: [4, 2, 0]       test labels: [False, True, True]
```

```
0.0
```

```
# f_2folds as defined above would typically be written using decorator syntax as follows
# we don't do that in these examples so we can reuse the toy objective function

@opportunity.cross_validated(x=data, y=labels, num_folds=2)
def f_2folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t test labels:\t" + str(y_test))
    return 0.0
```

If we use three folds instead of 2, we get 3 iterations in which the training set is twice the size of the test set.

```
f_3folds = opportunity.cross_validated(x=data, y=labels, num_folds=3) (f)
print("using 3 folds")
f_3folds()
```

```
using 3 folds

train data: [2, 1, 3, 0]      train labels: [True, True, False, True]
test data: [5, 4]   test labels: [False, False]

train data: [5, 4, 3, 0]      train labels: [False, False, False, True]
test data: [2, 1]   test labels: [True, True]

train data: [5, 4, 2, 1]      train labels: [False, False, True, True]
test data: [3, 0]   test labels: [False, True]
```

```
0.0
```

If we do two iterations of 3-fold cross-validation (denoted by 2x3 fold), two sets of folds are generated and evaluated.

```
f_2x3folds = opportunity.cross_validated(x=data, y=labels, num_folds=3, num_iter=2) (f)
print("using 2x3 folds")
f_2x3folds()
```

```
using 2x3 folds

train data: [4, 1, 5, 3]      train labels: [False, True, False, False]
test data: [0, 2]   test labels: [True, True]

train data: [0, 2, 5, 3]      train labels: [True, True, False, False]
test data: [4, 1]   test labels: [False, True]

train data: [0, 2, 4, 1]      train labels: [True, True, False, True]
```

```
test data: [5, 3]    test labels: [False, False]
train data: [0, 2, 1, 4]    train labels: [True, True, True, False]
test data: [5, 3]    test labels: [False, False]
train data: [5, 3, 1, 4]    train labels: [False, False, True, False]
test data: [0, 2]    test labels: [True, True]
train data: [5, 3, 0, 2]    train labels: [False, False, True, True]
test data: [1, 4]    test labels: [True, False]
```

```
0.0
```

### Using strata and clusters

Strata are defined as sets of instances that should be spread out across folds as much as possible (e.g. stratify patients by age). Clusters are sets of instances that must be put in a single fold (e.g. cluster measurements of the same patient).

Opportunity allows you to specify strata and/or clusters that must be accounted for while construct cross-validation folds. Not all instances have to belong to a stratum or clusters.

**Strata** We start by illustrating strata. Strata are specified as a list of lists of instances indices. Each list defines one stratum. We will reuse the toy data and objective function specified above. We will create 2 strata with 2 instances each. These instances will be spread across folds. We create two strata:  $\{0, 1\}$  and  $\{2, 3\}$ .

```
strata = [[0, 1], [2, 3]]
f_stratified = opportunity.cross_validated(x=data, y=labels, strata=strata, num_folds=3) (f)
f_stratified()
```

```
train data: [0, 4, 2, 5]    train labels: [True, False, True, False]
test data: [1, 3]    test labels: [True, False]
train data: [1, 3, 2, 5]    train labels: [True, False, True, False]
test data: [0, 4]    test labels: [True, False]
train data: [1, 3, 0, 4]    train labels: [True, False, True, False]
test data: [2, 5]    test labels: [True, False]
```

```
0.0
```

**Clusters** Clusters work similarly, except that now instances within a cluster are guaranteed to be placed within a single fold. The way to specify clusters is identical to strata. We create two clusters:  $\{0, 1\}$  and  $\{2, 3\}$ . These pairs will always occur in a single fold.

```
clusters = [[0, 1], [2, 3]]
f_clustered = opportunity.cross_validated(x=data, y=labels, clusters=clusters, num_folds=3) (f)
f_clustered()
```

```
train data: [0, 1, 2, 3]    train labels: [True, True, True, False]
test data: [4, 5]    test labels: [False, False]
train data: [4, 5, 2, 3]    train labels: [False, False, True, False]
test data: [0, 1]    test labels: [True, True]
```

```
train data: [4, 5, 0, 1]      train labels:  [False, False, True, True]
test data:  [2, 3]    test labels:    [True, False]
```

```
0.0
```

**Strata and clusters** Strata and clusters can be used together. Lets say we have the following configuration:

- 1 stratum: {0, 1, 2}
- 2 clusters: {0, 3}, {4, 5}

In this particular example, instances 1 and 2 will inevitably end up in a single fold, even though they are part of one stratum. This happens because the total data set has size 6, and 4 instances are already in clusters.

```
strata = [[0, 1, 2]]
clusters = [[0, 3], [4, 5]]
f_strata_clustered = optunity.cross_validated(x=data, y=labels, clusters=clusters, strata=strata, num_folds=3)
f_strata_clustered()
```

```
train data: [4, 5, 0, 3]      train labels:  [False, False, True, False]
test data:  [1, 2]    test labels:    [True, True]

train data: [1, 2, 0, 3]      train labels:  [True, True, True, False]
test data:  [4, 5]    test labels:    [False, False]

train data: [1, 2, 4, 5]      train labels:  [True, True, False, False]
test data:  [0, 3]    test labels:    [True, False]
```

```
0.0
```

## Aggregators

Aggregators are used to combine the scores per fold into a single result. The default approach used in cross-validation is to take the mean of all scores. In some cases, we might be interested in worst-case or best-case performance, the spread, ...

Opunity allows passing a custom callable to be used as aggregator.

The default aggregation in Opunity is to compute the mean across folds.

```
@optunity.cross_validated(x=data, num_folds=3)
def f(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

f(1)
```

```
4
1
2
```

```
2.3333333333333335
```

This can be replaced by any function, e.g. min or max.

```
@opportunity.cross_validated(x=data, num_folds=3, aggregator=max)
def fmax(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

fmax(1)
```

2  
5  
1

5

```
@opportunity.cross_validated(x=data, num_folds=3, aggregator=min)
def fmin(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

fmin(1)
```

3  
4  
5

3

**Retaining intermediate results** Often, it may be useful to retain all intermediate results, not just the final aggregated data. This is made possible via `opportunity.cross_validation.mean_and_list` aggregator. This aggregator computes the mean for internal use in cross-validation, but also returns a list of lists containing the full evaluation results.

```
@opportunity.cross_validated(x=data, num_folds=3,
                             aggregator=opportunity.cross_validation.mean_and_list)
def f_full(x_train, x_test, coeff):
    return x_test[0] * coeff

# evaluate f
mean_score, all_scores = f_full(1.0)
print(mean_score)
print(all_scores)
```

2.333333333333  
[0.0, 2.0, 5.0]

Note that a cross-validation based on the `mean_and_list` aggregator essentially returns a tuple of results. If the result is iterable, all solvers in Opportunity use the first element as the objective function value. You can let the cross-validation procedure return other useful statistics too, which you can access from the solver trace.

```
opt_coeff, info, _ = opportunity.minimize(f_full, coeff=[0, 1], num_evals=10)
print(opt_coeff)
print("call log")
for args, val in zip(info.call_log['args']['coeff'], info.call_log['values']):
    print(str(args) + '\t\t' + str(val))
```

```
{
  'coeff': 0.15771484375
  call log
  0.34521484375          (0.8055013020833334, [0.0, 0.6904296875, 1.72607421875])
  0.47021484375          (1.09716796875, [0.0, 0.9404296875, 2.35107421875])
  0.97021484375          (2.2638346354166665, [0.0, 1.9404296875, 4.85107421875])
  0.72021484375          (1.6805013020833333, [0.0, 1.4404296875, 3.60107421875])
  0.22021484375          (0.5138346354166666, [0.0, 0.4404296875, 1.10107421875])
  0.15771484375          (0.3680013020833333, [0.0, 0.3154296875, 0.78857421875])
  0.65771484375          (1.53466796875, [0.0, 1.3154296875, 3.28857421875])
  0.90771484375          (2.1180013020833335, [0.0, 1.8154296875, 4.53857421875])
  0.40771484375          (0.9513346354166666, [0.0, 0.8154296875, 2.03857421875])
  0.28271484375          (0.65966796875, [0.0, 0.5654296875, 1.41357421875])
}
```

### Cross-validation with scikit-learn

In this example we will show how to use cross-validation methods that are provided by scikit-learn in conjunction with Opportunity. To do this we provide Opportunity with the folds that scikit-learn produces in a specific format.

In supervised learning datasets often have unbalanced labels. When performing cross-validation with unbalanced data it is good practice to preserve the percentage of samples for each class across folds. To achieve this label balance we will use StratifiedKFold.

```

data = list(range(20))
labels = [1 if i%4==0 else 0 for i in range(20)]

@opportunity.cross_validated(x=data, y=labels, num_folds=5)
def unbalanced_folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\ntrain labels:\t" + str(y_train)) + '\n'
    print("test data:\t" + str(x_test) + "\ntest labels:\t" + str(y_test)) + '\n'
    return 0.0

unbalanced_folds()

```

```

train data: [16, 6, 4, 14, 0, 11, 19, 5, 9, 2, 12, 8, 7, 10, 18, 3]
train labels:      [1, 0, 1, 0, 1, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0]

```

```

test data:  [15, 1, 13, 17]
test labels:      [0, 0, 0, 0]

```

```

train data: [15, 1, 13, 17, 0, 11, 19, 5, 9, 2, 12, 8, 7, 10, 18, 3]
train labels:      [0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0, 0]

```

```

test data:  [16, 6, 4, 14]
test labels:      [1, 0, 1, 0]

```

```

train data: [15, 1, 13, 17, 16, 6, 4, 14, 9, 2, 12, 8, 7, 10, 18, 3]
train labels:      [0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 1, 1, 0, 0, 0, 0]

```

```

test data:  [0, 11, 19, 5]
test labels:      [1, 0, 0, 0]

```

```

train data: [15, 1, 13, 17, 16, 6, 4, 14, 0, 11, 19, 5, 7, 10, 18, 3]
train labels:      [0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0, 0, 0]

```

```
test data: [9, 2, 12, 8]
test labels: [0, 0, 1, 1]

train data: [15, 1, 13, 17, 16, 6, 4, 14, 0, 11, 19, 5, 9, 2, 12, 8]
train labels: [0, 0, 0, 0, 1, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 1]

test data: [7, 10, 18, 3]
test labels: [0, 0, 0, 0]
```

```
0.0
```

Notice above how the test label sets have a varying number of positive samples, some have none, some have one, and some have two.

```
from sklearn.cross_validation import StratifiedKFold

stratified_5folds = StratifiedKFold(labels, n_folds=5)
folds = [[list(test) for train, test in stratified_5folds]]

@opportunity.cross_validated(x=data, y=labels, folds=folds, num_folds=5)
def balanced_folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\ntrain labels:\t" + str(y_train)) + '\n'
    print("test data:\t" + str(x_test) + "\ntest labels:\t" + str(y_test)) + '\n'
    return 0.0

balanced_folds()
```

```
train data: [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]
```

```
test data: [0, 1, 2, 3]
test labels: [1, 0, 0, 0]
```

```
train data: [0, 1, 2, 3, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]
```

```
test data: [4, 5, 6, 7]
test labels: [1, 0, 0, 0]
```

```
train data: [0, 1, 2, 3, 4, 5, 6, 7, 12, 13, 14, 15, 16, 17, 18, 19]
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]
```

```
test data: [8, 9, 10, 11]
test labels: [1, 0, 0, 0]
```

```
train data: [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]
```

```
test data: [12, 13, 14, 15]
test labels: [1, 0, 0, 0]
```

```
train data: [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]
```

```
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]
test data: [16, 17, 18, 19]
test labels: [1, 0, 0, 0]
```

```
0.0
```

Now all of our train sets have four positive samples and our test sets have one positive sample.

To use predetermined folds, place a list of the test sample indices into a list. And then insert that list into another list. Why so many nested lists? Because you can perform multiple cross-validation runs by setting num\_iter appropriately and then append num\_iter lists of test samples to the outer most list. Note that the test samples for a given fold are the indices that you provide and then the train samples for that fold are all of the indices from all other test sets joined together. If not done carefully this may lead to duplicated samples in a train set and also samples that fall in both train and test sets of a fold if a datapoint is in multiple folds' test sets.

```
data = list(range(6))
labels = [True] * 3 + [False] * 3

fold1 = [[0, 3], [1, 4], [2, 5]]
fold2 = [[0, 5], [1, 4], [0, 3]] # notice what happens when the indices are not unique
folds = [fold1, fold2]

@opportunity.cross_validated(x=data, y=labels, folds=folds, num_folds=3, num_iter=2)
def multiple_iters(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t\t test labels:\t" + str(y_test))
    return 0.0

multiple_iters()
```

```
train data: [1, 4, 2, 5]      train labels: [True, False, True, False]
test data: [0, 3]              test labels: [True, False]

train data: [0, 3, 2, 5]      train labels: [True, False, True, False]
test data: [1, 4]              test labels: [True, False]

train data: [0, 3, 1, 4]      train labels: [True, False, True, False]
test data: [2, 5]              test labels: [True, False]

train data: [1, 4, 0, 3]      train labels: [True, False, True, False]
test data: [0, 5]              test labels: [True, False]

train data: [0, 5, 0, 3]      train labels: [True, False, True, False]
test data: [1, 4]              test labels: [True, False]

train data: [0, 5, 1, 4]      train labels: [True, False, True, False]
test data: [0, 3]              test labels: [True, False]
```

```
0.0
```

### Basic: nested cross-validation

In this notebook we will briefly illustrate how to use Opportunity for nested cross-validation.

Nested cross-validation is used to reliably estimate generalization performance of a learning pipeline (which may involve preprocessing, tuning, model selection, ...). Before starting this tutorial, we recommend making sure you are

reliable with basic cross-validation in Opportunity.

We will use a scikit-learn SVM to illustrate the key concepts on the MNIST data set.

```
import opportunity
import opportunity.cross_validation
import opportunity.metrics
import numpy as np
import sklearn.svm
```

We load the digits data set and will construct models to distinguish digits 6 from and 8.

```
from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 6
negative_digit = 8

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)
```

The basic nested cross-validation scheme involves two cross-validation routines:

- outer cross-validation: to estimate the generalization performance of the learning pipeline. We will use 5folds.
- inner cross-validation: to use while optimizing hyperparameters. We will use twice iterated 10-fold cross-validation.

Here, we have to take into account that we need to stratify the data based on the label, to ensure we don't run into situations where only one label is available in the train or testing splits. To do this, we use the `strata_by_labels` utility function.

We will use an SVM with RBF kernel and optimize gamma on an exponential grid  $10^{-5} < \gamma < 10^1$  and  $0 < C < 10$  on a linear grid.

```
# outer cross-validation to estimate performance of whole pipeline
@opportunity.cross_validated(x=data, y=labels, num_folds=5,
                             strata=opportunity.cross_validation.strata_by_labels(labels))
def nested_cv(x_train, y_train, x_test, y_test):

    # inner cross-validation to estimate performance of a set of hyperparameters
    @opportunity.cross_validated(x=x_train, y=y_train, num_folds=10, num_iter=2,
                                 strata=opportunity.cross_validation.strata_by_labels(y_train))
    def inner_cv(x_train, y_train, x_test, y_test, C, logGamma):
        # note that the x_train, ... variables in this function are not the same
        # as within nested_cv!
        model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
        predictions = model.decision_function(x_test)
        return opportunity.metrics.roc_auc(y_test, predictions)

    hpars, info, _ = opportunity.maximize(inner_cv, num_evals=100,
                                         C=[0, 10], logGamma=[-5, 1])
    print('')
    print('Hyperparameters: ' + str(hpars))
    print('Cross-validated AUROC after tuning: %1.3f' % info.optimum)
```

```

model = sklearn.svm.SVC(C=hpars['C'], gamma=10 ** hpars['logGamma']).fit(x_train, y_train)
predictions = model.decision_function(x_test)
return opportunity.metrics.roc_auc(y_test, predictions)

auc = nested_cv()
print('')
print('Nested AUROC: %1.3f' % auc)

```

```

Hyperparameters: {'logGamma': -3.8679410473451057, 'C': 0.6162109374999996}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -4.535231399331072, 'C': 0.4839113474508706}
Cross-validated AUROC after tuning: 0.999

```

```

Hyperparameters: {'logGamma': -4.0821875, 'C': 1.5395986549905802}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -3.078125, 'C': 6.015625}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -4.630859375, 'C': 3.173828125}
Cross-validated AUROC after tuning: 1.000

```

```
Nested AUROC: 0.999
```

If you want to explicitly retain statistics from the inner cross-validation procedure, such as the ones we printed below, we can do so by returning tuples in the outer cross-validation and using the `identity` aggregator.

```

# outer cross-validation to estimate performance of whole pipeline
@opportunity.cross_validated(x=data, y=labels, num_folds=5,
                               strata=opportunity.cross_validation.strata_by_labels(labels),
                               aggregator=opportunity.cross_validation.identity)
def nested_cv(x_train, y_train, x_test, y_test):

    # inner cross-validation to estimate performance of a set of hyperparameters
    @opportunity.cross_validated(x=x_train, y=y_train, num_folds=10, num_iter=2,
                                  strata=opportunity.cross_validation.strata_by_labels(y_train))
    def inner_cv(x_train, y_train, x_test, y_test, C, logGamma):
        # note that the x_train, ... variables in this function are not the same
        # as within nested_cv!
        model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
        predictions = model.decision_function(x_test)
        return opportunity.metrics.roc_auc(y_test, predictions)

    hpars, info, _ = opportunity.maximize(inner_cv, num_evals=100,
                                           C=[0, 10], logGamma=[-5, 1])
    model = sklearn.svm.SVC(C=hpars['C'], gamma=10 ** hpars['logGamma']).fit(x_train, y_train)
    predictions = model.decision_function(x_test)

    # return AUROC, optimized hyperparameters and AUROC during hyperparameter search
    return opportunity.metrics.roc_auc(y_test, predictions), hpars, info.optimum

nested_cv_result = nested_cv()

```

We can then process the results like this:

```

aucs, hpars, optima = zip(*nested_cv_result)

print("AUCs: " + str(aucs))

```

```

print('')
print("hpars: " + "\n".join(map(str, hpars)))
print('')
print("optima: " + str(optima))

mean_auc = sum(aucs) / len(aucs)
print('')
print("Mean AUC %1.3f" % mean_auc)

```

```

AUCs: (0.9992063492063492, 1.0, 1.0, 0.9976190476190476, 0.9984126984126984)

hpars: {'logGamma': -3.5753515625, 'C': 3.9048828125000004}
{'logGamma': -2.6765234375, 'C': 6.9193359375000005}
{'logGamma': -3.0538671875, 'C': 2.2935546875}
{'logGamma': -3.593515625, 'C': 4.4136718749999995}
{'logGamma': -3.337747403818736, 'C': 4.367953383541078}

optima: (0.9995032051282051, 0.9985177917320774, 0.9994871794871795, 0.9995238095238095, 0.9995032051282051)

Mean AUC 0.999

```

### Basic: minimizing a simple function

In this example we will use various Opportunity solvers to minimize the following simple 2d parabola:

$$f(x, y) = (x - x_{off})^2 + (y - y_{off})^2,$$

where  $x_{off}$  and  $y_{off}$  are randomly determined offsets.

```

def create_objective_function():
    xoff = random.random()
    yoff = random.random()
    def f(x, y):
        return (x - xoff)**2 + (y - yoff)**2
    return f

```

We start with the necessary imports.

```

# comment this line when running the notebook yourself
%matplotlib inline

import math
import opportunity
import random
import numpy as np
import matplotlib.pyplot as plt
import matplotlib.ticker as ticker

```

First we check which solvers are available. This is available via `opportunity.available_solvers`.

```

solvers = opportunity.available_solvers()
print('Available solvers: ' + ', '.join(solvers))

```

```
Available solvers: particle swarm, tpe, sobol, nelder-mead, random search, cma-es, grid search
```

To run an experiment, we start by generating an objective function.

```
f = create_objective_function()
```

Now we use every available solver to optimize the objective function within the box  $x \in ] -5, 5[$  and  $y \in ] -5, 5[$  and save the results.

```
logs = {}
for solver in solvers:
    pars, details, _ = optunity.minimize(f, num_evals=100, x=[-5, 5], y=[-5, 5], solver_name=solver)
    logs[solver] = np.array([details.call_log['args']['x'],
                           details.call_log['args']['y']])
```

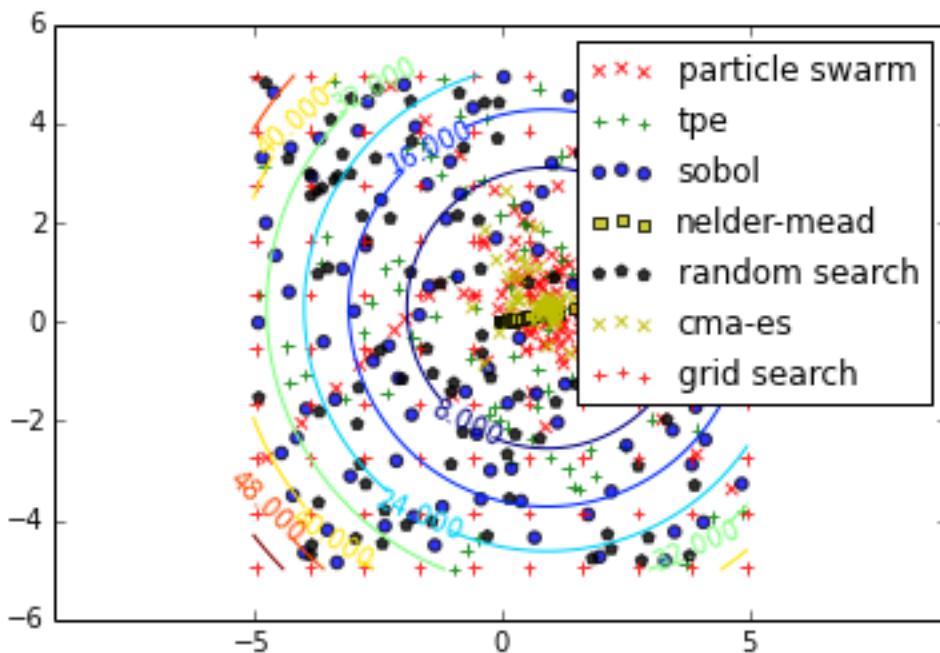
Finally, lets look at the results, that is the trace of each solver along with contours of the objective function.

```
# make sure different traces are somewhat visually separable
colors = ['r', 'g', 'b', 'y', 'k', 'y', 'r', 'g']
markers = ['x', '+', 'o', 's', 'p', 'x', '+', 'o']

# compute contours of the objective function
delta = 0.025
x = np.arange(-5.0, 5.0, delta)
y = np.arange(-5.0, 5.0, delta)
X, Y = np.meshgrid(x, y)
Z = f(X, Y)

CS = plt.contour(X, Y, Z)
plt.clabel(CS, inline=1, fontsize=10, alpha=0.5)
for i, solver in enumerate(solvers):
    plt.scatter(logs[solver][0,:], logs[solver][1,:], c=colors[i], marker=markers[i], alpha=0.80)

plt.xlim([-5, 5])
plt.ylim([-5, 5])
plt.axis('equal')
plt.legend(solvers)
plt.show()
```



Now let's see the performance of the solvers across in 100 repeated experiments. We will do 100 experiments for each solver and then report the resulting statistics. This may take a while to run.

```

optima = dict([(s, []) for s in solvers])
for i in range(100):
    f = create_objective_function()

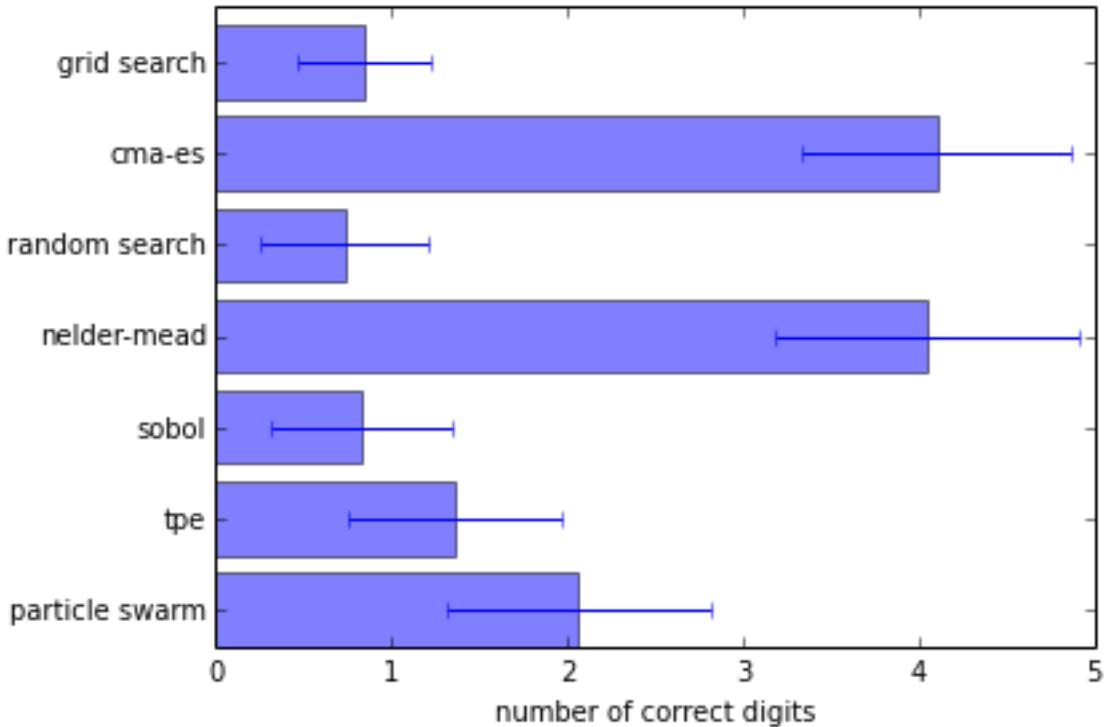
    for solver in solvers:
        pars, details, _ = optunity.minimize(f, num_evals=100, x=[-5, 5], y=[-5, 5],
                                             solver_name=solver)
        # the above line can be parallelized by adding `pmap=optunity.pmap`
        # however this is incompatible with IPython

        optima[solver].append(details.optimum)
        logs[solver] = np.array([details.call_log['args']['x'],
                               details.call_log['args']['y']])

from collections import OrderedDict
log_optima = OrderedDict()
means = OrderedDict()
std = OrderedDict()
for k, v in optima.items():
    log_optima[k] = [-math.log10(val) for val in v]
    means[k] = sum(log_optima[k]) / len(v)
    std[k] = np.std(log_optima[k])

plt.barh(np.arange(len(means)), means.values(), height=0.8, xerr=std.values(), alpha=0.5)
plt.xlabel('number of correct digits')
plt.yticks(np.arange(len(means))+0.4, list(means.keys()))
plt.tight_layout()
plt.show()

```



## Basic: Sobol sequences

In this example we will show the difference between a 2-d Sobol sequence and sampling uniformly at random in 2 dimensions. We will use the `sobol` and `random-search` solvers. The Sobol sequence has lower discrepancy, i.e., the generated samples are spread out better in the sampling space.

This example requires `matplotlib` to generate figures.

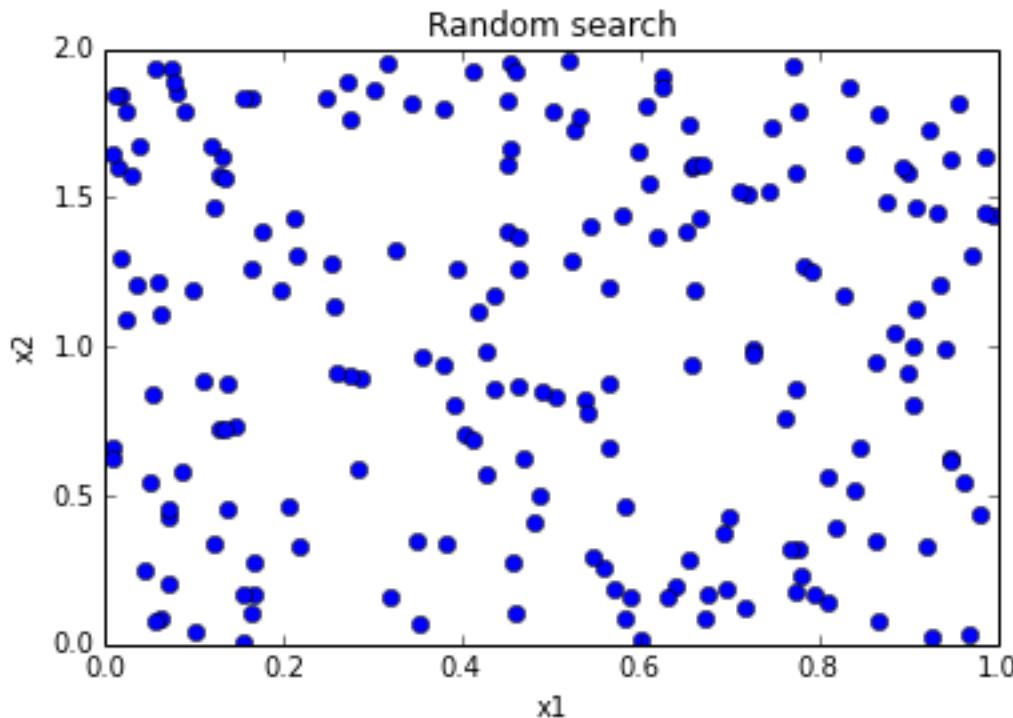
```
#uncomment the following line when you run this notebook interactively
%matplotlib inline
import matplotlib.pyplot as plt
import opportunity
import random
```

Basic configuration and a dummy objective function. We don't care about the returned function value, just the spread of samples.

```
def f(x1, x2):
    return 0.0
```

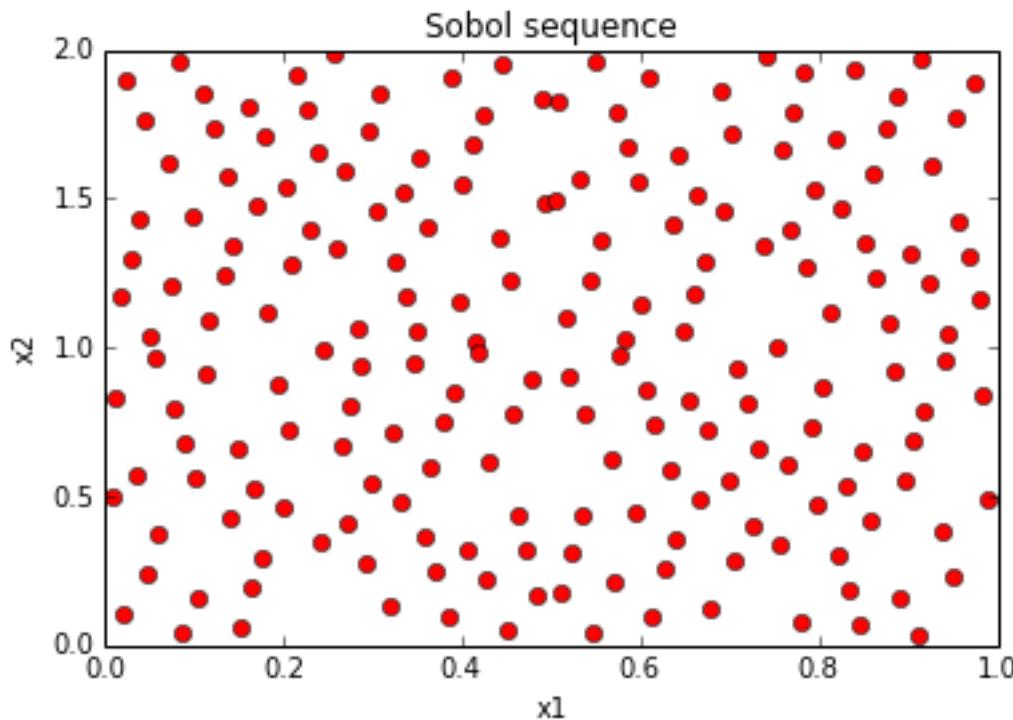
We run the random search solver and Sobol sequence solver.

```
_, info_random, _ = opportunity.minimize(f, num_evals=200, x1=[0, 1], x2=[0, 2], solver_name='random search')
plt.plot(info_random.call_log['args']['x1'], info_random.call_log['args']['x2'], 'bo')
plt.xlabel('x1')
plt.ylabel('x2')
plt.title('Random search')
plt.show()
```



```
_, info_sobol, _ = opportunity.minimize(f, num_evals=200, x1=[0, 1], x2=[0, 2], solver_name='sobol')
plt.plot(info_sobol.call_log['args']['x1'], info_sobol.call_log['args']['x2'], 'ro')
plt.xlabel('x1')
plt.ylabel('x2')
```

```
plt.title('Sobol sequence')
plt.show()
```



### Optimization response surface

In this example we will use Opportunity to optimize hyperparameters for a support vector machine classifier (SVC) in scikit-learn. We will construct a simple toy data set to illustrate that the response surface can be highly irregular for even simple learning tasks (i.e., the response surface is non-smooth, non-convex and has many local minima).

```
import opportunity
import opportunity.metrics

# comment this line if you are running the notebook
import sklearn.svm
import numpy as np
import math
import pandas

%matplotlib inline
from matplotlib import pylab as plt
from mpl_toolkits.mplot3d import Axes3D
```

Create a 2-D toy data set.

```
npos = 200
nneg = 200

delta = 2 * math.pi / npos
radius = 2
circle = np.array(([radius * math.sin(i * delta),
                  radius * math.cos(i * delta))
```

```

        for i in range(npos)) )

neg = np.random.randn(nneg, 2)
pos = np.random.randn(npos, 2) + circle

data = np.vstack((neg, pos))
labels = np.array([False] * nneg + [True] * npos)

```

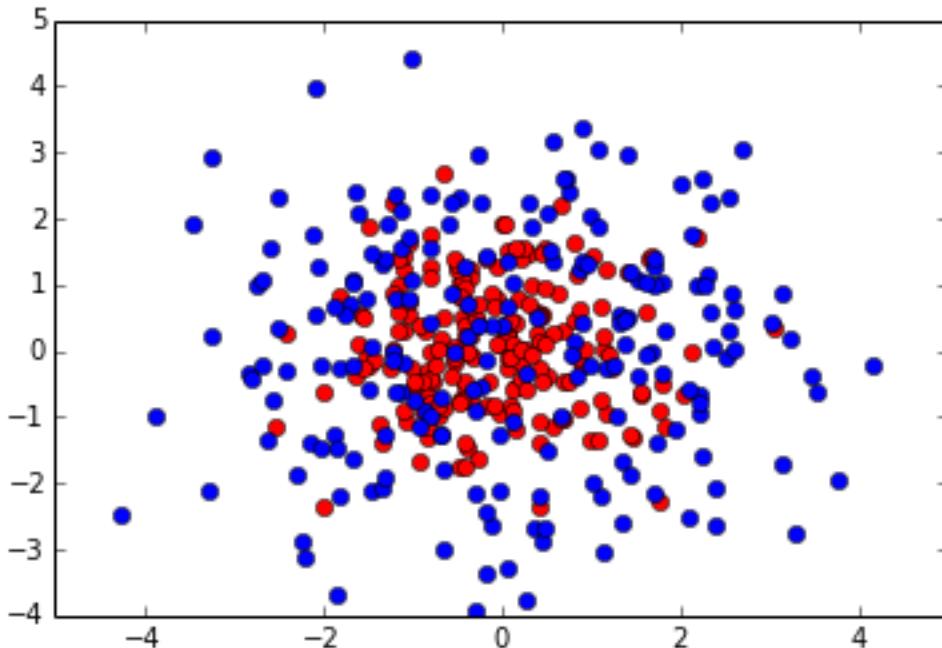
Lets have a look at our 2D data set.

```

plt.plot(neg[:,0], neg[:,1], 'ro')
plt.plot(pos[:,0], pos[:,1], 'bo')

```

```
[<matplotlib.lines.Line2D at 0x7f1f2a6a2b90>]
```



In order to use Optunity to optimize hyperparameters, we start by defining the objective function. We will use 5-fold cross-validated area under the ROC curve. We will regenerate folds for every iteration, which helps to minimize systematic bias originating from the fold partitioning.

We start by defining the objective function `svm_rbf_tuned_auroc()`, which accepts  $C$  and  $\gamma$  as arguments.

```

@optunity.cross_validated(x=data, y=labels, num_folds=5, regenerate_folds=True)
def svm_rbf_tuned_auroc(x_train, y_train, x_test, y_test, logC, logGamma):
    model = sklearn.svm.SVC(C=10 ** logC, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = optunity.metrics.roc_auc(y_test, decision_values)
    return auc

```

Now we can use Optunity to find the hyperparameters that maximize AUROC.

```

optimal_rbf_pars, info, _ = optunity.maximize(svm_rbf_tuned_auroc, num_evals=300, logC=[-4, 2], logGamma=[-4, 2])
# when running this outside of IPython we can parallelize via optunity.pmap
# optimal_rbf_pars, _, _ = optunity.maximize(svm_rbf_tuned_auroc, 150, C=[0, 10], gamma=[0, 0.1], pmap=True)

print("Optimal parameters: " + str(optimal_rbf_pars))

```

```
print("AUROC of tuned SVM with RBF kernel: %1.3f" % info.optimum)
```

```
Optimal parameters: {'logGamma': -1.7262599731822696, 'logC': 0.5460942232689681}
AUROC of tuned SVM with RBF kernel: 0.825
```

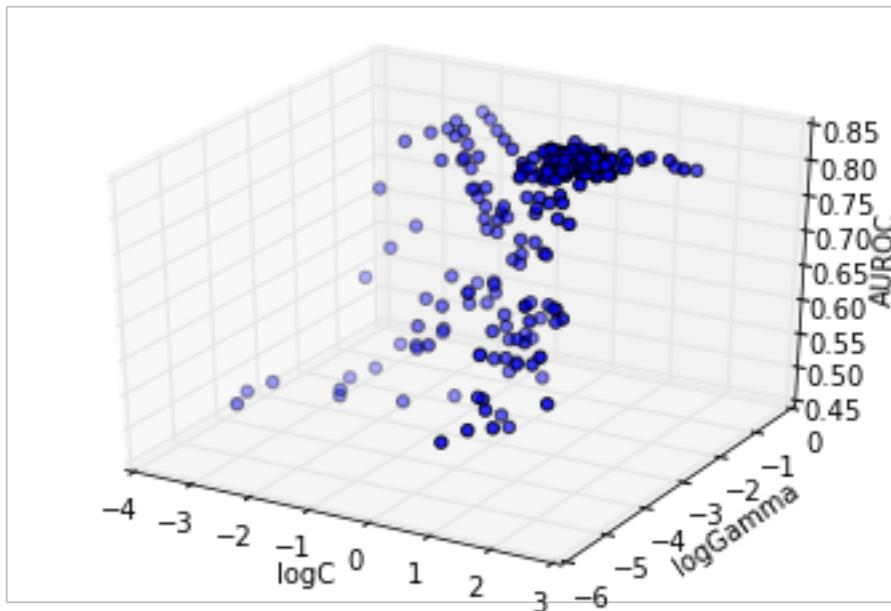
We can turn the call log into a pandas dataframe to efficiently inspect the solver trace.

```
df = optunity.call_log2dataframe(info.call_log)
```

The past function evaluations indicate that the response surface is filled with local minima, caused by finite sample effects. To see this, we can make surface plots.

```
cutoff = 0.5
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
           ys=df[df.value > cutoff]['logGamma'],
           zs=df[df.value > cutoff]['value'])
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')
```

```
<matplotlib.text.Text at 0x7f1f26cbed50>
```



The above plot shows the particles converge directly towards the optimum. At this granularity, the response surface appears smooth.

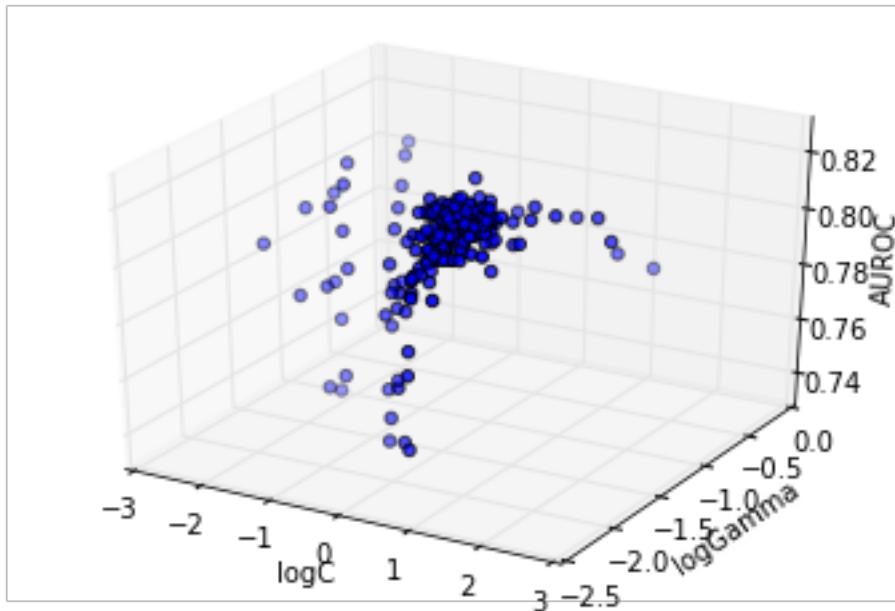
However, a more detailed analysis reveals this is not the case, as shown subsequently:

- showing the sub trace with score up to 90% of the optimum
- showing the sub trace with score up to 95% of the optimum
- showing the sub trace with score up to 99% of the optimum

```
cutoff = 0.9 * info.optimum
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
           ys=df[df.value > cutoff]['logGamma'],
           zs=df[df.value > cutoff]['value'])
```

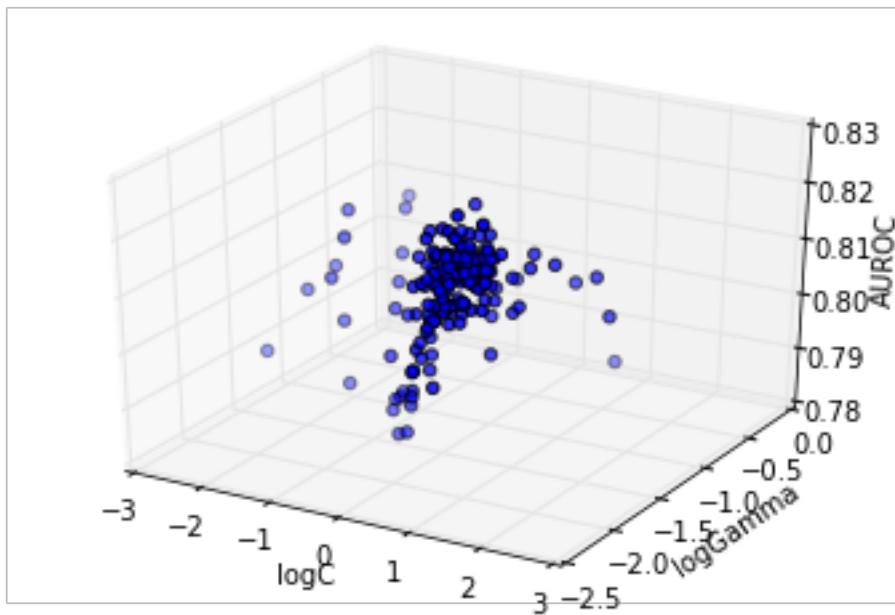
```
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')
```

```
<matplotlib.text.Text at 0x7f1f2a5c3190>
```



```
cutoff = 0.95 * info.optimum
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
           ys=df[df.value > cutoff]['logGamma'],
           zs=df[df.value > cutoff]['value'])
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')
```

```
<matplotlib.text.Text at 0x7f1f26ae5410>
```



Lets further examine the area close to the optimum, that is the 95% region. We will examine a 50x50 grid in this region.

```
minlogc = min(df[df.value > cutoff]['logC'])
maxlogc = max(df[df.value > cutoff]['logC'])
minloggamma = min(df[df.value > cutoff]['logGamma'])
maxloggamma = max(df[df.value > cutoff]['logGamma'])

_, info_new, _ = optunity.maximize(svm_rbf_tuned_auroc, num_evals=2500,
                                    logC=[minlogc, maxlogc],
                                    logGamma=[minloggamma, maxloggamma],
                                    solver_name='grid search')
```

Make a new data frame of the call log for easy manipulation.

```
df_new = optunity.call_log2dataframe(info_new.call_log)
```

Determine the region in which we will do a standard grid search to obtain contours.

```
reshape = lambda x: np.reshape(x, (50, 50))
logcs = reshape(df_new['logC'])
loggammams = reshape(df_new['logGamma'])
values = reshape(df_new['value'])
```

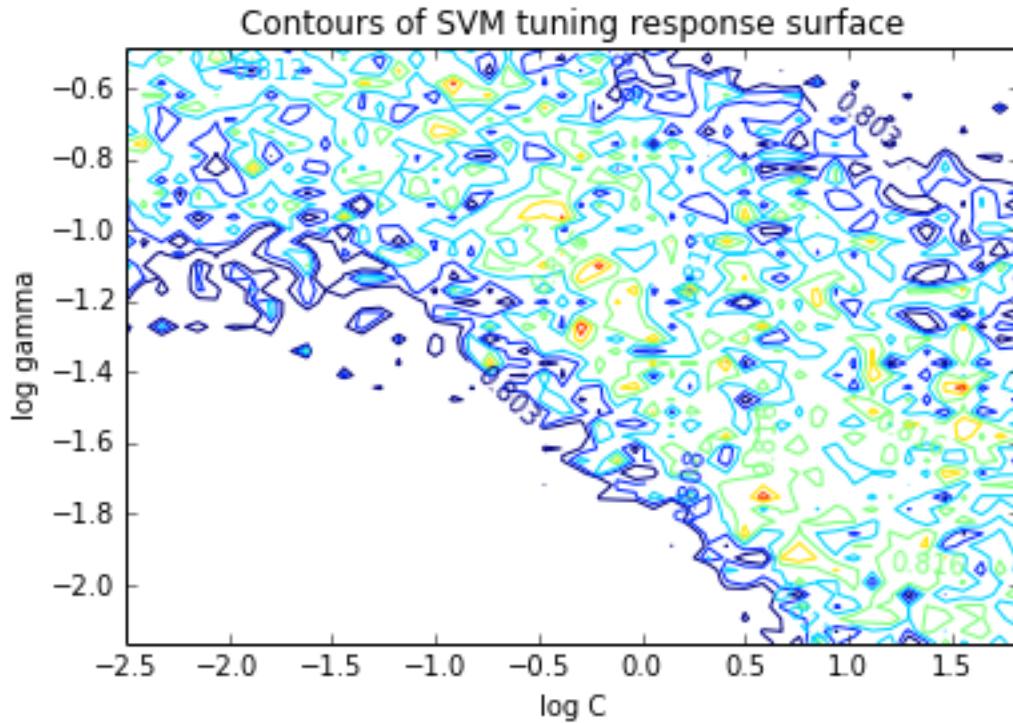
Now, lets look at a contour plot to see whether or not the response surface is smooth.

```
levels = np.arange(0.97, 1.0, 0.005) * info_new.optimum
CS = plt.contour(logcs, loggammams, values, levels=levels)
plt.clabel(CS, inline=1, fontsize=10)
plt.title('Contours of SVM tuning response surface')
plt.xlabel('log C')
plt.ylabel('log gamma')
```

```
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0]] + 1, xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
```

```
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
```

```
<matplotlib.text.Text at 0x7f1f26ac3090>
```



Clearly, this response surface is filled with local minima. This is a general observation in automated hyperparameter optimization, and is one of the key reasons we need robust solvers. If there were no local minima, a simple gradient-descent-like solver would do the trick.

### OpenCV: optical character recognition

We will use OpenCV (<http://www.opencv.org/>) for optical character recognition (OCR) using support vector machine (SVM) classifiers. This example is based on OpenCV's digit tutorial (available in `OPENCV_ROOT/samples/python2/digits.py`).

We start with the necessary imports.

```
import cv2
import numpy as np
import opportunity
import opportunity.metrics

# comment the line below when running the notebook yourself
%matplotlib inline
import matplotlib.pyplot as plt

# finally, we import some convenience functions to avoid bloating this tutorial
# these are available in the `OPTUNITY/notebooks` folder
import opencv_utility_functions as util
```

Load the data, which is available in the `OPTUNITY/notebooks` directory. This data file is a direct copy from OpenCV's example.

```
digits, labels = util.load_digits("opencv-digits.png")
```

```
loading "opencv-digits.png" ...
```

We preprocess the data and make a train/test split.

```
rand = np.random.RandomState(100)
shuffle = rand.permutation(len(digits))
digits, labels = digits[shuffle], labels[shuffle]

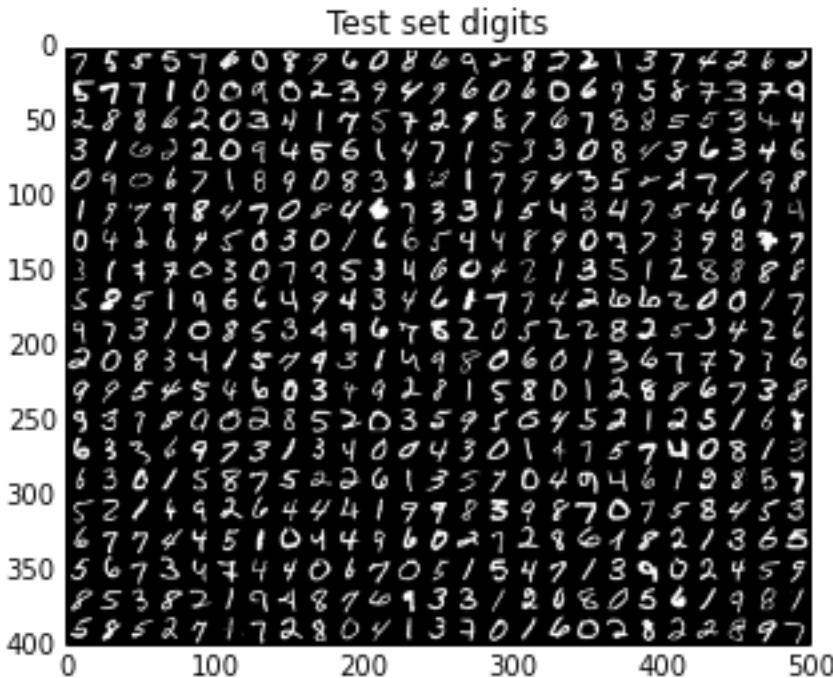
digits2 = map(util.deskew, digits)
samples = util.preprocess_hog(digits2)

train_n = int(0.9*len(samples))

digits_train, digits_test = np.split(digits2, [train_n])
samples_train, samples_test = np.split(samples, [train_n])
labels_train, labels_test = np.split(labels, [train_n])

test_set_img = util.mosaic(25, digits[train_n:])
plt.title('Test set digits')
plt.imshow(test_set_img, cmap='binary_r')
```

```
<matplotlib.image.AxesImage at 0x7fd3ee2e4990>
```



Now, it's time to construct classifiers. We will use a SVM classifier with an RBF kernel, i.e.  $\kappa(\mathbf{u}, \mathbf{v}) = \exp(-\gamma \|\mathbf{u} - \mathbf{v}\|^2)$ . Such an SVM has two hyperparameters that must be optimized, namely the misclassification penalty  $C$  and kernel parameter  $\gamma$ .

We start with an SVM with default parameters, which in this case means:  $C = 1$  and  $\gamma = 0.5$ .

```
model_default = util.SVM()
model_default.train(samples_train, labels_train)
vis_default, err_default = util.evaluate_model(model_default, digits_test, samples_test,
plt.title('unoptimized SVM test set performance')
plt.imshow(vis_default)
```

```

error: 4.20 %
confusion matrix:
[[50  0  0  0  0  0  0  0  0  0]
 [ 0 45  0  0  1  0  0  0  0  0]
 [ 0  0 43  1  0  0  0  0  3  1]
 [ 0  0  0 48  0  1  0  1  0  0]
 [ 0  0  0  0 52  0  2  0  1  1]
 [ 0  0  0  0  0 49  0  0  0  0]
 [ 0  0  0  0  0  0 48  0  0  0]
 [ 0  0  1  3  0  0  0 54  0  0]
 [ 0  1  0  0  0  2  0  0 50  0]
 [ 1  0  0  0  0  0  1  0 40]]]

```

```
<matplotlib.image.AxesImage at 0x7fd3ee288ad0>
```



Next, we will construct a model with optimized hyperparameters. First we need to build Optunity's objective function. We will use 5-fold cross-validated error rate as loss function, which we will minimize.

```

@optunity.cross_validated(x=samples_train, y=labels_train, num_folds=5)
def svm_error_rate(x_train, y_train, x_test, y_test, C, gamma):
    model = util.SVM(C=C, gamma=gamma)
    model.train(x_train, y_train)
    resp = model.predict(x_test)
    error_rate = (y_test != resp).mean()
    return error_rate

```

We will use Optunity's default solver to optimize the error rate given  $0 < C < 5$  and  $0 < \gamma < 10$  and up to 50 function evaluations. This may take a while.

```

optimal_parameters, details, _ = optunity.minimize(svm_error_rate, num_evals=50,
                                                 C=[0, 5], gamma=[0, 10])
# the above line can be parallelized by adding `pmap=optunity.pmap`
# however this is incompatible with IPython

```

```
print("Optimal parameters: C=%1.3f, gamma=%1.3f" % (optimal_parameters['C'], optimal_parameters['gamma']))
print("Cross-validated error rate: %1.3f" % details.optimum)
```

```
Optimal parameters: C=1.798, gamma=6.072
Cross-validated error rate: 0.025
```

Finally, we train a model with the optimized parameters and determine its test set performance.

```
model_opt = util.SVM(**optimal_parameters)
model_opt.train(samples_train, labels_train)
vis_opt, err_opt = util.evaluate_model(model_opt, digits_test, samples_test, labels_test)
plt.title('optimized SVM test set performance')
plt.imshow(vis_opt)
```

```
error: 2.80 %
confusion matrix:
[[50  0  0  0  0  0  0  0  0  0]
 [ 0 45  0  0  1  0  0  0  0  0]
 [ 0  0 44  1  0  0  0  1  1  1]
 [ 0  0  0 49  0  0  0  1  0  0]
 [ 0  0  1  0 53  0  2  0  0  0]
 [ 0  0  0  0  0 49  0  0  0  0]
 [ 0  0  0  0  0  0 48  0  0  0]
 [ 0  0  2  0  0  0  0 56  0  0]
 [ 0  1  0  0  0  1  0  0 51  0]
 [ 0  0  0  0  0  0  1  0 41]]
```

```
<matplotlib.image.AxesImage at 0x7fd3ee4da050>
```



```
print("Reduction in error rate by optimizing hyperparameters: %1.1f%%" % (100.0 - 100.0 * err_opt / 100.0))
```

```
Reduction in error rate by optimizing hyperparameters: 33.3%
```

## sklearn: automated learning method selection and tuning

In this tutorial we will show how to use Opportunity in combination with sklearn to classify the digit recognition data set available in sklearn. The cool part is that we will use Opportunity to choose the best approach from a set of available learning algorithms and optimize hyperparameters in one go. We will use the following learning algorithms:

- k-nearest neighbour
- SVM
- Naive Bayes
- Random Forest

For simplicity, we will focus on a binary classification task, namely digit 3 versus digit 9. We start with the necessary imports and create the data set.

```
import opportunity
import opportunity.metrics
import numpy as np

# k nearest neighbours
from sklearn.neighbors import KNeighborsClassifier
# support vector machine classifier
from sklearn.svm import SVC
# Naive Bayes
from sklearn.naive_bayes import GaussianNB
# Random Forest
from sklearn.ensemble import RandomForestClassifier

from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 3
negative_digit = 9

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)
```

For the SVM model we will let Opportunity optimize the kernel family, choosing from linear, polynomial and RBF. We start by creating a convenience functions for SVM training that handles this:

```
def train_svm(data, labels, kernel, C, gamma, degree, coef0):
    """A generic SVM training function, with arguments based on the chosen kernel."""
    if kernel == 'linear':
        model = SVC(kernel=kernel, C=C)
    elif kernel == 'poly':
        model = SVC(kernel=kernel, C=C, degree=degree, coef0=coef0)
    elif kernel == 'rbf':
        model = SVC(kernel=kernel, C=C, gamma=gamma)
    else:
        raise ValueError("Unknown kernel function: %s" % kernel)
    model.fit(data, labels)
    return model
```

Every learning algorithm has its own hyperparameters:

- k-NN:  $1 < n\_neighbors < 5$  the number of neighbours to use
- SVM: kernel family and misclassification penalty, we will make the penalty contingent on the family. Per kernel family, we have different hyperparameters:
  - polynomial kernel:  $0 < C < 50$ ,  $2 < degree < 5$  and  $0 < coef0 < 1$
  - linear kernel:  $0 < C < 2$ ,
  - RBF kernel:  $0 < C < 10$  and  $0 < gamma < 1$
- naive Bayes: no hyperparameters
- random forest:
  - $10 < n\_estimators < 30$ : number of trees in the forest
  - $5 < max\_features < 20$ : number of features to consider for each split

This translates into the following search space:

```
search = {'algorithm': {'k-nn': {'n_neighbors': [1, 5]},  
                      'SVM': {'kernel': {'linear': {'C': [0, 2]},  
                                 'rbf': {'gamma': [0, 1], 'C': [0, 10]},  
                                 'poly': {'degree': [2, 5], 'C': [0, 50], 'coef0': [0, 1]}  
                               }},  
                     'naive-bayes': None,  
                     'random-forest': {'n_estimators': [10, 30],  
                                      'max_features': [5, 20]}}  
}
```

We also need an objective function that can properly orchestrate everything. We will choose the best model based on area under the ROC curve in 5-fold cross-validation.

```
@opportunity.cross_validated(x=data, y=labels, num_folds=5)  
def performance(x_train, y_train, x_test, y_test,  
                 algorithm, n_neighbors=None, n_estimators=None, max_features=None,  
                 kernel=None, C=None, gamma=None, degree=None, coef0=None):  
    # fit the model  
    if algorithm == 'k-nn':  
        model = KNeighborsClassifier(n_neighbors=int(n_neighbors))  
        model.fit(x_train, y_train)  
    elif algorithm == 'SVM':  
        model = train_svm(x_train, y_train, kernel, C, gamma, degree, coef0)  
    elif algorithm == 'naive-bayes':  
        model = GaussianNB()  
        model.fit(x_train, y_train)  
    elif algorithm == 'random-forest':  
        model = RandomForestClassifier(n_estimators=int(n_estimators),  
                                       max_features=int(max_features))  
        model.fit(x_train, y_train)  
    else:  
        raise ArgumentError('Unknown algorithm: %s' % algorithm)  
  
    # predict the test set  
    if algorithm == 'SVM':  
        predictions = model.decision_function(x_test)  
    else:  
        predictions = model.predict_proba(x_test)[:, 1]
```

```
    return opportunity.metrics.roc_auc(y_test, predictions, positive=True)
```

Lets do a simple test run of this fancy objective function.

```
performance(algorithm='k-nn', n_neighbors=3)
```

```
0.9547920006472639
```

Seems okay! Now we can let Opportunity do its magic with a budget of 300 tries.

```
optimal_configuration, info, _ = opportunity.maximize_structured(performance,
                                                                search_space=search,
                                                                num_evals=300)

print(optimal_configuration)
print(info.optimum)
```

```
{'kernel': 'poly', 'C': 38.5498046875, 'algorithm': 'SVM', 'degree': 3.88525390625, 'n_neighbors': None
0.979302949566}
```

Finally, lets make the results a little bit more readable. All dictionary items in optimal\_configuration with value None can be removed.

```
solution = dict([(k, v) for k, v in optimal_configuration.items() if v is not None])
print('Solution\n=====')
print("\n".join(map(lambda x: "%s %s" % (x[0], str(x[1])), solution.items())))
```

```
Solution
=====
kernel      poly
C          38.5498046875
coef0       0.71826171875
degree      3.88525390625
algorithm   SVM
```

### sklearn: SVM classification

In this example we will use Opportunity to optimize hyperparameters for a support vector machine classifier (SVC) in scikit-learn. We will learn a model to distinguish digits 8 and 9 in the MNIST data set in two settings

- tune SVM with RBF kernel
- tune SVM with RBF, polynomial or linear kernel, that is choose the kernel function and its hyperparameters at once

```
import opportunity
import opportunity.metrics

# comment this line if you are running the notebook
import sklearn.svm
import numpy as np
```

Create the data set: we use the MNIST data set and will build models to distinguish digits 8 and 9.

```
from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 8
```

```

negative_digit = 9

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)

```

First, lets see the performance of an SVC with default hyperparameters.

```

# compute area under ROC curve of default parameters
@opportunity.cross_validated(x=data, y=labels, num_folds=5)
def svm_default_auroc(x_train, y_train, x_test, y_test):
    model = sklearn.svm.SVC().fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = opportunity.metrics.roc_auc(y_test, decision_values)
    return auc

svm_default_auroc()

```

```
0.7328666183635757
```

### Tune SVC with RBF kernel

In order to use Opportunity to optimize hyperparameters, we start by defining the objective function. We will use 5-fold cross-validated area under the ROC curve. For now, lets restrict ourselves to the RBF kernel and optimize  $C$  and  $\gamma$ .

We start by defining the objective function `svm_rbf_tuned_auroc()`, which accepts  $C$  and  $\gamma$  as arguments.

```

# we will make the cross-validation decorator once, so we can reuse it later for the other tuning tasks
# by reusing the decorator, we get the same folds etc.
cv_decorator = opportunity.cross_validated(x=data, y=labels, num_folds=5)

def svm_rbf_tuned_auroc(x_train, y_train, x_test, y_test, C, logGamma):
    model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = opportunity.metrics.roc_auc(y_test, decision_values)
    return auc

svm_rbf_tuned_auroc = cv_decorator(svm_rbf_tuned_auroc)
# this is equivalent to the more common syntax below
# @opportunity.cross_validated(x=data, y=labels, num_folds=5)
# def svm_rbf_tuned_auroc...

svm_rbf_tuned_auroc(C=1.0, logGamma=0.0)

```

```
0.5
```

Now we can use Opportunity to find the hyperparameters that maximize AUROC.

```

optimal_rbf_pars, info, _ = opportunity.maximize(svm_rbf_tuned_auroc, num_evals=150, C=[0, 10], logGamma=[0, 10])
# when running this outside of IPython we can parallelize via opportunity.pmap
# optimal_rbf_pars, _, _ = opportunity.maximize(svm_rbf_tuned_auroc, 150, C=[0, 10], gamma=[0, 0.1], pma

print("Optimal parameters: " + str(optimal_rbf_pars))
print("AUROC of tuned SVM with RBF kernel: %1.3f" % info.optimum)

```

```
Optimal parameters: {'logGamma': -3.0716796875000005, 'C': 3.3025997497032007}
AUROC of tuned SVM with RBF kernel: 0.987
```

We can turn the call log into a pandas dataframe to efficiently inspect the solver trace.

```
import pandas
df = opportunity.call_log2dataframe(info.call_log)
```

Lets look at the best 20 sets of hyperparameters to make sure the results are somewhat stable.

```
df.sort('value', ascending=False) [:10]
```

### Tune SVC without deciding the kernel in advance

In the previous part we choose to use an RBF kernel. Even though the RBF kernel is known to work well for a large variety of problems (and yielded good accuracy here), our choice was somewhat arbitrary.

We will now use Opportunity's conditional hyperparameter optimization feature to optimize over all kernel functions and their associated hyperparameters at once. This requires us to define the search space.

```
space = {'kernel': {'linear': {'C': [0, 2]},
                    'rbf': {'logGamma': [-5, 0], 'C': [0, 10]},
                    'poly': {'degree': [2, 5], 'C': [0, 5], 'coef0': [0, 2]}}
         }
```

We will also have to modify the objective function to cope with conditional hyperparameters. The reason we need to do this explicitly is because scikit-learn doesn't like dealing with None values for irrelevant hyperparameters (e.g. degree when using an RBF kernel). Opportunity will set all irrelevant hyperparameters in a given set to None.

```
def train_model(x_train, y_train, kernel, C, logGamma, degree, coef0):
    """A generic SVM training function, with arguments based on the chosen kernel."""
    if kernel == 'linear':
        model = sklearn.svm.SVC(kernel=kernel, C=C)
    elif kernel == 'poly':
        model = sklearn.svm.SVC(kernel=kernel, C=C, degree=degree, coef0=coef0)
    elif kernel == 'rbf':
        model = sklearn.svm.SVC(kernel=kernel, C=C, gamma=10 ** logGamma)
    else:
        raise ValueError("Unknown kernel function: %s" % kernel)
    model.fit(x_train, y_train)
    return model

def svm_tuned_auroc(x_train, y_train, x_test, y_test, kernel='linear', C=0, logGamma=0, degree=0, coef0=0):
    model = train_model(x_train, y_train, kernel, C, logGamma, degree, coef0)
    decision_values = model.decision_function(x_test)
    return opportunity.metrics.roc_auc(y_test, decision_values)

svm_tuned_auroc = cv_decorator(svm_tuned_auroc)
```

Now we are ready to go and optimize both kernel function and associated hyperparameters!

```
optimal_svm_pars, info, _ = opportunity.maximize_structured(svm_tuned_auroc, space, num_evals=150)
print("Optimal parameters" + str(optimal_svm_pars))
print("AUROC of tuned SVM: %1.3f" % info.optimum)
```

```
Optimal parameters{'kernel': 'rbf', 'C': 3.634209495387873, 'coef0': None, 'degree': None, 'logGamma': -3.0716796875000005}
AUROC of tuned SVM: 0.990
```

Again, we can have a look at the best sets of hyperparameters based on the call log.

```
df = opportunity.call_log2dataframe(info.call_log)
df.sort('value', ascending=False)
```

## sklearn: SVM regression

In this example we will show how to use Opportunity to tune hyperparameters for support vector regression, more specifically:

- measure empirical improvements through nested cross-validation
- optimizing hyperparameters for a given family of kernel functions
- determining the optimal model without choosing the kernel in advance

```
import math
import itertools
import opportunity
import opportunity.metrics
import sklearn.svm
```

We start by creating the data set. We use sklearn's diabetes data.

```
from sklearn.datasets import load_diabetes
diabetes = load_diabetes()
n = diabetes.data.shape[0]

data = diabetes.data
targets = diabetes.target
```

### Nested cross-validation

Nested cross-validation is used to estimate generalization performance of a full learning pipeline, which includes optimizing hyperparameters. We will use three folds in the outer loop.

When using default hyperparameters, there is no need for an inner cross-validation procedure. However, if we want to include tuning in the learning pipeline, the inner loop is used to determine generalization performance with optimized hyperparameters.

We start by measuring generalization performance with default hyperparameters.

```
# we explicitly generate the outer_cv decorator so we can use it twice
outer_cv = opportunity.cross_validated(x=data, y=targets, num_folds=3)

def compute_mse_standard(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and default hyperparameters.
    """
    model = sklearn.svm.SVR().fit(x_train, y_train)
    predictions = model.predict(x_test)
    return opportunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_standard = outer_cv(compute_mse_standard)
```

`compute_mse_standard()` returns a three-fold cross-validation estimate of MSE for an SVR with default hyperparameters.

```
compute_mse_standard()
```

```
6190.481497665955
```

We will create a function that returns MSE based on optimized hyperparameters, where we choose a polynomial kernel in advance.

```
def compute_mse_poly_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @opportunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, C, degree, coef0):
        model = sklearn.svm.SVR(C=C, degree=degree, coef0=coef0, kernel='poly').fit(x_train, y_train)
        predictions = model.predict(x_test)
        return opportunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = opportunity.minimize(tune_cv, 150, C=[1000, 20000], degree=[2, 5],
                                                coef0=[0, 1])

    print("optimal hyperparameters: " + str(optimal_pars))

    tuned_model = sklearn.svm.SVR(kernel='poly', **optimal_pars).fit(x_train, y_train)
    predictions = tuned_model.predict(x_test)
    return opportunity.metrics.mse(y_test, predictions)

    # wrap with outer cross-validation
compute_mse_poly_tuned = outer_cv(compute_mse_poly_tuned)
```

`compute_mse_poly_tuned()` returns a three-fold cross-validation estimate of MSE for an SVR with RBF kernel with tuned hyperparameters  $1000 < C < 20000$ ,  $2 < \text{degree} < 5$  and  $0 < \text{coef0} < 1$  with a budget of 150 function evaluations. Each tuple of hyperparameters is evaluated using twice-iterated 5-fold cross-validation.

```
compute_mse_poly_tuned()
```

```
optimal hyperparameters: {'C': 12078.673881034498, 'coef0': 0.5011052085197018, 'degree': 4.608902814}
optimal hyperparameters: {'C': 14391.165364583334, 'coef0': 0.1731315104166666, 'degree': 2.3582617}
optimal hyperparameters: {'C': 11713.456382191061, 'coef0': 0.49836486667796476, 'degree': 4.6160779}
```

```
3047.035965991627
```

The polynomial kernel yields pretty good results when optimized, but maybe we can do even better with an RBF kernel.

```
def compute_mse_rbf_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @opportunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, C, gamma):
        model = sklearn.svm.SVR(C=C, gamma=gamma).fit(x_train, y_train)
        predictions = model.predict(x_test)
        return opportunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = opportunity.minimize(tune_cv, 150, C=[1, 100], gamma=[0, 50])
    print("optimal hyperparameters: " + str(optimal_pars))

    tuned_model = sklearn.svm.SVR(**optimal_pars).fit(x_train, y_train)
    predictions = tuned_model.predict(x_test)
```

```

    return optunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_rbf_tuned = outer_cv(compute_mse_rbf_tuned)

```

`compute_mse_rbf_tuned()` returns a three-fold cross-validation estimate of MSE for an SVR with RBF kernel with tuned hyperparameters  $1 < C < 100$  and  $0 < \gamma < 5$  with a budget of 150 function evaluations. Each tuple of hyperparameters is evaluated using twice-iterated 5-fold cross-validation.

```
compute_mse_rbf_tuned()
```

```

optimal hyperparameters: {'C': 21.654003906250026, 'gamma': 16.536188056152554}
optimal hyperparameters: {'C': 80.89867187499999, 'gamma': 3.2346692538501784}
optimal hyperparameters: {'C': 19.35431640625002, 'gamma': 22.083848774716085}

```

```
2990.8572696483493
```

Woop! Seems like an RBF kernel is a good choice. An optimized RBF kernel leads to a 50% reduction in MSE compared to the default configuration.

### Determining the kernel family during tuning

In the previous part we've seen that the choice of kernel and its parameters significantly impact performance. However, testing every kernel family separately is cumbersome. It's better to let Opportunity do the work for us.

Opportunity can optimize conditional search spaces, here the kernel family and depending on which family the hyperparameterization ( $\gamma$ , degree,  $\text{coef0}$ , ...). We start by defining the search space (we will try the linear, polynomial and RBF kernel).

```

space = {'kernel': {'linear': {'C': [0, 100]},
                   'rbf': {'gamma': [0, 50], 'C': [1, 100]},
                   'poly': {'degree': [2, 5], 'C': [1000, 20000], 'coef0': [0, 1]}}
        }

```

Now we do nested cross-validation again.

```

def compute_mse_all_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @optunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, kernel, C, gamma, degree, coef0):
        if kernel == 'linear':
            model = sklearn.svm.SVR(kernel=kernel, C=C)
        elif kernel == 'poly':
            model = sklearn.svm.SVR(kernel=kernel, C=C, degree=degree, coef0=coef0)
        elif kernel == 'rbf':
            model = sklearn.svm.SVR(kernel=kernel, C=C, gamma=gamma)
        else:
            raise ArgumentError("Unknown kernel function: %s" % kernel)
        model.fit(x_train, y_train)

        predictions = model.predict(x_test)
        return optunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = optunity.minimize_structured(tune_cv, num_evals=150, search_space=space)

```

```
# remove hyperparameters with None value from optimal_pars
for k, v in optimal_pars.items():
    if v is None: del optimal_pars[k]
print("optimal hyperparameters: " + str(optimal_pars))

tuned_model = sklearn.svm.SVR(**optimal_pars).fit(x_train, y_train)
predictions = tuned_model.predict(x_test)
return optunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_all_tuned = outer_cv(compute_mse_all_tuned)
```

And now the kernel family will be optimized along with its hyperparameterization.

```
compute_mse_all_tuned()
```

```
optimal hyperparameters: {'kernel': 'rbf', 'C': 33.70116043112164, 'gamma': 16.32317353448437}
optimal hyperparameters: {'kernel': 'rbf', 'C': 58.11404170763237, 'gamma': 26.45349823062099}
optimal hyperparameters: {'kernel': 'poly', 'C': 14964.421875843143, 'coef0': 0.5127175861493205, 'd'
```

```
3107.625560844859
```

It looks like the RBF and polynomial kernel are competitive for this problem.

Additional examples:

## OpenCV

### Optimizing a simple 2D parabola

In this example, we will use Optunity to optimize a very simple function, namely a two-dimensional parabola.

More specifically, the objective function is  $f(x, y) = -x^2 - y^2$ .

The full code in Python:

```
import optunity

def f(x, y):
    return -x**2 - y**2

optimal_pars, details, _ = optunity.maximize(f, num_evals=100, x=[-5, 5], y=[-5, 5])
```

For such simple functions we would use different solvers in practice, but the main idea remains.

To get a basic understanding of the way various solvers in Optunity work, we can optimize this function with all solvers and plot the resulting call logs. The source code for this example is available in `bin/examples/python/parabola.py`. Below are plots of the traces and precision of various solvers in optimizing a 2D parabola. These results are averaged over 200 runs, in each run the solvers got a budget of 100 function evaluations with the box  $x = (-5, 5), y = -(5, 5)$ .

Note that both Nelder-Mead simplex and CMA-ES started from an initial solution close to the optimal one. The results on this toy problem do not generalize directly to a real tuning problem.

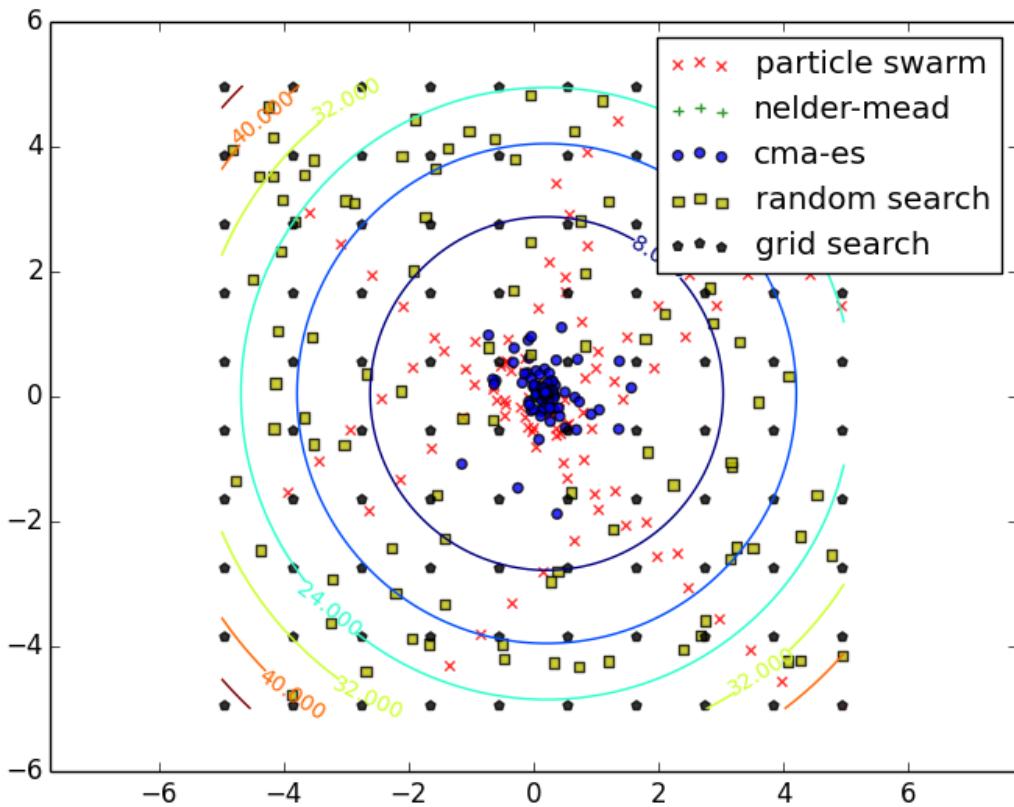


Fig. 5.7: Solver traces showing the evolution of their search for a single run on the 2D parabola.

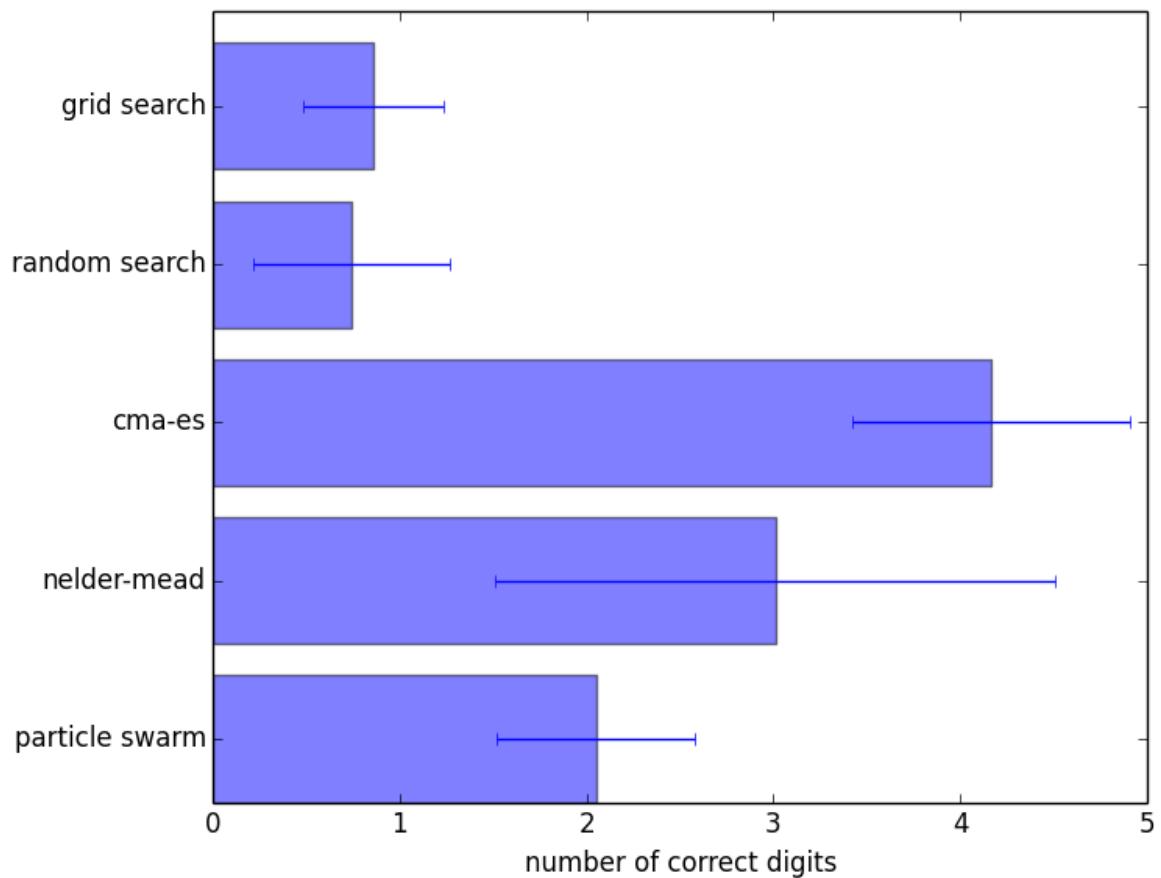


Fig. 5.8: Log10 of the precision of each solver's optimum, averaged across all runs.

## scikit-learn

### Support vector machine classification (SVC)

You can find an executable version of this example in `bin/examples/python/sklearn/svc.py` in your Opportunity release.

In this example, we will train an SVC with RBF kernel using scikit-learn. In this case, we have to tune two hyperparameters:  $C$  and  $\gamma$ . We will use twice iterated 10-fold cross-validation to test a pair of hyperparameters.

In this example, we will use `opportunity.maximize()`.

```
import opportunity
import opportunity.metrics
import sklearn.svm

# score function: twice iterated 10-fold cross-validated accuracy
@opportunity.cross_validated(x=data, y=labels, num_folds=10, num_iter=2)
def svm_auc(x_train, y_train, x_test, y_test, logC, logGamma):
    model = sklearn.svm.SVC(C=10 ** logC, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    return opportunity.metrics.roc_auc(y_test, decision_values)

# perform tuning
hps, _, _ = opportunity.maximize(svm_auc, num_evals=200, logC=[-5, 2], logGamma=[-5, 1])

# train model on the full training set with tuned hyperparameters
optimal_model = sklearn.svm.SVC(C=10 ** hps['logC'], gamma=10 ** hps['logGamma']).fit(data, labels)
```

### Support vector machine regression (SVR)

You can find an executable version of this example in `bin/examples/python/sklearn/svc.py` in your Opportunity release.

In this example, we will train an SVC with RBF kernel using scikit-learn. In this case, we have to tune two hyperparameters:  $C$  and  $\gamma$ . We will use twice iterated 10-fold cross-validation to test a pair of hyperparameters.

In this example, we will use `opportunity.maximize()`.

## Spark

### Logistic regression with Spark and MLlib

In this example, we will train a linear logistic regression model using Spark and MLlib. In this case, we have to tune one hyperparameter: `regParam` for L2 regularization. We will use 5-fold cross-validation to find optimal hyperparameters.

In this example, we will use `opportunity.maximize()`, which by default uses Particle Swarms. Assumes the data is located in file `sample_svm_data.txt`, change it if necessary.

```
# To run Spark with ipython use:
# IPYTHON=1 .bin/pyspark

from pyspark.mllib.classification import LogisticRegressionWithSGD
from pyspark.mllib.regression import LabeledPoint
from numpy import array

import opportunity
```

```
# Load and parse the data
def parsePoint(line):
    values = [float(x) for x in line.split(' ')]
    return LabeledPoint(values[0], values[1:])

data = sc.textFile("sample_svm_data.txt")
parsedData = data.map(parsePoint).cache()

# cross-validation using opportunity
@opportunity.cross_validated(x=parsedData, num_folds=5, num_iter=1)
def logistic_l2_accuracy(x_train, x_test, regParam):
    # cache data to get reasonable speeds for methods like LogisticRegression and SVM
    xc = x_train.cache()
    # training logistic regression with L2 regularization
    model = LogisticRegressionWithSGD.train(xc, regParam=regParam, regType="l2")
    # making prediction on x_test
    yhat = x_test.map(lambda p: (p.label, model.predict(p.features)))
    # returning accuracy on x_test
    return yhat.filter(lambda (v, p): v == p).count() / float(x_test.count())

# using default maximize (particle swarm) with 10 evaluations, regularization between 0 and 10
optimal_pars, _, _ = opportunity.maximize(logistic_l2_accuracy, num_evals=10, regParam=[0, 10])

# training model with all data for the best parameters
model = LogisticRegressionWithSGD.train(parsedData, regType="l2", **optimal_pars)

# prediction (in real application you would use here newData instead of parsedData)
yhat = parsedData.map(lambda p: (p.label, model.predict(p.features)))
```

## Theano

These examples use Opportunity in combination with Theano. Theano is available at <https://github.com/Theano/Theano>.

### Logistic regression

In this example we will use Theano to train logistic regression models on a simple two-dimensional data set. We will use Opportunity to tune the degree of regularization and step sizes (learning rate). This example requires Theano and NumPy.

We start with the necessary imports:

```
import numpy
from numpy.random import multivariate_normal
rng = numpy.random

import theano
import theano.tensor as T

import opportunity
import opportunity.metrics
```

The next step is defining our data set. We will generate a random 2-dimensional data set. The generative procedure for the targets is as follows:  $1 + 2 * x_1 + 3 * x_2 + \text{a noise term}$ . We assign binary class labels based on whether or not the target value is higher than the mean target:

```

N = 200
feats = 2
noise_level = 1
data = multivariate_normal((0.0, 0.0), numpy.array([[1.0, 0.0], [0.0, 1.0]]), N)
noise = noise_level * numpy.random.randn(N)
targets = 1 + 2 * data[:,0] + 3 * data[:,1] + noise

median_target = numpy.median(targets)
labels = numpy.array(map(lambda t: 1 if t > median_target else 0, targets))

```

The next thing we need is a training function for LR models, based on Theano's example:

```

training_steps = 2000

def train_lr(x_train, y_train, regularization=0.01, step=0.1):
    x = T.matrix("x")
    y = T.vector("y")
    w = theano.shared(rng.randn(feats), name="w")
    b = theano.shared(0., name="b")

    # Construct Theano expression graph
    p_1 = 1 / (1 + T.exp(-T.dot(x, w) - b))                      # Probability that target = 1
    prediction = p_1
    xent = -y * T.log(p_1) - (1-y) * T.log(1-p_1)                  # Cross-entropy loss function
    cost = xent.mean() + regularization * (w ** 2).sum()           # The cost to minimize
    gw, gb = T.grad(cost, [w, b])                                    # Compute the gradient of the cost
    # (we shall return to this in a following section of this tutorial)

    # Compile
    train = theano.function(
        inputs=[x,y],
        outputs=[prediction, xent],
        updates=((w, w - step * gw), (b, b - step * gb)))
    predict = theano.function(inputs=[x], outputs=prediction)

    # Train
    for i in range(training_steps):
        train(x_train, y_train)
    return predict, w, b

```

Now that we know how to train, we can define a modeling strategy with default and tuned hyperparameters:

```

def lr_untuned(x_train, y_train, x_test, y_test):
    predict, w, b = train_lr(x_train, y_train)
    yhat = predict(x_test)
    loss = optunity.metrics.logloss(y_test, yhat)
    brier = optunity.metrics.brier(y_test, yhat)
    return loss, brier

def lr_tuned(x_train, y_train, x_test, y_test):
    @optunity.cross_validated(x=x_train, y=y_train, num_folds=3)
    def inner_cv(x_train, y_train, x_test, y_test, regularization, step):
        predict, _, _ = train_lr(x_train, y_train,
                                regularization=regularization, step=step)
        yhat = predict(x_test)
        return optunity.metrics.logloss(y_test, yhat)

    pars, _, _ = optunity.minimize(inner_cv, num_evals=50,

```

```

        regularization=[0.001, 0.05],
        step=[0.01, 0.2])
predict, w, b = train_lr(x_train, y_train, **pars)
yhat = predict(x_test)
loss = opportunity.metrics.logloss(y_test, yhat)
brier = opportunity.metrics.brier(y_test, yhat)
return loss, brier

```

Note that both modeling functions (train, predict, score) return two score measures (log loss and Brier score). We will evaluate both modeling approaches using cross-validation and report both performance measures (see [Cross-validation](#)). The cross-validation decorator:

```

outer_cv = opportunity.cross_validated(x=data, y=labels, num_folds=3,
                                         aggregator=opportunity.cross_validation.list_mean)
lr_untuned = outer_cv(lr_untuned)
lr_tuned = outer_cv(lr_tuned)

```

At this point, *lr\_untuned* and *lr\_tuned* will return a 3-fold cross-validation estimate of [*logloss*, *Brier*] when evaluated.

**Full code** This example is available in detail in `<opportunity>/bin/examples/python/theano/logistic_regression.py`. Typical output of this script will look like:

```

true model: 1 + 2 * x1 + 3 * x2

evaluating untuned LR model
+ model: -0.18 + 1.679 * x1 + 2.045 * x2
++ log loss in test fold: 0.08921125198
++ Brier loss in test fold: 0.0786225946458

+ model: -0.36 + 1.449 * x1 + 2.247 * x2
++ log loss in test fold: 0.08217097905
++ Brier loss in test fold: 0.070741583014

+ model: -0.48 + 1.443 * x1 + 2.187 * x2
++ log loss in test fold: 0.10545356515
++ Brier loss in test fold: 0.0941325050801

evaluating tuned LR model
+ model: -0.66 + 2.354 * x1 + 3.441 * x2
++ log loss in test fold: 0.07508872472
++ Brier loss in test fold: 0.0718020866519

+ model: -0.44 + 2.648 * x1 + 3.817 * x2
++ log loss in test fold: 0.0718891792875
++ Brier loss in test fold: 0.0638209513581

+ model: -0.45 + 2.689 * x1 + 3.858 * x2
++ log loss in test fold: 0.06380803593
++ Brier loss in test fold: 0.0590374290183

Log loss (lower is better):
untuned: 0.0922785987325000
tuned: 0.070261979980

Brier loss (lower is better):
untuned: 0.0811655609133
tuned: 0.0648868223427

```

## 5.3.2 MATLAB

### LS-SVMlab

These examples show use-cases of Opportunity in combination with the LS-SVMlab software, available at <http://www.esat.kuleuven.be/sista/lssvmlab/>.

#### Least-squares SVM classification

You can find a MATLAB script for this example in `<opportunity>/wrappers/matlab/example_lssvmlab/demo_classification.m` in your Opportunity release.

In this example, we will perform nonlinear classification using LS-SVM with RBF kernel using the LS-SVMlab toolbox. In this case, we have to tune two hyperparameters: *gam* (regularization) and *sig2* (kernel bandwidth). We will 10-fold cross-validation to test a pair of hyperparameters.

In this example, we will use `opportunity.minimize()` to minimize misclassification rate.

Start off by defining the objective function, we will generate a double helix for classification:

```
r = 0.5;
h = 1;
helix = @(x, theta, n) [x + r*cos(theta + linspace(0, 4*pi, n)); ...
    r*sin(theta + linspace(0, 4*pi, n)); linspace(0, h, n)]';

noise = 0.1;
ntr = 200;
pos = helix(0, 0, ntr);
neg = helix(0, pi, ntr);
X = [pos; neg] + noise*randn(2*ntr, 3);
Y = [ones(ntr, 1); -1*ones(ntr, 1)];
```

Now that we have training data, we can use LS-SVMlab and Opportunity for tuning. In this example we use LS-SVMlab's (verbose) functional API.

Tuning with LS-SVMlab's built in procedure (in this case a combination of coupled simulated annealing and Nelder-Mead simplex):

```
[lssvm_gam, lssvm_sig2] = tunelssvm({X,Y,'c',[],[],'RBF_kernel'}, 'simplex',...
    'leaveoneoutlssvm', {'misclass'});
[alpha_lssvm,b_lssvm] = trainlssvm({X,Y,type,lssvm_gam,lssvm_sig2,'RBF_kernel'});
```

To do the same with Opportunity, we must first define an objective function, lets say `demo_classification_misclass.m`:

```
function [ misclass ] = demo_classification_misclass( x_train, y_train, x_test, y_test, pars )

% train model
[alpha,b] = trainlssvm({x_train,y_train,'classification',pars.gam,pars.sig2,'RBF_kernel'});

% predict test data
Yt = simlssvm({x_train,y_train,'classification',pars.gam,pars.sig2,'RBF_kernel','preprocess'}, {alpha,b});

% compute misclassification rate
misclass = sum(Yt ~= y_test)/numel(Yt);
end
```

The objective function seems quite involved but it's essentially training a model, predicting test data and evaluating. The code to perform the actual tuning (we use 10-fold cross-validation):

```
obj_fun = opportunity.cross_validate(@demo_classification_misclass, X, 'y', Y, 'num_folds', 10);
opt_pars = opportunity.minimize(obj_fun, 100, 'gam', [1, 100], 'sig2', [0.01, 2]);
[alpha_opportunity,b_opportunity] = trainlssvm({X,Y,type,opt_pars.gam,opt_pars.sig2,'RBF_kernel'});
```

Now we have two LS-SVM models, one tuned with LS-SVM's own algorithms and one using Opportunity. For this particular problem, their performance is close to equivalent. The figures below shows test set predictions for both models.

**LS-SVMlab-tuned test predictions**

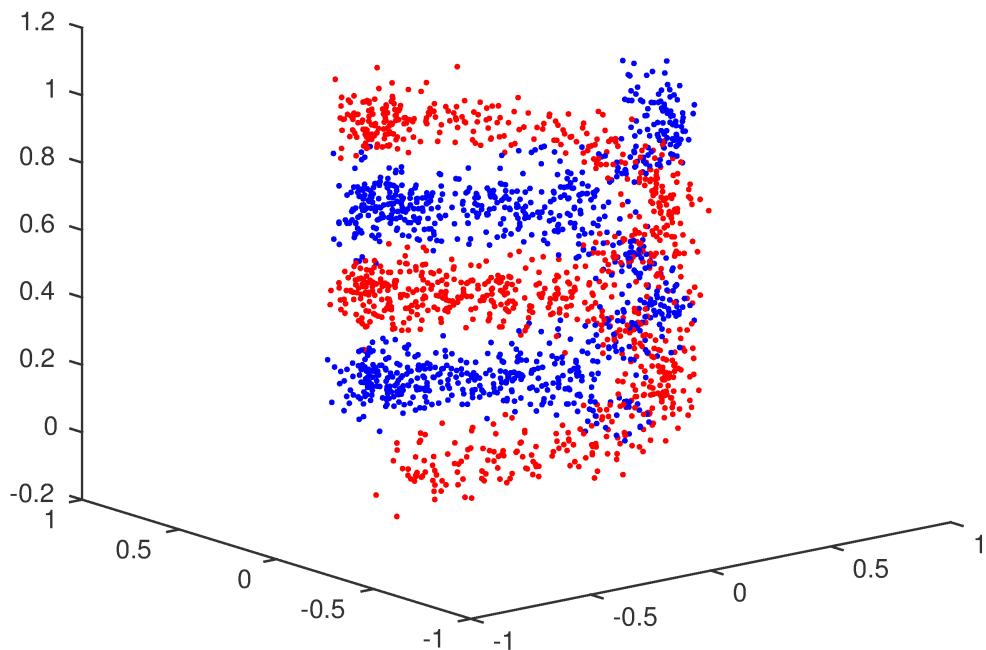


Fig. 5.9: Test set predictions of LS-SVM after tuning with LS-SVMlab.

### Least-squares SVM regression

You can find a MATLAB script for this example in `<opportunity>/wrappers/matlab/example_lssvmlab/demo_regression.m` in your Opportunity release.

In this example, we will perform nonlinear regression using LS-SVM with RBF kernel using the LS-SVMlab toolbox. In this case, we have to tune two hyperparameters: *gam* (regularization) and *sig2* (kernel bandwidth). We will 10-fold cross-validation to test a pair of hyperparameters.

In this example, we will use `opportunity.minimize()` to minimize mean squared error (MSE).

Start off by defining the objective function:

```
fun = @(X) sinc(X) + 0.03 * sin(15*pi * X);
```

This is a simple sinc function with a high frequency sine wave imposed. Now lets generate some noisy data:

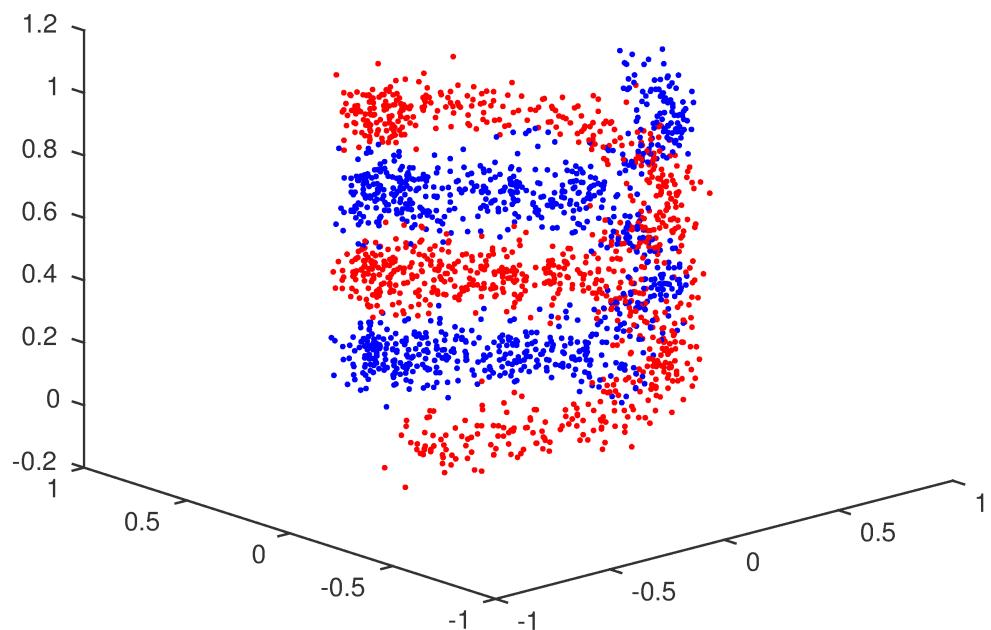
**Opportunity-tuned test predictions**

Fig. 5.10: Test set predictions of LS-SVM after tuning with Opportunity.

```
X = randn(200, 1);
Y = fun(X)+0.1*randn(length(X), 1);
```

Now that we have the function we want to estimate, and data to use for it, we can use LS-SVMlab and Optunity for tuning. In this example we use LS-SVMlab's (verbose) functional API.

Tuning with LS-SVMlab's built in procedure (in this case a combination of coupled simulated annealing and Nelder-Mead simplex):

```
[lssvm_gam, lssvm_sig2] = tunelssvm({X,Y,'f',[],[],'RBF_kernel'}, ...
    'simplex', 'leaveoneoutlssvm', {'mse'});
type = 'function estimation';
[alpha_lssvm,b_lssvm] = trainlssvm({X,Y,type,lssvm_gam,lssvm_sig2,'RBF_kernel'});
```

To do the same with Optunity, we must first define an objective function, lets say *demo\_regression\_mse.m*:

```
function [ mse ] = demo_regression_mse( x_train, y_train, x_test, y_test, pars )
% train model
type = 'function estimation';
[alpha,b] = trainlssvm({x_train,y_train,type,pars.gam,pars.sig2,'RBF_kernel'});

% predict test data
Yt = simlssvm({x_train,y_train,type,pars.gam,pars.sig2, ...
    'RBF_kernel','preprocess'}, {alpha,b}, x_test);

% compute mse
mse = sum((y_test - Yt).^2) / numel(y_test);
end
```

The objective function seems quite involved but it's essentially training a model, predicting test data and evaluating. The code to perform the actual tuning (we use 10-fold cross-validation):

```
obj_fun = optunity.cross_validate(@demo_regression_mse, X, 'y', Y, 'num_folds', 10);
opt_pars = optunity.minimize(obj_fun, 100, 'gam', [1, 30], 'sig2', [0.01, 1]);
[alpha_optunity,b_optunity] = trainlssvm({X, Y, type, ...
    opt_pars.gam, opt_pars.sig2, 'RBF_kernel'});
```

Now we have two LS-SVM models, one tuned with LS-SVM's own algorithms and one using Optunity. For this particular problem, their performance is close to equivalent. The figure below shows test set predictions.

### 5.3.3 R

#### Optimizing a simple 2D parabola

In this example, we will use Optunity in R to maximize a very simple function, namely a two-dimensional parabola.

More specifically, the objective function is  $f(x, y) = -x^2 - y^2$ .

The full code in R:

```
library(optunity)

f <- function(x,y) -x^2 - y^2
opt <- particle_swarm(f, x=c(-5, 5), y=c(-5, 5))
```

In this simple example we used particle swarms optimization with default settings (number of particles 5, number of generations 10).

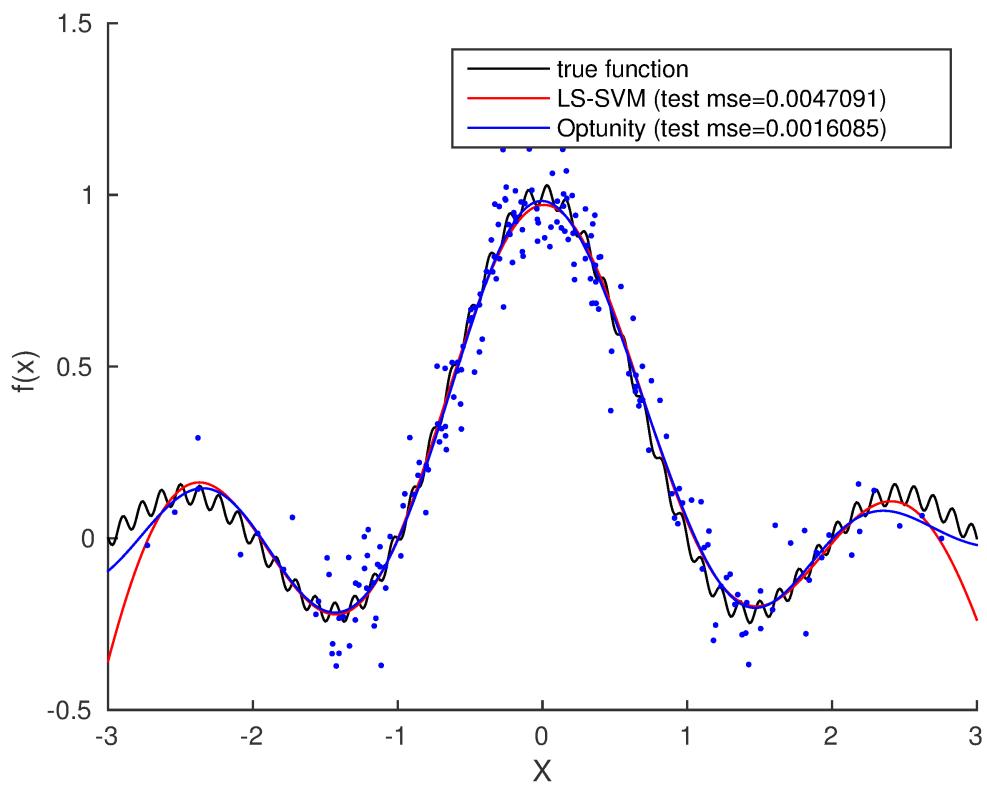


Fig. 5.11: Test set predictions and the true function.

An example with 10 particles and 15 generations:

```
opt <- particle_swarm(f, x=c(-5, 5), y=c(-5, 5), num_particles=10, num_generations=15)
```

In addition to *particle\_swarm* the R interface has *grid\_search*, *random\_search*, *nelder\_mead*. For examples with them use R's internal help, e.g. *?random\_search*.

### Ridge Regression

In this example, we will train a (linear) ridge regression. In this case, we have to tune one hyperparameters: *logC* (regularization). We will use twice iterated 5-fold cross-validation to test a hyperparameter and minimize mean squared error (*mean\_se*):

```
library(optunity)

## artificial data
N <- 50
x <- matrix(runif(N*5), N, 5)
y <- x[,1] + 0.5*x[,2] + 0.1*runif(N)

## ridge regression
regr <- function(x, y, xtest, ytest, logC) {
  ## regularization matrix
  C = diag(x=exp(logC), ncol(x))
  beta = solve(t(x) %*% x + C, t(x) %*% y)
  xtest %*% beta
}
cv <- cv.setup(x, y, score=mean_se, num_folds = 5, num_iter = 2)
res <- cv.particle_swarm(cv, regr, logC = c(-5, 5), maximize = FALSE)

## optimal value for logC:
res$solution$logC
```

Here we used default settings for *cv.particle\_swarm*, see *?cv.particle\_swarm* in R for details.

### SVM (e1071)

In this example, we will train an SVM with RBF kernel and tune its hyperparameters, i.e. **cost** and **gamma** in log-scale. We will use particle swarms to maximize AUC of Precision-Recall (**auc\_pr**) of twice iterated 5-fold cross-validation:

```
library(optunity)
library(e1071)

## artificial data
N <- 100
x <- matrix(runif(N*5), N, 5)
y <- x[,1] + 0.5*x[,2]^2 + 0.1*runif(N) + sin(x[,3]) > 1.0

## SVM train and predict
svm.rbf <- function(x, y, xtest, ytest, logCost, logGamma) {
  m <- svm(x, y, kernel="radial", cost=exp(logCost), gamma=exp(logGamma),
            type="C-classification", probability=TRUE)
  yscore <- attr(predict(m, xtest, probability= T), "probabilities") [, "TRUE"]
  return(yscore)
}

cv <- cv.setup(x, y, score=auc_pr, num_folds = 5, num_iter = 2)
```

```
res <- cv.particle_swarm(cv, svm.rbf, logCost = c(-5, 5), logGamma = c(-5, 5), maximize = TRUE)

## optimal value for logC:
res$solution

## auc_pr reached inside CV
res$optimum
```

Here we used default settings for `cv.particle_swarm`, see `?cv.particle_swarm` in R for details.

## 5.4 Notebooks

Here you can find a variety of notebooks to illustrate various features of Opportunity. All of these examples are available in the *OPPORTUNITY/notebooks* folder.

### 5.4.1 Basic: cross-validation

This notebook explores the main elements of Opportunity's cross-validation facilities, including:

- standard cross-validation
- using strata and clusters while constructing folds
- using different aggregators

We recommend perusing the related documentation for more details.

Nested cross-validation is available as a separate notebook.

```
import opportunity
import opportunity.cross_validation
```

We start by generating some toy data containing 6 instances which we will partition into folds.

```
data = list(range(6))
labels = [True] * 3 + [False] * 3
```

#### Standard cross-validation

Each function to be decorated with cross-validation functionality must accept the following arguments: - `x_train`: training data - `x_test`: test data - `y_train`: training labels (required only when `y` is specified in the cross-validation decorator) - `y_test`: test labels (required only when `y` is specified in the cross-validation decorator)

These arguments will be set implicitly by the cross-validation decorator to match the right folds. Any remaining arguments to the decorated function remain as free parameters that must be set later on.

Lets start with the basics and look at Opportunity's cross-validation in action. We use an objective function that simply prints out the train and test data in every split to see what's going on.

```
def f(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t test labels:\t" + str(y_test))
    return 0.0
```

We start with 2 folds, which leads to equally sized train and test partitions.

```
f_2folds = optunity.cross_validated(x=data, y=labels, num_folds=2)(f)
print("using 2 folds")
f_2folds()
```

```
using 2 folds

train data: [4, 2, 0]      train labels: [False, True, True]
test data: [3, 1, 5]       test labels: [False, True, False]

train data: [3, 1, 5]      train labels: [False, True, False]
test data: [4, 2, 0]       test labels: [False, True, True]
```

```
0.0
```

```
# f_2folds as defined above would typically be written using decorator syntax as follows
# we don't do that in these examples so we can reuse the toy objective function

@optunity.cross_validated(x=data, y=labels, num_folds=2)
def f_2folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t test labels:\t" + str(y_test))
    return 0.0
```

If we use three folds instead of 2, we get 3 iterations in which the training set is twice the size of the test set.

```
f_3folds = optunity.cross_validated(x=data, y=labels, num_folds=3)(f)
print("using 3 folds")
f_3folds()
```

```
using 3 folds

train data: [2, 1, 3, 0]      train labels: [True, True, False, True]
test data: [5, 4]   test labels: [False, False]

train data: [5, 4, 3, 0]      train labels: [False, False, False, True]
test data: [2, 1]   test labels: [True, True]

train data: [5, 4, 2, 1]      train labels: [False, False, True, True]
test data: [3, 0]   test labels: [False, True]
```

```
0.0
```

If we do two iterations of 3-fold cross-validation (denoted by 2x3 fold), two sets of folds are generated and evaluated.

```
f_2x3folds = optunity.cross_validated(x=data, y=labels, num_folds=3, num_iter=2)(f)
print("using 2x3 folds")
f_2x3folds()
```

```
using 2x3 folds

train data: [4, 1, 5, 3]      train labels: [False, True, False, False]
test data: [0, 2]   test labels: [True, True]

train data: [0, 2, 5, 3]      train labels: [True, True, False, False]
test data: [4, 1]   test labels: [False, True]

train data: [0, 2, 4, 1]      train labels: [True, True, False, True]
test data: [5, 3]   test labels: [False, False]
```

```

train data: [0, 2, 1, 4]      train labels:  [True, True, True, False]
test data: [5, 3]    test labels:    [False, False]

train data: [5, 3, 1, 4]      train labels:  [False, False, True, False]
test data: [0, 2]    test labels:    [True, True]

train data: [5, 3, 0, 2]      train labels:  [False, False, True, True]
test data: [1, 4]    test labels:    [True, False]

```

```
0.0
```

## Using strata and clusters

Strata are defined as sets of instances that should be spread out across folds as much as possible (e.g. stratify patients by age). Clusters are sets of instances that must be put in a single fold (e.g. cluster measurements of the same patient).

Opportunity allows you to specify strata and/or clusters that must be accounted for while constructing cross-validation folds. Not all instances have to belong to a stratum or clusters.

### Strata

We start by illustrating strata. Strata are specified as a list of lists of instances indices. Each list defines one stratum. We will reuse the toy data and objective function specified above. We will create 2 strata with 2 instances each. These instances will be spread across folds. We create two strata: {0, 1} and {2, 3}.

```

strata = [[0, 1], [2, 3]]
f_stratified = opportunity.cross_validated(x=data, y=labels, strata=strata, num_folds=3) (f)
f_stratified()

```

```

train data: [0, 4, 2, 5]      train labels:  [True, False, True, False]
test data: [1, 3]    test labels:    [True, False]

train data: [1, 3, 2, 5]      train labels:  [True, False, True, False]
test data: [0, 4]    test labels:    [True, False]

train data: [1, 3, 0, 4]      train labels:  [True, False, True, False]
test data: [2, 5]    test labels:    [True, False]

```

```
0.0
```

### Clusters

Clusters work similarly, except that now instances within a cluster are guaranteed to be placed within a single fold. The way to specify clusters is identical to strata. We create two clusters: {0, 1} and {2, 3}. These pairs will always occur in a single fold.

```

clusters = [[0, 1], [2, 3]]
f_clustered = opportunity.cross_validated(x=data, y=labels, clusters=clusters, num_folds=3) (f)
f_clustered()

```

```

train data: [0, 1, 2, 3]      train labels:  [True, True, True, False]
test data: [4, 5]    test labels:    [False, False]

train data: [4, 5, 2, 3]      train labels:  [False, False, True, False]

```

```
test data: [0, 1]    test labels:  [True, True]
train data: [4, 5, 0, 1]    train labels:  [False, False, True, True]
test data: [2, 3]    test labels:  [True, False]
```

```
0.0
```

### Strata and clusters

Strata and clusters can be used together. Lets say we have the following configuration:

- 1 stratum: {0, 1, 2}
- 2 clusters: {0, 3}, {4, 5}

In this particular example, instances 1 and 2 will inevitably end up in a single fold, even though they are part of one stratum. This happens because the total data set has size 6, and 4 instances are already in clusters.

```
strata = [[0, 1, 2]]
clusters = [[0, 3], [4, 5]]
f_strata_clustered = opunity.cross_validated(x=data, y=labels, clusters=clusters, strata=strata, num_folds=3)
f_strata_clustered()
```

```
train data: [4, 5, 0, 3]    train labels:  [False, False, True, False]
test data: [1, 2]    test labels:  [True, True]

train data: [1, 2, 0, 3]    train labels:  [True, True, True, False]
test data: [4, 5]    test labels:  [False, False]

train data: [1, 2, 4, 5]    train labels:  [True, True, False, False]
test data: [0, 3]    test labels:  [True, False]
```

```
0.0
```

### Aggregators

Aggregators are used to combine the scores per fold into a single result. The default approach used in cross-validation is to take the mean of all scores. In some cases, we might be interested in worst-case or best-case performance, the spread, ...

Opunity allows passing a custom callable to be used as aggregator.

The default aggregation in Opunity is to compute the mean across folds.

```
@opunity.cross_validated(x=data, num_folds=3)
def f(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

f(1)
```

```
4
1
2
```

```
2.3333333333333333
```

This can be replaced by any function, e.g. min or max.

```
@optunity.cross_validated(x=data, num_folds=3, aggregator=max)
def fmax(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

fmax(1)
```

```
2
5
1
```

```
5
```

```
@optunity.cross_validated(x=data, num_folds=3, aggregator=min)
def fmin(x_train, x_test):
    result = x_test[0]
    print(result)
    return result

fmin(1)
```

```
3
4
5
```

```
3
```

## Retaining intermediate results

Often, it may be useful to retain all intermediate results, not just the final aggregated data. This is made possible via `optunity.cross_validation.mean_and_list` aggregator. This aggregator computes the mean for internal use in cross-validation, but also returns a list of lists containing the full evaluation results.

```
@optunity.cross_validated(x=data, num_folds=3,
                           aggregator=optunity.cross_validation.mean_and_list)
def f_full(x_train, x_test, coeff):
    return x_test[0] * coeff

# evaluate f
mean_score, all_scores = f_full(1.0)
print(mean_score)
print(all_scores)
```

```
2.333333333333
[0.0, 2.0, 5.0]
```

Note that a cross-validation based on the `mean_and_list` aggregator essentially returns a tuple of results. If the result is iterable, all solvers in Optunity use the first element as the objective function value. You can let the cross-validation procedure return other useful statistics too, which you can access from the solver trace.

```
opt_coeff, info, _ = optunity.minimize(f_full, coeff=[0, 1], num_evals=10)
print(opt_coeff)
```

```
print("call log")
for args, val in zip(info.call_log['args']['coeff'], info.call_log['values']):
    print(str(args) + '\t\t' + str(val))
```

```
{'coeff': 0.15771484375}
call log
0.34521484375      (0.8055013020833334, [0.0, 0.6904296875, 1.72607421875])
0.47021484375      (1.09716796875, [0.0, 0.9404296875, 2.35107421875])
0.97021484375      (2.2638346354166665, [0.0, 1.9404296875, 4.85107421875])
0.72021484375      (1.6805013020833333, [0.0, 1.4404296875, 3.60107421875])
0.22021484375      (0.5138346354166666, [0.0, 0.4404296875, 1.10107421875])
0.15771484375      (0.3680013020833333, [0.0, 0.3154296875, 0.78857421875])
0.65771484375      (1.53466796875, [0.0, 1.3154296875, 3.28857421875])
0.90771484375      (2.1180013020833335, [0.0, 1.8154296875, 4.53857421875])
0.40771484375      (0.9513346354166666, [0.0, 0.8154296875, 2.03857421875])
0.28271484375      (0.65966796875, [0.0, 0.5654296875, 1.41357421875])
```

### Cross-validation with scikit-learn

In this example we will show how to use cross-validation methods that are provided by scikit-learn in conjunction with Opportunity. To do this we provide Opportunity with the folds that scikit-learn produces in a specific format.

In supervised learning datasets often have unbalanced labels. When performing cross-validation with unbalanced data it is good practice to preserve the percentage of samples for each class across folds. To achieve this label balance we will use StratifiedKFold.

```
data = list(range(20))
labels = [1 if i%4==0 else 0 for i in range(20)]

@opportunity.cross_validated(x=data, y=labels, num_folds=5)
def unbalanced_folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\ntrain labels:\t" + str(y_train) + '\n')
    print("test data:\t" + str(x_test) + "\ntest labels:\t" + str(y_test)) + '\n'
    return 0.0

unbalanced_folds()
```

```
train data: [16, 6, 4, 14, 0, 11, 19, 5, 9, 2, 12, 8, 7, 10, 18, 3]
train labels: [1, 0, 1, 0, 1, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0]
```

```
test data: [15, 1, 13, 17]
test labels: [0, 0, 0, 0]
```

```
train data: [15, 1, 13, 17, 0, 11, 19, 5, 9, 2, 12, 8, 7, 10, 18, 3]
train labels: [0, 0, 0, 0, 1, 0, 0, 0, 0, 1, 1, 0, 0, 0, 0]
```

```
test data: [16, 6, 4, 14]
test labels: [1, 0, 1, 0]
```

```
train data: [15, 1, 13, 17, 16, 6, 4, 14, 9, 2, 12, 8, 7, 10, 18, 3]
train labels: [0, 0, 0, 0, 1, 0, 1, 0, 0, 0, 1, 1, 0, 0, 0]
```

```
test data: [0, 11, 19, 5]
test labels: [1, 0, 0, 0]
```

```

train data: [15, 1, 13, 17, 16, 6, 4, 14, 0, 11, 19, 5, 7, 10, 18, 3]
train labels:      [0, 0, 0, 0, 1, 0, 1, 0, 1, 0, 0, 0, 0, 0, 0, 0]

test data:  [9, 2, 12, 8]
test labels:      [0, 0, 1, 1]

train data: [15, 1, 13, 17, 16, 6, 4, 14, 0, 11, 19, 5, 9, 2, 12, 8]
train labels:      [0, 0, 0, 0, 1, 0, 1, 0, 1, 0, 0, 0, 0, 0, 1, 1]

test data:  [7, 10, 18, 3]
test labels:      [0, 0, 0, 0]

```

```
0.0
```

Notice above how the test label sets have a varying number of positive samples, some have none, some have one, and some have two.

```

from sklearn.cross_validation import StratifiedKFold

stratified_5folds = StratifiedKFold(labels, n_folds=5)
folds = [[list(test) for train, test in stratified_5folds]]

@opportunity.cross_validated(x=data, y=labels, folds=folds, num_folds=5)
def balanced_folds(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\ntrain labels:\t" + str(y_train)) + '\n'
    print("test data:\t" + str(x_test) + "\ntest labels:\t" + str(y_test)) + '\n'
    return 0.0

balanced_folds()

```

```

train data: [4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
train labels:      [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]

test data:  [0, 1, 2, 3]
test labels:      [1, 0, 0, 0]

train data: [0, 1, 2, 3, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19]
train labels:      [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]

test data:  [4, 5, 6, 7]
test labels:      [1, 0, 0, 0]

train data: [0, 1, 2, 3, 4, 5, 6, 7, 12, 13, 14, 15, 16, 17, 18, 19]
train labels:      [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]

test data:  [8, 9, 10, 11]
test labels:      [1, 0, 0, 0]

train data: [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 16, 17, 18, 19]
train labels:      [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]

test data:  [12, 13, 14, 15]

```

```
test labels: [1, 0, 0, 0]

train data: [0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15]
train labels: [1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0, 1, 0, 0, 0]

test data: [16, 17, 18, 19]
test labels: [1, 0, 0, 0]
```

0.0

Now all of our train sets have four positive samples and our test sets have one positive sample.

To use predetermined folds, place a list of the test sample indices into a list. And then insert that list into another list. Why so many nested lists? Because you can perform multiple cross-validation runs by setting num\_iter appropriately and then append num\_iter lists of test samples to the outer most list. Note that the test samples for a given fold are the indices that you provide and then the train samples for that fold are all of the indices from all other test sets joined together. If not done carefully this may lead to duplicated samples in a train set and also samples that fall in both train and test sets of a fold if a datapoint is in multiple folds' test sets.

```
data = list(range(6))
labels = [True] * 3 + [False] * 3

fold1 = [[0, 3], [1, 4], [2, 5]]
fold2 = [[0, 5], [1, 4], [0, 3]] # notice what happens when the indices are not unique
folds = [fold1, fold2]

@opportunity.cross_validated(x=data, y=labels, folds=folds, num_folds=3, num_iter=2)
def multiple_iters(x_train, y_train, x_test, y_test):
    print("")
    print("train data:\t" + str(x_train) + "\t train labels:\t" + str(y_train))
    print("test data:\t" + str(x_test) + "\t\t test labels:\t" + str(y_test))
    return 0.0

multiple_iters()
```

```
train data: [1, 4, 2, 5]      train labels: [True, False, True, False]
test data: [0, 3]              test labels: [True, False]

train data: [0, 3, 2, 5]      train labels: [True, False, True, False]
test data: [1, 4]              test labels: [True, False]

train data: [0, 3, 1, 4]      train labels: [True, False, True, False]
test data: [2, 5]              test labels: [True, False]

train data: [1, 4, 0, 3]      train labels: [True, False, True, False]
test data: [0, 5]              test labels: [True, False]

train data: [0, 5, 0, 3]      train labels: [True, False, True, False]
test data: [1, 4]              test labels: [True, False]

train data: [0, 5, 1, 4]      train labels: [True, False, True, False]
test data: [0, 3]              test labels: [True, False]
```

0.0

## 5.4.2 Basic: nested cross-validation

In this notebook we will briefly illustrate how to use Opportunity for nested cross-validation.

Nested cross-validation is used to reliably estimate generalization performance of a learning pipeline (which may involve preprocessing, tuning, model selection, ...). Before starting this tutorial, we recommend making sure you are reliable with basic cross-validation in Opportunity.

We will use a scikit-learn SVM to illustrate the key concepts on the MNIST data set.

```
import opportunity
import opportunity.cross_validation
import opportunity.metrics
import numpy as np
import sklearn.svm
```

We load the digits data set and will construct models to distinguish digits 6 from and 8.

```
from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 6
negative_digit = 8

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)
```

The basic nested cross-validation scheme involves two cross-validation routines:

- outer cross-validation: to estimate the generalization performance of the learning pipeline. We will use 5folds.
- inner cross-validation: to use while optimizing hyperparameters. We will use twice iterated 10-fold cross-validation.

Here, we have to take into account that we need to stratify the data based on the label, to ensure we don't run into situations where only one label is available in the train or testing splits. To do this, we use the `strata_by_labels` utility function.

We will use an SVM with RBF kernel and optimize gamma on an exponential grid  $10^{-5} < \gamma < 10^1$  and  $0 < C < 10$  on a linear grid.

```
# outer cross-validation to estimate performance of whole pipeline
@opportunity.cross_validated(x=data, y=labels, num_folds=5,
                               strata=opportunity.cross_validation.strata_by_labels(labels))
def nested_cv(x_train, y_train, x_test, y_test):

    # inner cross-validation to estimate performance of a set of hyperparameters
    @opportunity.cross_validated(x=x_train, y=y_train, num_folds=10, num_iter=2,
                                  strata=opportunity.cross_validation.strata_by_labels(y_train))
    def inner_cv(x_train, y_train, x_test, y_test, C, logGamma):
        # note that the x_train, ... variables in this function are not the same
```

```

# as within nested_cv!
model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
predictions = model.decision_function(x_test)
return optunity.metrics.roc_auc(y_test, predictions)

hpars, info, _ = optunity.maximize(inner_cv, num_evals=100,
                                    C=[0, 10], logGamma=[-5, 1])
print('')
print('Hyperparameters: ' + str(hpars))
print('Cross-validated AUROC after tuning: %1.3f' % info.optimum)
model = sklearn.svm.SVC(C=hpars['C'], gamma=10 ** hpars['logGamma']).fit(x_train, y_train)
predictions = model.decision_function(x_test)
return optunity.metrics.roc_auc(y_test, predictions)

auc = nested_cv()
print('')
print('Nested AUROC: %1.3f' % auc)

```

```

Hyperparameters: {'logGamma': -3.8679410473451057, 'C': 0.6162109374999996}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -4.535231399331072, 'C': 0.4839113474508706}
Cross-validated AUROC after tuning: 0.999

```

```

Hyperparameters: {'logGamma': -4.0821875, 'C': 1.5395986549905802}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -3.078125, 'C': 6.015625}
Cross-validated AUROC after tuning: 1.000

```

```

Hyperparameters: {'logGamma': -4.630859375, 'C': 3.173828125}
Cross-validated AUROC after tuning: 1.000

```

```
Nested AUROC: 0.999
```

If you want to explicitly retain statistics from the inner cross-validation procedure, such as the ones we printed below, we can do so by returning tuples in the outer cross-validation and using the `identity` aggregator.

```

# outer cross-validation to estimate performance of whole pipeline
@optunity.cross_validated(x=data, y=labels, num_folds=5,
                          strata=optunity.cross_validation.strata_by_labels(labels),
                          aggregator=optunity.cross_validation.identity)
def nested_cv(x_train, y_train, x_test, y_test):

    # inner cross-validation to estimate performance of a set of hyperparameters
    @optunity.cross_validated(x=x_train, y=y_train, num_folds=10, num_iter=2,
                              strata=optunity.cross_validation.strata_by_labels(y_train))
    def inner_cv(x_train, y_train, x_test, y_test, C, logGamma):
        # note that the x_train, ... variables in this function are not the same
        # as within nested_cv!
        model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
        predictions = model.decision_function(x_test)
        return optunity.metrics.roc_auc(y_test, predictions)

    hpars, info, _ = optunity.maximize(inner_cv, num_evals=100,
                                      C=[0, 10], logGamma=[-5, 1])
    model = sklearn.svm.SVC(C=hpars['C'], gamma=10 ** hpars['logGamma']).fit(x_train, y_train)
    predictions = model.decision_function(x_test)

```

```
# return AUROC, optimized hyperparameters and AUROC during hyperparameter search
return optunity.metrics.roc_auc(y_test, predictions), hpars, info.optimum

nested_cv_result = nested_cv()
```

We can then process the results like this:

```
aucs, hpars, optima = zip(*nested_cv_result)

print("AUCs: " + str(aucs))
print('')
print("hpars: " + "\n".join(map(str, hpars)))
print('')
print("optima: " + str(optima))

mean_auc = sum(aucs) / len(aucs)
print('')
print("Mean AUC %.3f" % mean_auc)
```

```
AUCs: (0.9992063492063492, 1.0, 1.0, 0.9976190476190476, 0.9984126984126984)

hpars: {'logGamma': -3.5753515625, 'C': 3.9048828125000004}
{'logGamma': -2.6765234375, 'C': 6.9193359375000005}
{'logGamma': -3.0538671875, 'C': 2.2935546875}
{'logGamma': -3.593515625, 'C': 4.4136718749999995}
{'logGamma': -3.337747403818736, 'C': 4.367953383541078}

optima: (0.9995032051282051, 0.9985177917320774, 0.9994871794871795, 0.9995238095238095, 0.9995032051282051)

Mean AUC 0.999
```

### 5.4.3 Basic: minimizing a simple function

In this example we will use various Opportunity solvers to minimize the following simple 2d parabola:

$$f(x, y) = (x - x_{off})^2 + (y - y_{off})^2,$$

where  $x_{off}$  and  $y_{off}$  are randomly determined offsets.

```
def create_objective_function():
    xoff = random.random()
    yoff = random.random()
    def f(x, y):
        return (x - xoff)**2 + (y - yoff)**2
    return f
```

We start with the necessary imports.

```
# comment this line when running the notebook yourself
%matplotlib inline

import math
import optunity
import random
import numpy as np
import matplotlib.pyplot as plt
import matplotlib.ticker as ticker
```

First we check which solvers are available. This is available via `opportunity.available_solvers`.

```
solvers = opportunity.available_solvers()
print('Available solvers: ' + ', '.join(solvers))
```

```
Available solvers: particle swarm, tpe, sobol, nelder-mead, random search, cma-es, grid search
```

To run an experiment, we start by generating an objective function.

```
f = create_objective_function()
```

Now we use every available solver to optimize the objective function within the box  $x \in ] -5, 5[$  and  $y \in ] -5, 5[$  and save the results.

```
logs = {}
for solver in solvers:
    pars, details, _ = opportunity.minimize(f, num_evals=100, x=[-5, 5], y=[-5, 5], solver_name=solver)
    logs[solver] = np.array([details.call_log['args']['x'],
                           details.call_log['args']['y']])
```

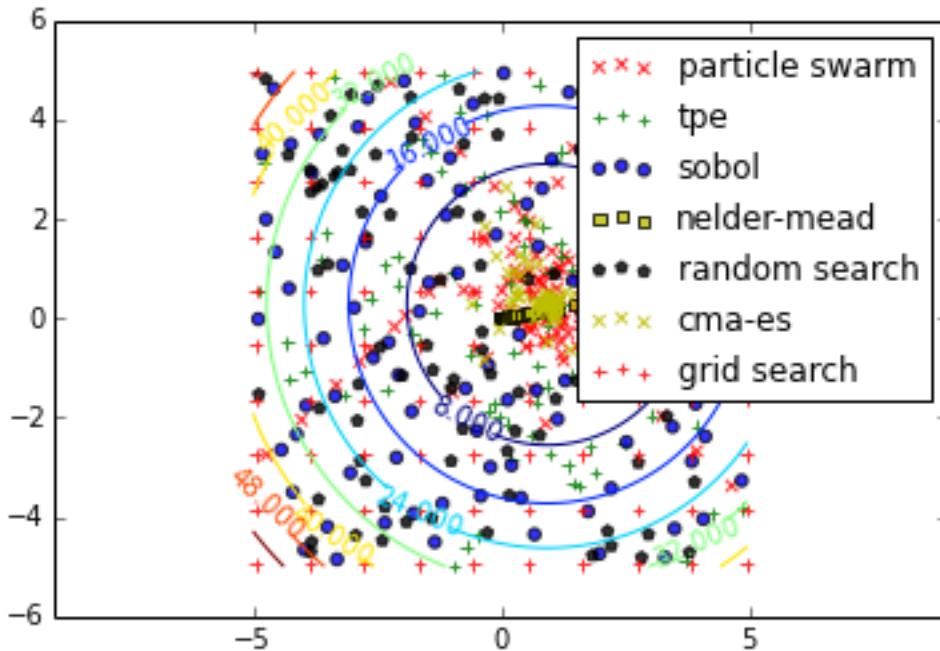
Finally, lets look at the results, that is the trace of each solver along with contours of the objective function.

```
# make sure different traces are somewhat visually separable
colors = ['r', 'g', 'b', 'y', 'k', 'y', 'r', 'g']
markers = ['x', '+', 'o', 's', 'p', 'x', '+', 'o']

# compute contours of the objective function
delta = 0.025
x = np.arange(-5.0, 5.0, delta)
y = np.arange(-5.0, 5.0, delta)
X, Y = np.meshgrid(x, y)
Z = f(X, Y)

CS = plt.contour(X, Y, Z)
plt.clabel(CS, inline=1, fontsize=10, alpha=0.5)
for i, solver in enumerate(solvers):
    plt.scatter(logs[solver][0,:], logs[solver][1,:], c=colors[i], marker=markers[i], alpha=0.80)

plt.xlim([-5, 5])
plt.ylim([-5, 5])
plt.axis('equal')
plt.legend(solvers)
plt.show()
```



Now lets see the performance of the solvers across in 100 repeated experiments. We will do 100 experiments for each solver and then report the resulting statistics. This may take a while to run.

```

optima = dict([(s, []) for s in solvers])
for i in range(100):
    f = create_objective_function()

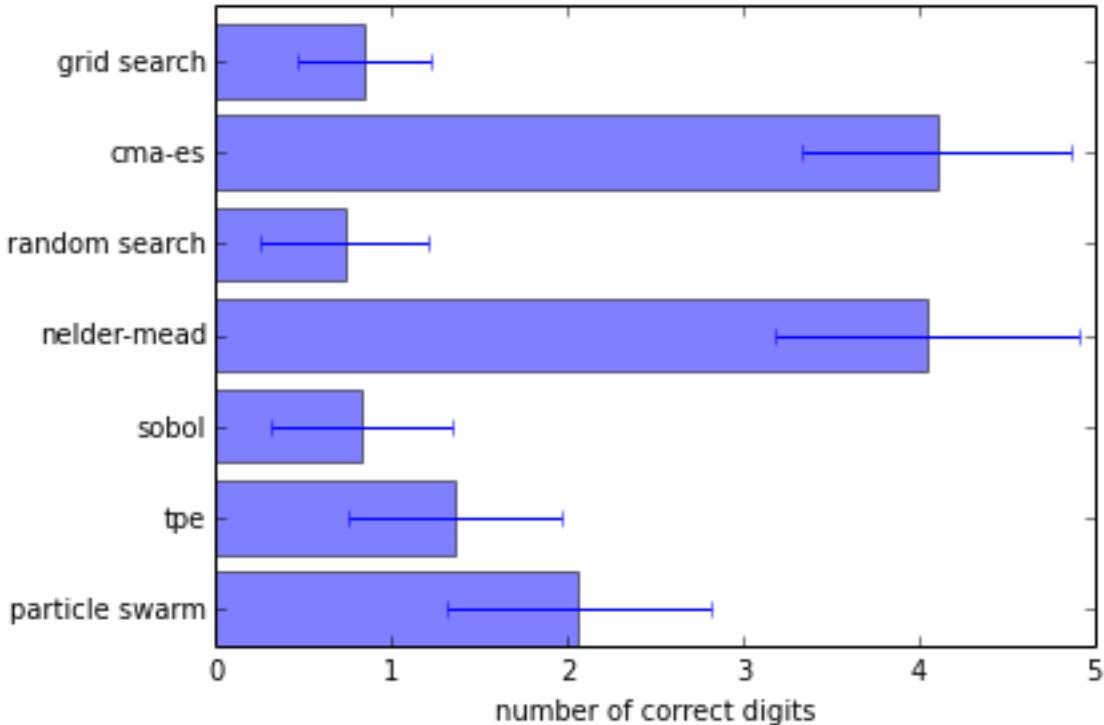
    for solver in solvers:
        pars, details, _ = optunity.minimize(f, num_evals=100, x=[-5, 5], y=[-5, 5],
                                              solver_name=solver)
        # the above line can be parallelized by adding `pmap=optunity.pmap`
        # however this is incompatible with IPython

        optima[solver].append(details.optimum)
        logs[solver] = np.array([details.call_log['args']['x'],
                               details.call_log['args']['y']])

from collections import OrderedDict
log_optima = OrderedDict()
means = OrderedDict()
std = OrderedDict()
for k, v in optima.items():
    log_optima[k] = [-math.log10(val) for val in v]
    means[k] = sum(log_optima[k]) / len(v)
    std[k] = np.std(log_optima[k])

plt.barh(np.arange(len(means)), means.values(), height=0.8, xerr=std.values(), alpha=0.5)
plt.xlabel('number of correct digits')
plt.yticks(np.arange(len(means))+0.4, list(means.keys()))
plt.tight_layout()
plt.show()

```



#### 5.4.4 Basic: Sobol sequences

In this example we will show the difference between a 2-d Sobol sequence and sampling uniformly at random in 2 dimensions. We will use the `sobol` and `random-search` solvers. The Sobol sequence has lower discrepancy, i.e., the generated samples are spread out better in the sampling space.

This example requires matplotlib to generate figures.

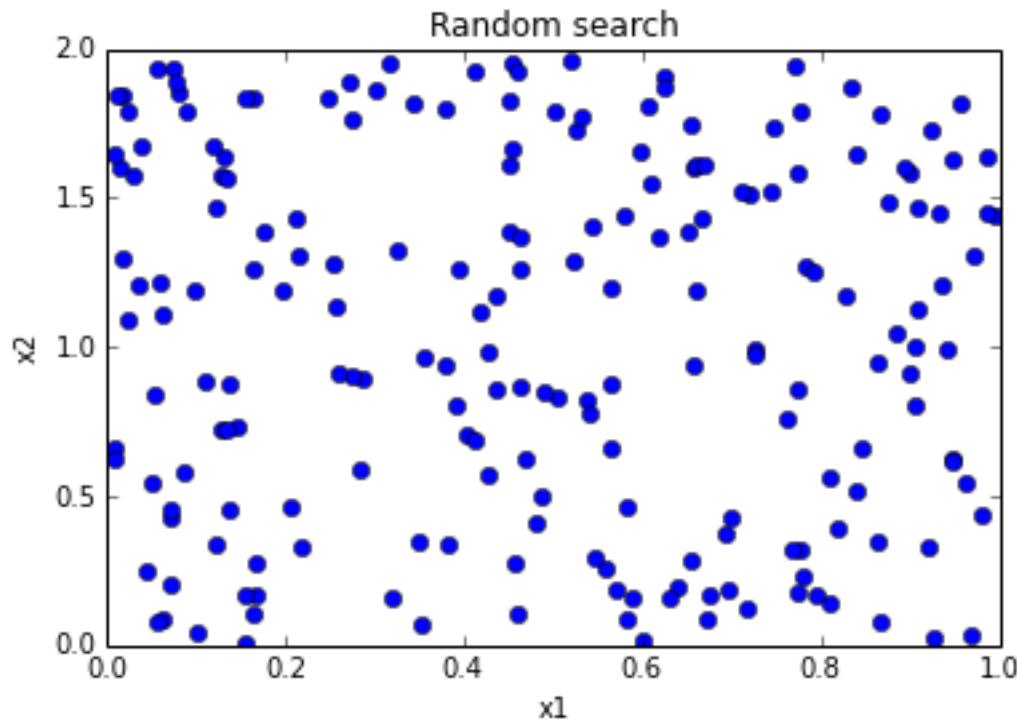
```
#uncomment the following line when you run this notebook interactively
%matplotlib inline
import matplotlib.pyplot as plt
import optunity
import random
```

Basic configuration and a dummy objective function. We don't care about the returned function value, just the spread of samples.

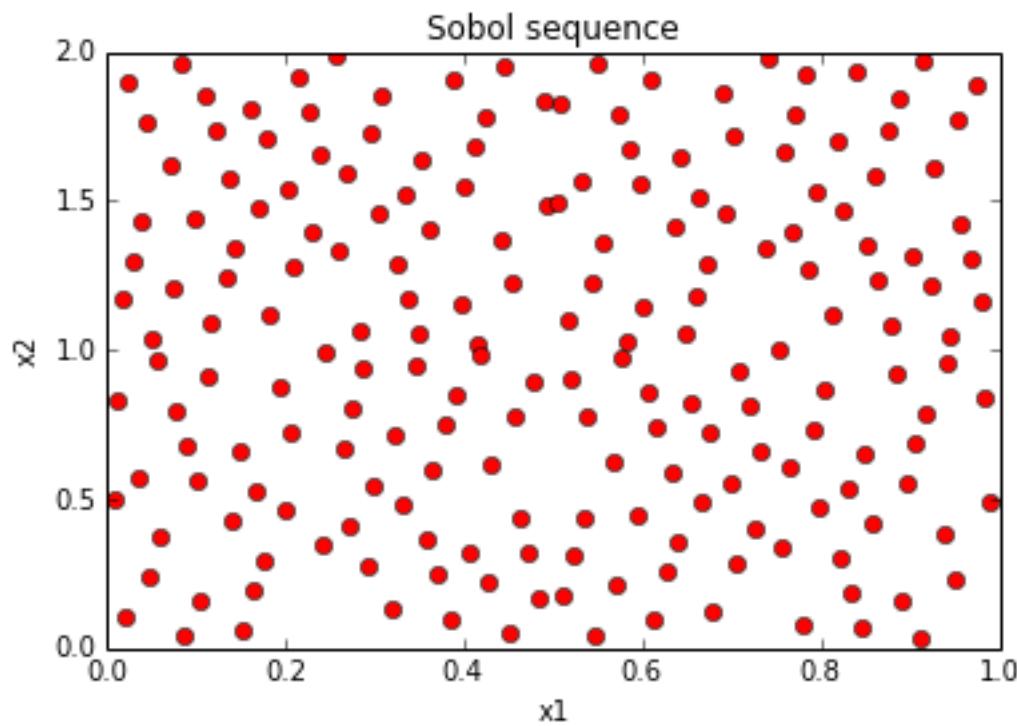
```
def f(x1, x2):
    return 0.0
```

We run the random search solver and Sobol sequence solver.

```
_, info_random, _ = optunity.minimize(f, num_evals=200, x1=[0, 1], x2=[0, 2], solver_name='random search')
plt.plot(info_random.call_log['args']['x1'], info_random.call_log['args']['x2'], 'bo')
plt.xlabel('x1')
plt.ylabel('x2')
plt.title('Random search')
plt.show()
```



```
_, info_sobol, _ = optunity.minimize(f, num_evals=200, x1=[0, 1], x2=[0, 2], solver_name='sobol')
plt.plot(info_sobol.call_log['args']['x1'], info_sobol.call_log['args']['x2'], 'ro')
plt.xlabel('x1')
plt.ylabel('x2')
plt.title('Sobol sequence')
plt.show()
```



## 5.4.5 Optimization response surface

In this example we will use Optunity to optimize hyperparameters for a support vector machine classifier (SVC) in scikit-learn. We will construct a simple toy data set to illustrate that the response surface can be highly irregular for even simple learning tasks (i.e., the response surface is non-smooth, non-convex and has many local minima).

```
import optunity
import optunity.metrics

# comment this line if you are running the notebook
import sklearn.svm
import numpy as np
import math
import pandas

%matplotlib inline
from matplotlib import pylab as plt
from mpl_toolkits.mplot3d import Axes3D
```

Create a 2-D toy data set.

```
npos = 200
nneg = 200

delta = 2 * math.pi / npos
radius = 2
circle = np.array(([radius * math.sin(i * delta),
                   radius * math.cos(i * delta))
                  for i in range(npos)]))

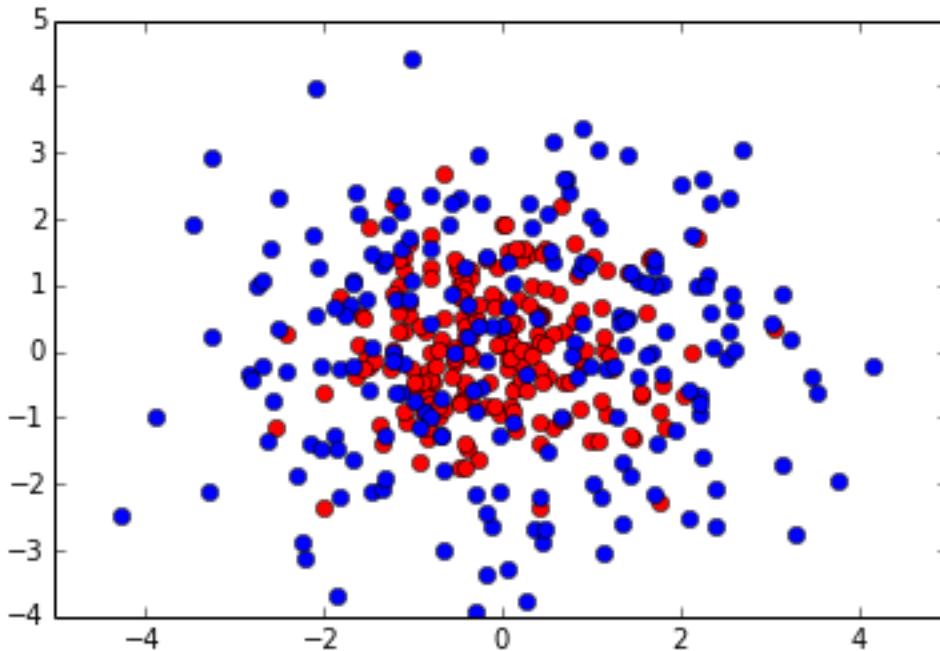
neg = np.random.randn(nneg, 2)
pos = np.random.randn(npos, 2) + circle

data = np.vstack((neg, pos))
labels = np.array([False] * nneg + [True] * npos)
```

Lets have a look at our 2D data set.

```
plt.plot(neg[:, 0], neg[:, 1], 'ro')
plt.plot(pos[:, 0], pos[:, 1], 'bo')
```

```
[<matplotlib.lines.Line2D at 0x7f1f2a6a2b90>]
```



In order to use Opportunity to optimize hyperparameters, we start by defining the objective function. We will use 5-fold cross-validated area under the ROC curve. We will regenerate folds for every iteration, which helps to minimize systematic bias originating from the fold partitioning.

We start by defining the objective function `svm_rbf_tuned_auroc()`, which accepts  $C$  and  $\gamma$  as arguments.

```
@opportunity.cross_validated(x=data, y=labels, num_folds=5, regenerate_folds=True)
def svm_rbf_tuned_auroc(x_train, y_train, x_test, y_test, logC, logGamma):
    model = sklearn.svm.SVC(C=10 ** logC, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = opportunity.metrics.roc_auc(y_test, decision_values)
    return auc
```

Now we can use Opportunity to find the hyperparameters that maximize AUROC.

```
optimal_rbf_pars, info, _ = opportunity.maximize(svm_rbf_tuned_auroc, num_evals=300, logC=[-4, 2], logGamma=[0, 0.1])
# when running this outside of IPython we can parallelize via opportunity.pmap
# optimal_rbf_pars, _, _ = opportunity.maximize(svm_rbf_tuned_auroc, 150, C=[0, 10], gamma=[0, 0.1], pmap=True)

print("Optimal parameters: " + str(optimal_rbf_pars))
print("AUROC of tuned SVM with RBF kernel: %1.3f" % info.optimum)
```

```
Optimal parameters: {'logGamma': -1.7262599731822696, 'logC': 0.5460942232689681}
AUROC of tuned SVM with RBF kernel: 0.825
```

We can turn the call log into a pandas dataframe to efficiently inspect the solver trace.

```
df = opportunity.call_log2dataframe(info.call_log)
```

The past function evaluations indicate that the response surface is filled with local minima, caused by finite sample effects. To see this, we can make surface plots.

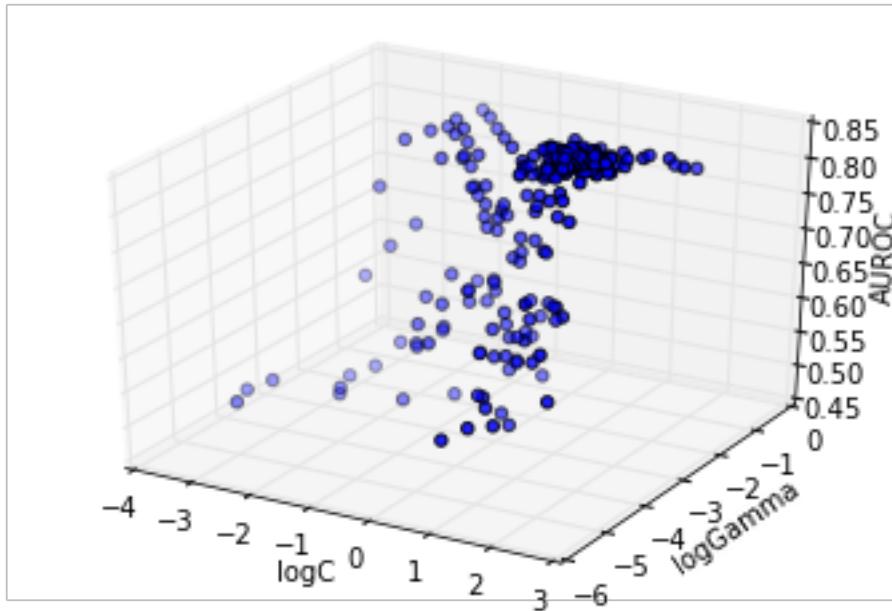
```
cutoff = 0.5
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
```

```

ys=df[df.value > cutoff]['logGamma'],
zs=df[df.value > cutoff]['value'])
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')

```

```
<matplotlib.text.Text at 0x7f1f26cbed50>
```



The above plot shows the particles converge directly towards the optimum. At this granularity, the response surface appears smooth.

However, a more detailed analysis reveals this is not the case, as shown subsequently:

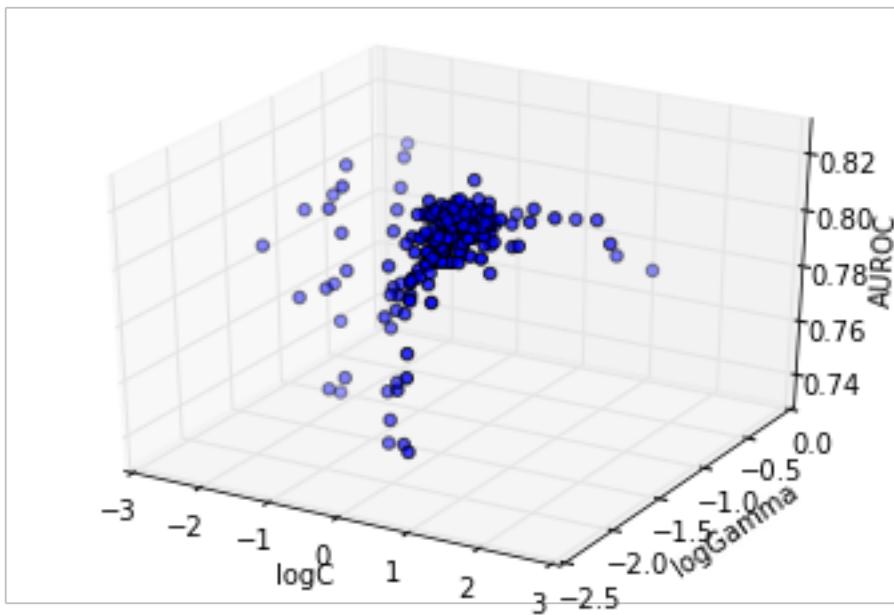
- showing the sub trace with score up to 90% of the optimum
- showing the sub trace with score up to 95% of the optimum
- showing the sub trace with score up to 99% of the optimum

```

cutoff = 0.9 * info.optimum
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
           ys=df[df.value > cutoff]['logGamma'],
           zs=df[df.value > cutoff]['value'])
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')

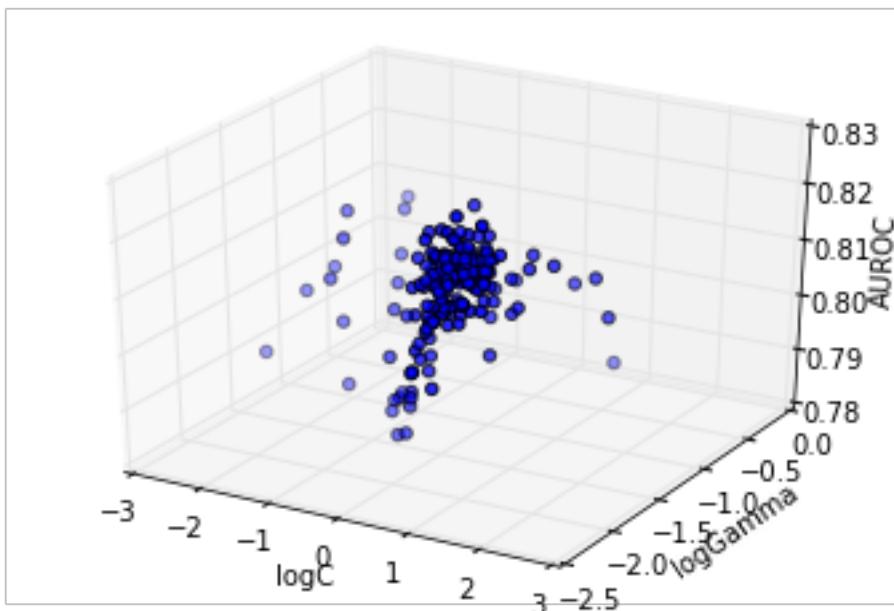
```

```
<matplotlib.text.Text at 0x7f1f2a5c3190>
```



```
cutoff = 0.95 * info.optimum
fig = plt.figure()
ax = fig.add_subplot(111, projection='3d')
ax.scatter(xs=df[df.value > cutoff]['logC'],
           ys=df[df.value > cutoff]['logGamma'],
           zs=df[df.value > cutoff]['value'])
ax.set_xlabel('logC')
ax.set_ylabel('logGamma')
ax.set_zlabel('AUROC')
```

<matplotlib.text.Text at 0x7f1f26ae5410>



Lets further examine the area close to the optimum, that is the 95% region. We will examine a 50x50 grid in this region.

```

minlogc = min(df[df.value > cutoff]['logC'])
maxlogc = max(df[df.value > cutoff]['logC'])
minloggamma = min(df[df.value > cutoff]['logGamma'])
maxloggamma = max(df[df.value > cutoff]['logGamma'])

_, info_new, _ = optunity.maximize(svm_rbf_tuned_auroc, num_evals=2500,
                                    logC=[minlogc, maxlogc],
                                    logGamma=[minloggamma, maxloggamma],
                                    solver_name='grid search')

```

Make a new data frame of the call log for easy manipulation.

```
df_new = optunity.call_log2dataframe(info_new.call_log)
```

Determine the region in which we will do a standard grid search to obtain contours.

```

reshape = lambda x: np.reshape(x, (50, 50))
logcs = reshape(df_new['logC'])
loggammams = reshape(df_new['logGamma'])
values = reshape(df_new['value'])

```

Now, lets look at a contour plot to see whether or not the response surface is smooth.

```

levels = np.arange(0.97, 1.0, 0.005) * info_new.optimum
CS = plt.contour(logcs, loggammams, values, levels=levels)
plt.clabel(CS, inline=1, fontsize=10)
plt.title('Contours of SVM tuning response surface')
plt.xlabel('log C')
plt.ylabel('log gamma')

```

```

/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
    nlc.append(np.r_[xy2, lc[I[1]:]])

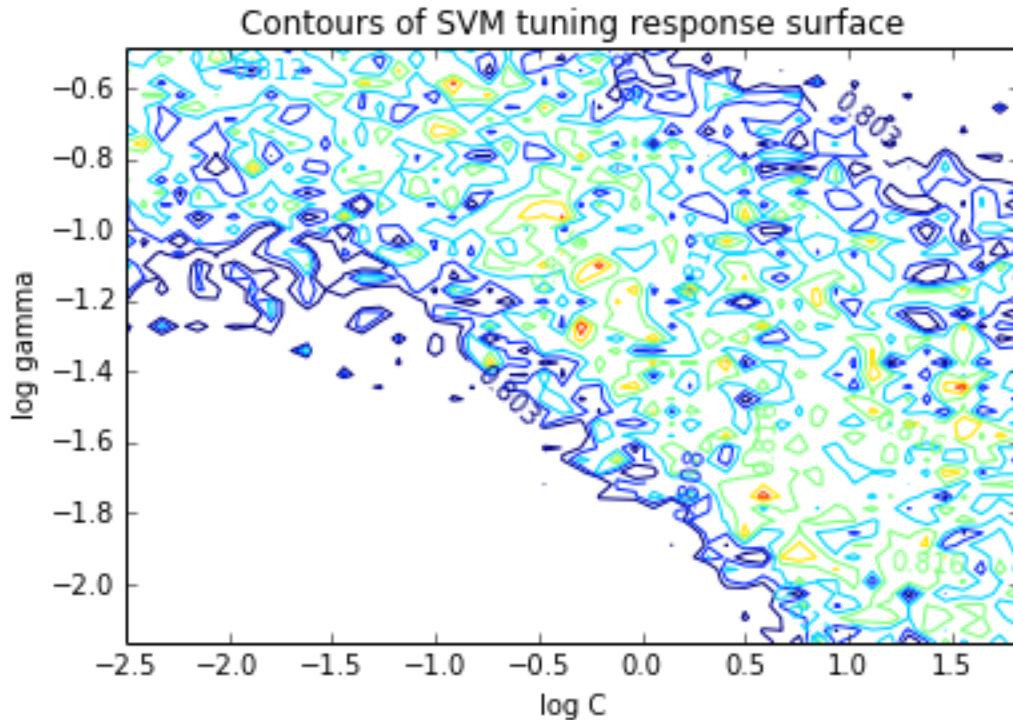
```

```

nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:483: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[lc[:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:485: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:]])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])
/usr/lib/pymodules/python2.7/matplotlib/contour.py:479: DeprecationWarning: using a non-integer number
  nlc.append(np.r_[xy2, lc[I[1]:I[0] + 1], xy1])

```

<matplotlib.text.Text at 0x7f1f26ac3090>



Clearly, this response surface is filled with local minima. This is a general observation in automated hyperparameter optimization, and is one of the key reasons we need robust solvers. If there were no local minima, a simple gradient-descent-like solver would do the trick.

### 5.4.6 OpenCV: optical character recognition

We will use OpenCV (<http://www.opencv.org/>) for optical character recognition (OCR) using support vector machine (SVM) classifiers. This example is based on OpenCV's digit tutorial (available in `OPENCV_ROOT/samples/python2/digits.py`).

We start with the necessary imports.

```
import cv2
import numpy as np
import opportunity
import opportunity.metrics

# comment the line below when running the notebook yourself
%matplotlib inline
import matplotlib.pyplot as plt

# finally, we import some convenience functions to avoid bloating this tutorial
# these are available in the `OPTUNITY/notebooks` folder
import opencv_utility_functions as util
```

Load the data, which is available in the `OPPORTUNITY/notebooks` directory. This data file is a direct copy from OpenCV's example.

```
digits, labels = util.load_digits("opencv-digits.png")
```

```
loading "opencv-digits.png" ...
```

We preprocess the data and make a train/test split.

```
rand = np.random.RandomState(100)
shuffle = rand.permutation(len(digits))
digits, labels = digits[shuffle], labels[shuffle]

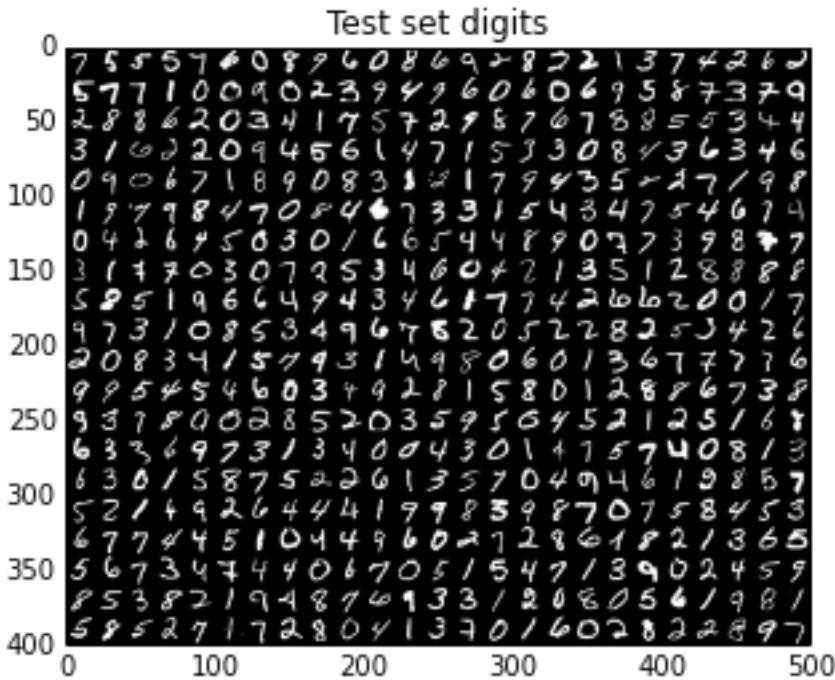
digits2 = map(util.deskew, digits)
samples = util.preprocess_hog(digits2)

train_n = int(0.9*len(samples))

digits_train, digits_test = np.split(digits2, [train_n])
samples_train, samples_test = np.split(samples, [train_n])
labels_train, labels_test = np.split(labels, [train_n])

test_set_img = util.mosaic(25, digits[train_n:])
plt.title('Test set digits')
plt.imshow(test_set_img, cmap='binary_r')
```

```
<matplotlib.image.AxesImage at 0x7fd3ee2e4990>
```



Now, it's time to construct classifiers. We will use a SVM classifier with an RBF kernel, i.e.  $\kappa(\mathbf{u}, \mathbf{v}) = \exp(-\gamma\|\mathbf{u} - \mathbf{v}\|^2)$ . Such an SVM has two hyperparameters that must be optimized, namely the misclassification penalty  $C$  and kernel parameter  $\gamma$ .

We start with an SVM with default parameters, which in this case means:  $C = 1$  and  $\gamma = 0.5$ .

```
model_default = util.SVM()
model_default.train(samples_train, labels_train)
vis_default, err_default = util.evaluate_model(model_default, digits_test, samples_test,
plt.title('unoptimized SVM test set performance')
plt.imshow(vis_default)
```

```
error: 4.20 %
confusion matrix:
[[50  0  0  0  0  0  0  0  0  0]
 [ 0 45  0  0  1  0  0  0  0  0]
 [ 0  0 43  1  0  0  0  0  3  1]
 [ 0  0  0 48  0  1  0  1  0  0]
 [ 0  0  0  0 52  0  2  0  1  1]
 [ 0  0  0  0  0 49  0  0  0  0]
 [ 0  0  0  0  0  0 48  0  0  0]
 [ 0  0  1  3  0  0  0 54  0  0]
 [ 0  1  0  0  0  2  0  0 50  0]
 [ 1  0  0  0  0  0  1  0 40]]
```

```
<matplotlib.image.AxesImage at 0x7fd3ee288ad0>
```



Next, we will construct a model with optimized hyperparameters. First we need to build Opportunity's objective function. We will use 5-fold cross-validated error rate as loss function, which we will minimize.

```
@opportunity.cross_validated(x=samples_train, y=labels_train, num_folds=5)
def svm_error_rate(x_train, y_train, x_test, y_test, C, gamma):
    model = util.SVM(C=C, gamma=gamma)
    model.train(x_train, y_train)
    resp = model.predict(x_test)
    error_rate = (y_test != resp).mean()
    return error_rate
```

We will use Opportunity's default solver to optimize the error rate given  $0 < C < 5$  and  $0 < \gamma < 10$  and up to 50 function evaluations. This may take a while.

```
optimal_parameters, details, _ = opportunity.minimize(svm_error_rate, num_evals=50,
                                                     C=[0, 5], gamma=[0, 10])
# the above line can be parallelized by adding `pmap=opportunity.pmap`
# however this is incompatible with IPython

print("Optimal parameters: C=%1.3f, gamma=%1.3f" % (optimal_parameters['C'], optimal_parameters['gamma']))
print("Cross-validated error rate: %1.3f" % details.optimum)
```

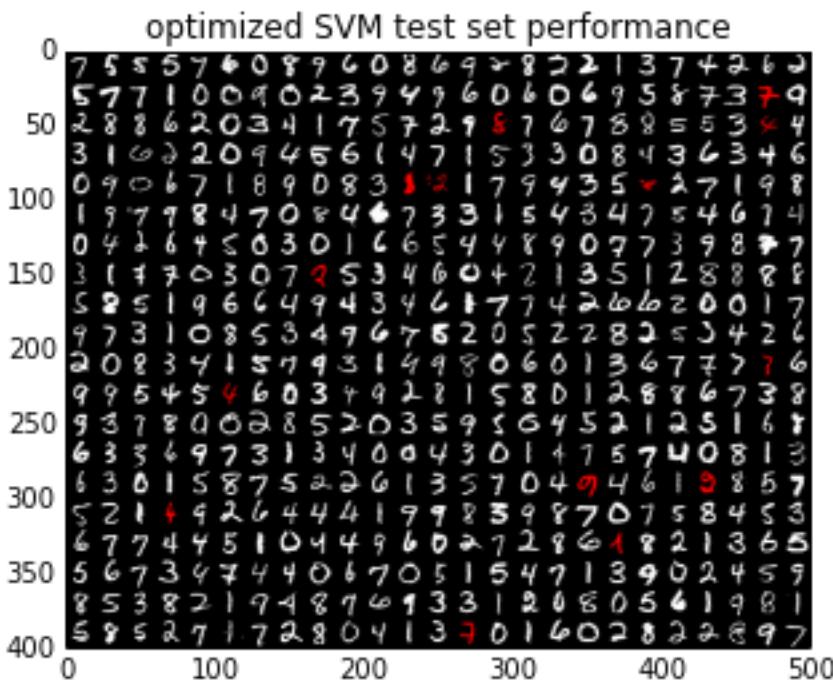
```
Optimal parameters: C=1.798, gamma=6.072
Cross-validated error rate: 0.025
```

Finally, we train a model with the optimized parameters and determine its test set performance.

```
model_opt = util.SVM(**optimal_parameters)
model_opt.train(samples_train, labels_train)
vis_opt, err_opt = util.evaluate_model(model_opt, digits_test, samples_test, labels_test)
plt.title('optimized SVM test set performance')
plt.imshow(vis_opt)
```

```
error: 2.80 %
confusion matrix:
[[50  0  0  0  0  0  0  0  0  0]
 [ 0 45  0  0  1  0  0  0  0  0]
 [ 0  0 44  1  0  0  0  1  1  1]
 [ 0  0  0 49  0  0  0  1  0  0]
 [ 0  0  1  0 53  0  2  0  0  0]
 [ 0  0  0  0  0 49  0  0  0  0]
 [ 0  0  0  0  0  0 48  0  0  0]
 [ 0  0  2  0  0  0  0 56  0  0]
 [ 0  1  0  0  0  1  0  0 51  0]
 [ 0  0  0  0  0  0  1  0 41]]
```

```
<matplotlib.image.AxesImage at 0x7fd3ee4da050>
```



```
print("Reduction in error rate by optimizing hyperparameters: %1.1f%%" % (100.0 - 100.0 * err_opt / err))
```

```
Reduction in error rate by optimizing hyperparameters: 33.3%
```

## 5.4.7 sklearn: automated learning method selection and tuning

In this tutorial we will show how to use Opportunity in combination with sklearn to classify the digit recognition data set available in sklearn. The cool part is that we will use Opportunity to choose the best approach from a set of available learning algorithms and optimize hyperparameters in one go. We will use the following learning algorithms:

- k-nearest neighbour
- SVM
- Naive Bayes
- Random Forest

For simplicity, we will focus on a binary classification task, namely digit 3 versus digit 9. We start with the necessary imports and create the data set.

```
import opportunity
import opportunity.metrics
import numpy as np

# k nearest neighbours
from sklearn.neighbors import KNeighborsClassifier
# support vector machine classifier
from sklearn.svm import SVC
# Naive Bayes
from sklearn.naive_bayes import GaussianNB
# Random Forest
from sklearn.ensemble import RandomForestClassifier

from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 3
negative_digit = 9

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)
```

For the SVM model we will let Opportunity optimize the kernel family, choosing from linear, polynomial and RBF. We start by creating a convenience functions for SVM training that handles this:

```
def train_svm(data, labels, kernel, C, gamma, degree, coef0):
    """A generic SVM training function, with arguments based on the chosen kernel."""
    if kernel == 'linear':
        model = SVC(kernel=kernel, C=C)
    elif kernel == 'poly':
        model = SVC(kernel=kernel, C=C, degree=degree, coef0=coef0)
    elif kernel == 'rbf':
        model = SVC(kernel=kernel, C=C, gamma=gamma)
    else:
        raise ArgumentError("Unknown kernel function: %s" % kernel)
    model.fit(data, labels)
    return model
```

Every learning algorithm has its own hyperparameters:

- k-NN:  $1 < n\_neighbors < 5$  the number of neighbours to use
- SVM: kernel family and misclassification penalty, we will make the penalty contingent on the family. Per kernel family, we have different hyperparameters:
  - polynomial kernel:  $0 < C < 50$ ,  $2 < degree < 5$  and  $0 < coef0 < 1$
  - linear kernel:  $0 < C < 2$ ,
  - RBF kernel:  $0 < C < 10$  and  $0 < gamma < 1$
- naive Bayes: no hyperparameters

- random forest:
  - $10 < n\_estimators < 30$ : number of trees in the forest
  - $5 < max\_features < 20$ : number of features to consider for each split

This translates into the following search space:

```
search = {'algorithm': {'k-nn': {'n_neighbors': [1, 5]},  
                      'SVM': {'kernel': {'linear': {'C': [0, 2]},  
                               'rbf': {'gamma': [0, 1], 'C': [0, 10]},  
                               'poly': {'degree': [2, 5], 'C': [0, 50], 'coef0': [0, 1]}  
                           }  
                     },  
                      'naive-bayes': None,  
                      'random-forest': {'n_estimators': [10, 30],  
                                       'max_features': [5, 20]}}  
        }  
    }
```

We also need an objective function that can properly orchestrate everything. We will choose the best model based on area under the ROC curve in 5-fold cross-validation.

```
@opportunity.cross_validated(x=data, y=labels, num_folds=5)  
def performance(x_train, y_train, x_test, y_test,  
                 algorithm, n_neighbors=None, n_estimators=None, max_features=None,  
                 kernel=None, C=None, gamma=None, degree=None, coef0=None):  
    # fit the model  
    if algorithm == 'k-nn':  
        model = KNeighborsClassifier(n_neighbors=int(n_neighbors))  
        model.fit(x_train, y_train)  
    elif algorithm == 'SVM':  
        model = train_svm(x_train, y_train, kernel, C, gamma, degree, coef0)  
    elif algorithm == 'naive-bayes':  
        model = GaussianNB()  
        model.fit(x_train, y_train)  
    elif algorithm == 'random-forest':  
        model = RandomForestClassifier(n_estimators=int(n_estimators),  
                                       max_features=int(max_features))  
        model.fit(x_train, y_train)  
    else:  
        raise ArgumentError('Unknown algorithm: %s' % algorithm)  
  
    # predict the test set  
    if algorithm == 'SVM':  
        predictions = model.decision_function(x_test)  
    else:  
        predictions = model.predict_proba(x_test)[:, 1]  
  
    return opportunity.metrics.roc_auc(y_test, predictions, positive=True)
```

Lets do a simple test run of this fancy objective function.

```
performance(algorithm='k-nn', n_neighbors=3)
```

```
0.9547920006472639
```

Seems okay! Now we can let Opportunity do its magic with a budget of 300 tries.

```
optimal_configuration, info, _ = opportunity.maximize_structured(performance,  
                                                               search_space=search,
```

```

        num_evals=300)
print(optimal_configuration)
print(info.optimum)

{'kernel': 'poly', 'C': 38.5498046875, 'algorithm': 'SVM', 'degree': 3.88525390625, 'n_neighbors': None,
 0.979302949566

```

Finally, lets make the results a little bit more readable. All dictionary items in optimal\_configuration with value None can be removed.

```

solution = dict([(k, v) for k, v in optimal_configuration.items() if v is not None])
print('Solution\n=====')
print("\n".join(map(lambda x: "%s %t %s" % (x[0], str(x[1])), solution.items())))

```

```

Solution
=====
kernel      poly
C          38.5498046875
coef0       0.71826171875
degree      3.88525390625
algorithm   SVM

```

## 5.4.8 sklearn: SVM classification

In this example we will use Optunity to optimize hyperparameters for a support vector machine classifier (SVC) in scikit-learn. We will learn a model to distinguish digits 8 and 9 in the MNIST data set in two settings

- tune SVM with RBF kernel
- tune SVM with RBF, polynomial or linear kernel, that is choose the kernel function and its hyperparameters at once

```

import optunity
import optunity.metrics

# comment this line if you are running the notebook
import sklearn.svm
import numpy as np

```

Create the data set: we use the MNIST data set and will build models to distinguish digits 8 and 9.

```

from sklearn.datasets import load_digits
digits = load_digits()
n = digits.data.shape[0]

positive_digit = 8
negative_digit = 9

positive_idx = [i for i in range(n) if digits.target[i] == positive_digit]
negative_idx = [i for i in range(n) if digits.target[i] == negative_digit]

# add some noise to the data to make it a little challenging
original_data = digits.data[positive_idx + negative_idx, ...]
data = original_data + 5 * np.random.randn(original_data.shape[0], original_data.shape[1])
labels = [True] * len(positive_idx) + [False] * len(negative_idx)

```

First, lets see the performance of an SVC with default hyperparameters.

```
# compute area under ROC curve of default parameters
@opportunity.cross_validated(x=data, y=labels, num_folds=5)
def svm_default_auroc(x_train, y_train, x_test, y_test):
    model = sklearn.svm.SVC().fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = opportunity.metrics.roc_auc(y_test, decision_values)
    return auc

svm_default_auroc()
```

```
0.7328666183635757
```

## Tune SVC with RBF kernel

In order to use Opportunity to optimize hyperparameters, we start by defining the objective function. We will use 5-fold cross-validated area under the ROC curve. For now, lets restrict ourselves to the RBF kernel and optimize  $C$  and  $\gamma$ .

We start by defining the objective function `svm_rbf_tuned_auroc()`, which accepts  $C$  and  $\gamma$  as arguments.

```
#we will make the cross-validation decorator once, so we can reuse it later for the other tuning tasks
# by reusing the decorator, we get the same folds etc.
cv_decorator = opportunity.cross_validated(x=data, y=labels, num_folds=5)

def svm_rbf_tuned_auroc(x_train, y_train, x_test, y_test, C, logGamma):
    model = sklearn.svm.SVC(C=C, gamma=10 ** logGamma).fit(x_train, y_train)
    decision_values = model.decision_function(x_test)
    auc = opportunity.metrics.roc_auc(y_test, decision_values)
    return auc

svm_rbf_tuned_auroc = cv_decorator(svm_rbf_tuned_auroc)
# this is equivalent to the more common syntax below
# @opportunity.cross_validated(x=data, y=labels, num_folds=5)
# def svm_rbf_tuned_auroc...

svm_rbf_tuned_auroc(C=1.0, logGamma=0.0)
```

```
0.5
```

Now we can use Opportunity to find the hyperparameters that maximize AUROC.

```
optimal_rbf_pars, info, _ = opportunity.maximize(svm_rbf_tuned_auroc, num_evals=150, C=[0, 10], logGamma=[0, 1])
# when running this outside of IPython we can parallelize via opportunity.pmap
# optimal_rbf_pars, _, _ = opportunity.maximize(svm_rbf_tuned_auroc, 150, C=[0, 10], gamma=[0, 0.1], pmapping=True)

print("Optimal parameters: " + str(optimal_rbf_pars))
print("AUROC of tuned SVM with RBF kernel: %1.3f" % info.optimum)
```

```
Optimal parameters: {'logGamma': -3.0716796875000005, 'C': 3.3025997497032007}
AUROC of tuned SVM with RBF kernel: 0.987
```

We can turn the call log into a pandas dataframe to efficiently inspect the solver trace.

```
import pandas
df = opportunity.call_log2dataframe(info.call_log)
```

Lets look at the best 20 sets of hyperparameters to make sure the results are somewhat stable.

```
df.sort('value', ascending=False) [:10]
```

## Tune SVC without deciding the kernel in advance

In the previous part we choose to use an RBF kernel. Even though the RBF kernel is known to work well for a large variety of problems (and yielded good accuracy here), our choice was somewhat arbitrary.

We will now use Optunity's conditional hyperparameter optimization feature to optimize over all kernel functions and their associated hyperparameters at once. This requires us to define the search space.

```
space = {'kernel': {'linear': {'C': [0, 2]},  
                  'rbf': {'logGamma': [-5, 0], 'C': [0, 10]},  
                  'poly': {'degree': [2, 5], 'C': [0, 5], 'coef0': [0, 2]}  
                }  
}
```

We will also have to modify the objective function to cope with conditional hyperparameters. The reason we need to do this explicitly is because scikit-learn doesn't like dealing with `None` values for irrelevant hyperparameters (e.g. `degree` when using an RBF kernel). Optunity will set all irrelevant hyperparameters in a given set to `None`.

```
def train_model(x_train, y_train, kernel, C, logGamma, degree, coef0):  
    """A generic SVM training function, with arguments based on the chosen kernel."""  
    if kernel == 'linear':  
        model = sklearn.svm.SVC(kernel=kernel, C=C)  
    elif kernel == 'poly':  
        model = sklearn.svm.SVC(kernel=kernel, C=C, degree=degree, coef0=coef0)  
    elif kernel == 'rbf':  
        model = sklearn.svm.SVC(kernel=kernel, C=C, gamma=10 ** logGamma)  
    else:  
        raise ValueError("Unknown kernel function: %s" % kernel)  
    model.fit(x_train, y_train)  
    return model  
  
def svm_tuned_auroc(x_train, y_train, x_test, y_test, kernel='linear', C=0, logGamma=0, degree=0, coef0=0):  
    model = train_model(x_train, y_train, kernel, C, logGamma, degree, coef0)  
    decision_values = model.decision_function(x_test)  
    return optunity.metrics.roc_auc(y_test, decision_values)  
  
svm_tuned_auroc = cv_decorator(svm_tuned_auroc)
```

Now we are ready to go and optimize both kernel function and associated hyperparameters!

```
optimal_svm_pars, info, _ = optunity.maximize_structured(svm_tuned_auroc, space, num_evals=150)  
print("Optimal parameters" + str(optimal_svm_pars))  
print("AUROC of tuned SVM: %.3f" % info.optimum)
```

```
Optimal parameters{'kernel': 'rbf', 'C': 3.634209495387873, 'coef0': None, 'degree': None, 'logGamma': 0.0}  
AUROC of tuned SVM: 0.990
```

Again, we can have a look at the best sets of hyperparameters based on the call log.

```
df = optunity.call_log2dataframe(info.call_log)  
df.sort('value', ascending=False)
```

## 5.4.9 sklearn: SVM regression

In this example we will show how to use Opportunity to tune hyperparameters for support vector regression, more specifically:

- measure empirical improvements through nested cross-validation
- optimizing hyperparameters for a given family of kernel functions
- determining the optimal model without choosing the kernel in advance

```
import math
import itertools
import opportunity
import opportunity.metrics
import sklearn.svm
```

We start by creating the data set. We use sklearn's diabetes data.

```
from sklearn.datasets import load_diabetes
diabetes = load_diabetes()
n = diabetes.data.shape[0]

data = diabetes.data
targets = diabetes.target
```

### Nested cross-validation

Nested cross-validation is used to estimate generalization performance of a full learning pipeline, which includes optimizing hyperparameters. We will use three folds in the outer loop.

When using default hyperparameters, there is no need for an inner cross-validation procedure. However, if we want to include tuning in the learning pipeline, the inner loop is used to determine generalization performance with optimized hyperparameters.

We start by measuring generalization performance with default hyperparameters.

```
# we explicitly generate the outer_cv decorator so we can use it twice
outer_cv = opportunity.cross_validated(x=data, y=targets, num_folds=3)

def compute_mse_standard(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and default hyperparameters.
    """
    model = sklearn.svm.SVR().fit(x_train, y_train)
    predictions = model.predict(x_test)
    return opportunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_standard = outer_cv(compute_mse_standard)
```

`compute_mse_standard()` returns a three-fold cross-validation estimate of MSE for an SVR with default hyperparameters.

```
compute_mse_standard()
```

```
6190.481497665955
```

We will create a function that returns MSE based on optimized hyperparameters, where we choose a polynomial kernel in advance.

```

def compute_mse_poly_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @opportunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, C, degree, coef0):
        model = sklearn.svm.SVR(C=C, degree=degree, coef0=coef0, kernel='poly').fit(x_train, y_train)
        predictions = model.predict(x_test)
        return opportunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = opportunity.minimize(tune_cv, 150, C=[1000, 20000], degree=[2, 5],
                                                coef0=[0, 1])
    print("optimal hyperparameters: " + str(optimal_pars))

    tuned_model = sklearn.svm.SVR(kernel='poly', **optimal_pars).fit(x_train, y_train)
    predictions = tuned_model.predict(x_test)
    return opportunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_poly_tuned = outer_cv(compute_mse_poly_tuned)

```

`compute_mse_poly_tuned()` returns a three-fold cross-validation estimate of MSE for an SVR with RBF kernel with tuned hyperparameters  $1000 < C < 20000$ ,  $2 < \text{degree} < 5$  and  $0 < \text{coef0} < 1$  with a budget of 150 function evaluations. Each tuple of hyperparameters is evaluated using twice-iterated 5-fold cross-validation.

```
compute_mse_poly_tuned()
```

```

optimal hyperparameters: {'C': 12078.673881034498, 'coef0': 0.5011052085197018, 'degree': 4.608902814}
optimal hyperparameters: {'C': 14391.165364583334, 'coef0': 0.17313151041666666, 'degree': 2.358261771}
optimal hyperparameters: {'C': 11713.456382191061, 'coef0': 0.49836486667796476, 'degree': 4.61607799}

```

```
3047.035965991627
```

The polynomial kernel yields pretty good results when optimized, but maybe we can do even better with an RBF kernel.

```

def compute_mse_rbf_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @opportunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, C, gamma):
        model = sklearn.svm.SVR(C=C, gamma=gamma).fit(x_train, y_train)
        predictions = model.predict(x_test)
        return opportunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = opportunity.minimize(tune_cv, 150, C=[1, 100], gamma=[0, 50])
    print("optimal hyperparameters: " + str(optimal_pars))

    tuned_model = sklearn.svm.SVR(**optimal_pars).fit(x_train, y_train)
    predictions = tuned_model.predict(x_test)
    return opportunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_rbf_tuned = outer_cv(compute_mse_rbf_tuned)

```

`compute_mse_rbf_tuned()` returns a three-fold cross-validation estimate of MSE for an SVR with RBF kernel with tuned hyperparameters  $1 < C < 100$  and  $0 < \gamma < 5$  with a budget of 150 function evaluations. Each tuple of

hyperparameters is evaluated using twice-iterated 5-fold cross-validation.

```
compute_mse_rbf_tuned()
```

```
optimal hyperparameters: {'C': 21.654003906250026, 'gamma': 16.536188056152554}
optimal hyperparameters: {'C': 80.89867187499999, 'gamma': 3.2346692538501784}
optimal hyperparameters: {'C': 19.35431640625002, 'gamma': 22.083848774716085}
```

```
2990.8572696483493
```

Woop! Seems like an RBF kernel is a good choice. An optimized RBF kernel leads to a 50% reduction in MSE compared to the default configuration.

## Determining the kernel family during tuning

In the previous part we've seen that the choice of kernel and its parameters significantly impact performance. However, testing every kernel family separately is cumbersome. It's better to let Opportunity do the work for us.

Opportunity can optimize conditional search spaces, here the kernel family and depending on which family the hyperparameterization ( $\gamma$ , degree, coef0, ...). We start by defining the search space (we will try the linear, polynomial and RBF kernel).

```
space = {'kernel': {'linear': {'C': [0, 100]},
                    'rbf': {'gamma': [0, 50], 'C': [1, 100]},
                    'poly': {'degree': [2, 5], 'C': [1000, 20000], 'coef0': [0, 1]}}
                }
```

Now we do nested cross-validation again.

```
def compute_mse_all_tuned(x_train, y_train, x_test, y_test):
    """Computes MSE of an SVR with RBF kernel and optimized hyperparameters."""

    # define objective function for tuning
    @opportunity.cross_validated(x=x_train, y=y_train, num_iter=2, num_folds=5)
    def tune_cv(x_train, y_train, x_test, y_test, kernel, C, gamma, degree, coef0):
        if kernel == 'linear':
            model = sklearn.svm.SVR(kernel=kernel, C=C)
        elif kernel == 'poly':
            model = sklearn.svm.SVR(kernel=kernel, C=C, degree=degree, coef0=coef0)
        elif kernel == 'rbf':
            model = sklearn.svm.SVR(kernel=kernel, C=C, gamma=gamma)
        else:
            raise ArgumentError("Unknown kernel function: %s" % kernel)
        model.fit(x_train, y_train)

        predictions = model.predict(x_test)
        return opportunity.metrics.mse(y_test, predictions)

    # optimize parameters
    optimal_pars, _, _ = opportunity.minimize_structured(tune_cv, num_evals=150, search_space=space)

    # remove hyperparameters with None value from optimal pars
    for k, v in optimal_pars.items():
        if v is None: del optimal_pars[k]
    print("optimal hyperparameters: " + str(optimal_pars))

    tuned_model = sklearn.svm.SVR(**optimal_pars).fit(x_train, y_train)
```

```
predictions = tuned_model.predict(x_test)
return opportunity.metrics.mse(y_test, predictions)

# wrap with outer cross-validation
compute_mse_all_tuned = outer_cv(compute_mse_all_tuned)
```

And now the kernel family will be optimized along with its hyperparameterization.

```
compute_mse_all_tuned()
```

```
optimal hyperparameters: {'kernel': 'rbf', 'C': 33.70116043112164, 'gamma': 16.32317353448437}
optimal hyperparameters: {'kernel': 'rbf', 'C': 58.11404170763237, 'gamma': 26.45349823062099}
optimal hyperparameters: {'kernel': 'poly', 'C': 14964.421875843143, 'coef0': 0.5127175861493205, 'd
```

```
3107.625560844859
```

It looks like the RBF and polynomial kernel are competitive for this problem.

To contribute examples, please send us a pull request on [Github](#).

## 5.5 Using Other Environments

### 5.5.1 MATLAB

In this page we briefly discuss the MATLAB wrapper, which provides most of Opportunity's functionality. For a general overview, we recommend reading the [User Guide](#).

For installation instructions, please refer to [Installing Opportunity](#).

Whenever the Python API requires a dictionary, we use a MATLAB *struct*. As MATLAB has no keyword arguments, we use *varargs* with the convention of *<name>, <value>*. Please refer to each function's help for details on how to use it.

**Warning:** Since MATLAB does not have keyword arguments, every objective function you define should accept a single struct argument, whose fields represent hyperparameters.

Example, to model a function  $f(x, y) = x + y$ :

```
f = @(pars) pars.x + pars.y;
```

For MATLAB, the following main features are provided:

- *Manual*
- *Optimizing hyperparameters*
- *Cross-validation*

The file *<opportunity>/wrappers/matlab/opportunity\_example.m* contains code for all the functionality that is available in MATLAB.

---

**Note:** If you experience any issues with the MATLAB wrapper, create a global variable *DEBUG\_OPTUNITY* and set its value to true. This will show all debugging output:

```
global DEBUG_OPTUNITY
DEBUG_OPTUNITY=true;
```

Please submit this output as an issue at <https://github.com/claesem/optunity/issues/new> to obtain help.

If MATLAB hangs while using Optunity there is a communication issue. This should not occur, if you encounter this please file an issue at <https://github.com/claesem/optunity/issues/new>. Currently the only way out of this is to kill MATLAB.

---

## Manual

To obtain the manual of all solvers and a list of available solver names, use `optunity.manual()`. If you specify a solver name, its manual will be printed out.

You can verify whether or not a certain solver, for instance `cma-es`, is available like this:

```
solvers = optunity.manual();
available = any(arrayfun(@(x) strcmp(x, 'cma-es'), solvers));
```

## Optimizing hyperparameters

The MATLAB wrapper offers `optunity.maximize()`, `optunity.minimize()` and `optunity.optimize()`. These provide the same functionality as their Python equivalents.

The following code fragment shows how to optimize a simple function  $f$  with Random Search within the box  $-4 < x < 4$  and  $-5 < y < 5$  and 200 evaluations:

```
offx = rand();
offy = rand();
f = @(pars) - (offx+pars.x)^2 - (offy+pars.y)^2;

[max_solution, max_details, max_solver] = optunity.maximize(f, 200, ...
    'solver_name', 'random search', 'x', [-4, 4], 'y', [-5, 5]);
```

If you want to use `optunity.optimize()`, you must create the solver in advance using `optunity.make_solver()`. This will return an `optunity.Solver` object or return an error message:

```
f = @(pars) pars.x + pars.y
rnd_solver = optunity.make_solver('random search', 'x', [-5, 5], ...
    'y', [-5, 5], 'num_evals', 400);
[rnd_solution, rnd_details] = optunity.optimize(rnd_solver, f);
```

## Differences between Python and MATLAB version of `optimize`

The MATLAB version has an extra argument `constraints` where you can specify domain constraints (the MATLAB wrapper has no equivalent of `optunity.wrap_constraints()`). Constraints are communicated as a struct with the correct field names and corresponding values (more info at [Domain constraints](#)).

As an example, to add constraints  $x < 3$  and  $y \geq 0$ , we use the following code:

```
constraints = struct('ub_o', struct('x', 3), 'lb_c', struct('y', 0));
```

Information about the solvers can be obtained at [Solver overview](#). To learn about the specific parameter names for each solver, check out the `optunity.solvers`.

## Cross-validation

Two functions are provided for cross-validation:

- `opportunity.generate_folds()`: generates a set of cross-validation folds
- `opportunity.cross_validate()`: function decorator, to perform cross-validation with specified function

Both functions can deal with strata and clusters. You can specify these as a cell array of vectors of indices:

```
strata = {[1,2,3], [6,7,8,9]};
folds = opportunity.generate_folds(20, 'num_folds', 10, 'num_iter', 2, 'strata', strata);
```

Cross-validation folds are returned as a matrix of `num_instances * num_iter` with entries ranging from 1 to `num_folds` to indicate the fold each instance belongs to per iteration.

`opportunity.cross_validate()` requires a function handle as its first argument. This is the function that will be decorated, which must have the following first arguments: `x_train` and `x_test` (if unsupervised) or `x_train, y_train, x_test, y_test`.

As an example, assume we have a function `opportunity_cv_fun(x_train, x_test, pars)`:

```
function [ result ] = cv_fun( x_train, x_test, pars )
    disp('training set:');
    disp(x_train');
    disp('test set:');
    disp(x_test');

    result = -pars.x^2 - pars.y^2;
end
```

This must be decorated with cross-validation, for instance:

```
x = (1:10)';
cvf = opportunity.cross_validate(@opportunity_cv_fun, x);

% evaluate the function: this will return a cross-validation result
performance = cvf(struct('x',1,'y',2));
```

**Warning:** After decorating with cross-validation, the objective function should have a single argument, namely a struct of hyperparameters.

## 5.5.2 GNU Octave

In this page we briefly discuss the Octave wrapper, which provides most of Opportunity's functionality. For a general overview, we recommend reading the [User Guide](#).

For installation instructions, please refer to [Installing Opportunity](#).

Whenever the Python API requires a dictionary, we use an Octave *struct*. As Octave has no keyword arguments, we use *varargs* with the convention of `<name>, <value>`. Please refer to each function's help for details on how to use it.

**Warning:** Since Octave does not have keyword arguments, every objective function you define should accept a single struct argument, whose fields represent hyperparameters.

Example, to model a function  $f(x,y) = x + y$ :

```
f = @(pars) pars.x + pars.y;
```

For octave, the following main features are provided:

- *Manual*
- *Optimizing hyperparameters*
- *Cross-validation*

The file `<optunity>/wrappers/octave/optunity/example/optunity_example.m` contains code for all the functionality that is available in octave.

**Note:** If you experience any issues with the octave wrapper, create a global variable `DEBUG_OPTUNITY` and set its value to true. This will show all debugging output:

```
global DEBUG_OPTUNITY
DEBUG_OPTUNITY=true;
```

Please submit this output as an issue at <https://github.com/claesennm/optunity/issues/new> to obtain help.

If octave hangs while using Optunity there is a communication issue. This should not occur, if you encounter this please file an issue at <https://github.com/claesennm/optunity/issues/new>. Currently the only way out of this is to kill octave.

## Manual

To obtain the manual of all solvers and a list of available solver names, use `manual()`. If you specify a solver name, its manual will be printed out.

You can verify whether or not a certain solver, for instance `cma-es`, is available like this:

```
solvers = manual();
available = any(arrayfun(@(x) strcmp(x, 'cma-es'), solvers));
```

## Optimizing hyperparameters

The octave wrapper offers `maximize()`, `minimize()` and `optimize()`. These provide the same functionality as their Python equivalents.

The following code fragment shows how to optimize a simple function  $f$  with Random Search within the box  $-4 < x < 4$  and  $-5 < y < 5$  and 200 evaluations:

```
offx = rand();
offy = rand();
f = @(pars) - (offx+pars.x)^2 - (offy+pars.y)^2;

[max_solution, max_details, max_solver] = maximize(f, 200, ...
    'solver_name', 'random search', 'x', [-4, 4], 'y', [-5, 5]);
```

If you want to use `optimize()`, you must create the solver in advance using `make_solver()`. This will return a *Solver* object or return an error message:

```
f = @(pars) pars.x + pars.y
rnd_solver = make_solver('random search', 'x', [-5, 5], ...
    'y', [-5, 5], 'num_evals', 400);
[rnd_solution, rnd_details] = optimize(rnd_solver, f);
```

### Differences between Python and octave version of *optimize*

The octave version has an extra argument *constraints* where you can specify domain constraints (the octave wrapper has no equivalent of `opportunity.wrap_constraints()`). Constraints are communicated as a struct with the correct field names and corresponding values (more info at [Domain constraints](#)).

As an example, to add constraints  $x < 3$  and  $y \geq 0$ , we use the following code:

```
constraints = struct('ub_o', struct('x', 3), 'lb_c', struct('y', 0));
```

Information about the solvers can be obtained at [Solver overview](#). To learn about the specific parameter names for each solver, check out the `opportunity.solvers`.

### Cross-validation

Two functions are provided for cross-validation:

- `generate_folds()`: generates a set of cross-validation folds
- `cross_validate()`: function decorator, to perform cross-validation with specified function

Both functions can deal with strata and clusters. You can specify these as a cell array of vectors of indices:

```
strata = {[1,2,3], [6,7,8,9]};  
folds = generate_folds(20, 'num_folds', 10, 'num_iter', 2, 'strata', strata);
```

Cross-validation folds are returned as a matrix of *num\_instances* \* *num\_iter* with entries ranging from 1 to *num\_folds* to indicate the fold each instance belongs to per iteration.

`cross_validate()` requires a function handle as its first argument. This is the function that will be decorated, which must have the following first arguments: *x\_train* and *x\_test* (if unsupervised) or *x\_train*, *y\_train*, *x\_test*, *y\_test*.

As an example, assume we have a function `opportunity_cv_fun(x_train, x_test, pars)`:

```
function [ result ] = opportunity_cv_fun( x_train, x_test, pars )  
disp('training set:');  
disp(x_train');  
disp('test set:');  
disp(x_test');  
  
result = -pars.x^2 - pars.y^2;  
end
```

This must be decorated with cross-validation, for instance:

```
x = (1:10)';  
cvf = cross_validate(@opportunity_cv_fun, x);  
  
% evaluate the function: this will return a cross-validation result  
performance = cvf(struct('x',1,'y',2));
```

**Warning:** After decorating with cross-validation, the objective function should have a single argument, namely a struct of hyperparameters.

### 5.5.3 R

In this page we briefly discuss the R wrapper, which provides most of Opportunity's functionality. For a general overview, we recommend reading the [User Guide](#).

For installation instructions, please refer to [Installing Opportunity](#). To use the package has to be loaded like all R packages by:

```
library(optunity)
```

## Manual

All functions in R wrapper have documentation available in R's help system, e.g., `?cv.particle_swarm` gives help for cross-validation (CV) using **particle swarms**. To see the list of available functions type `optunity::<TAB><TAB>` in R's command line.

For R following main functions are available:

- `cv.setup` for **setting up CV**, specifying data and the number of folds
- `cv.particle_swarm`, `cv.nelder_mead`, `cv.grid_search`, `cv.random_search` for **optimizing hyperparameters** in CV setting.
- `auc_roc` and `auc_pr` for calculating **AUC** of ROC and precision-recall curve.
- `early_rie` and `early_bedroc` for **early discovery** metrics.
- `mean_se` and `mean_ae` for regression **loss functions**, mean squared error and mean absolute error.
- `particle_swarm`, `nelder_mead`, `grid_search`, `random_search` for minimizing (or maximizing) regular functionals.
- `cv.run` performs a **cross-validation** evaluation for the given values of the hyperparameters (no optimization).

General workflow is to first create a CV object using `cv.setup` and then use it to run CV, by using functions like `cv.particle_swarms`, `cv.random_search` etc.

## Examples

Please see the example pages, [Ridge Regression](#) and [SVM \(e1071\)](#). The R help pages for each function contain examples showing how to use them.

## 5.6 Opportunity API

## opportunity submodules

### Key features

#### 5.6.1 Solvers

The `opportunity.solvers` package bundles all solvers and related logic available in Opportunity.

### Module contents

Module to take care of registering solvers for use in the main Opportunity API.

Main classes in this module:

- `GridSearch`
- `RandomSearch`
- `NelderMead`
- `ParticleSwarm`
- `CMA_ES`
- `TPE`
- `Sobol`

**Warning:** `CMA_ES` require `DEAP` and `NumPy`.

**Warning:** `TPE` require `Hyperopt` and `NumPy`.

Module author: Marc Claesen

### Submodules

This package contains the following submodules:

#### opportunity.solvers.CMAES module

```
class opportunity.solvers.CMAES.CMA_ES (num_generations,      sigma=1.0,      Lambda=None,
                                         **kwargs)
```

Bases: `opportunity.solvers.util.Solver`

Please refer to `CMA-ES` for details about this algorithm.

This solver uses an implementation available in the DEAP library [[DEAP2012](#)].

**Warning:** This solver has dependencies on `DEAP` and `NumPy` and will be unavailable if these are not met.

blah

**Warning:** This solver is not explicitly constrained. The box constraints that are given are used for initialization, but solver may leave the specified region during iterations. If this is unacceptable, you must manually constrain the domain of the objective function prior to using this solver (cfr. [Domain constraints](#)).

`lambda_`

`maximize (f, pmap=<built-in function map>)`

Maximizes `f`.

#### Parameters

- `f (callable)` – the objective function
- `pmap (callable)` – the map() function to use

#### Returns

- the arguments which optimize `f`
- an optional solver report, can be `None`

`minimize (f, pmap=<built-in function map>)`

Minimizes `f`.

#### Parameters

- `f (callable)` – the objective function
- `pmap (callable)` – the map() function to use

#### Returns

- the arguments which optimize `f`

## 5.6.9 Optunity

### Provides

1. Routines to efficiently optimize hyperparameters
2. Function decorators to implicitly log evaluations, constrain the domain and more.
3. Facilities for k-fold cross-validation to estimate generalization performance.

### Available modules

**solvers** contains all officially supported solvers

**functions** a variety of useful function decorators for hyperparameter tuning

**cross\_validation** k-fold cross-validation

### Available subpackages

**tests** regression test suite

**solvers** solver implementations and auxillary functions

### Utilities

**\_\_version\_\_** Optunity version string

**\_\_revision\_\_** Optunity revision string

**\_\_author\_\_** Main authors of the package

**optunity.manual** (*solver\_name=None*)

Prints the manual of requested solver.

**Parameters** **solver\_name** – (optional) name of the solver to request a manual from. If none is specified, a general manual is printed.

Raises `KeyError` if `solver_name` is not registered.

**optunity.maximize** (*f*, *num\_evals=50*, *solver\_name=None*, *pmap=<built-in function map>*, *\*\*kwargs*)

Basic function maximization routine. Maximizes *f* within the given box constraints.

#### Parameters

- **f** – the function to be maximized
- **num\_evals** – number of permitted function evaluations
- **solver\_name** (*string*) – name of the solver to use (optional)
- **pmap** (*callable*) – the map function to use
- **kwargs** – box constraints, a dict of the following form `{'parameter_name': [lower_bound, upper_bound], ...}`

**Returns** retrieved maximum, extra information and solver info

This function will implicitly choose an appropriate solver and its initialization based on `num_evals` and the box constraints.

```
optunity.minimize(f, num_evals=50, solver_name=None, pmap=<built-in function map>, **kwargs)
```

Basic function minimization routine. Minimizes f within the given box constraints.

### Parameters

- **f** – the function to be minimized
- **num\_evals** – number of permitted function evaluations
- **solver\_name** (*string*) – name of the solver to use (optional)
- **pmap** (*callable*) – the map function to use
- **kwargs** – box constraints, a dict of the following form {‘parameter\_name’: [lower\_bound, upper\_bound], ...}

### Returns

retrieved minimum, extra information and solver info

This function will implicitly choose an appropriate solver and its initialization based on num\_evals and the box constraints.

```
optunity.optimize(solver, func, maximize=True, max_evals=0, pmap=<built-in function map>, decoder=None)
```

Optimizes func with given solver.

### Parameters

- **solver** – the solver to be used, for instance a result from *optunity.make\_solver()*
- **func** (*callable*) – the objective function
- **maximize** (*bool*) – maximize or minimize?
- **max\_evals** (*int*) – maximum number of permitted function evaluations
- **pmap** (*function*) – the map() function to use, to vectorize use *optunity.pmap()*

Returns the solution and a namedtuple with further details.

### Result details includes the following:

**optimum** optimal function value f(solution)

**stats** statistics about the solving process

**call\_log** the call log

**report** solver report, can be None

### Statistics gathered while solving a problem:

**num\_evals** number of function evaluations

**time** wall clock time needed to solve

```
optunity.wrap_call_log(f, call_dict)
```

Wraps an existing call log (as dictionary) around f.

This allows you to communicate known function values to solvers. (currently available solvers do not use this info)

```
optunity.wrap_constraints(f, default=None, ub_o=None, ub_c=None, lb_o=None, lb_c=None,
                           range_oo=None, range_co=None, range_oc=None, range_cc=None,
                           custom=None)
```

Decorates f with given input domain constraints.

### Parameters

- **f** (*callable*) – the function that will be constrained

- **default** (*number*) – function value to default to in case of constraint violations
- **ub\_o** (*dict*) – open upper bound constraints, e.g.  $x < c$
- **ub\_c** (*dict*) – closed upper bound constraints, e.g.  $x \leq c$
- **lb\_o** (*dict*) – open lower bound constraints, e.g.  $x > c$
- **lb\_c** (*dict*) – closed lower bound constraints, e.g.  $x \geq c$
- **range\_oo** (*dict with 2-element lists as values ([lb, ub])*) – range constraints (open lb and open ub)  $lb < x < ub$
- **range\_co** (*dict with 2-element lists as values ([lb, ub])*) – range constraints (closed lb and open ub)  $lb \leq x < ub$
- **range\_oc** (*dict with 2-element lists as values ([lb, ub])*) – range constraints (open lb and closed ub)  $lb < x \leq ub$
- **range\_cc** (*dict with 2-element lists as values ([lb, ub])*) – range constraints (closed lb and closed ub)  $lb \leq x \leq ub$
- **custom** (*list of constraints*) – custom, user-defined constraints

\*custom constraints are binary functions that yield False in case of violations.

```
>>> def f(x):
...     return x
>>> fc = wrap_constraints(f, default=-1, range_oc={'x': [0, 1]})
>>> fc(x=0.5)
0.5
>>> fc(x=1)
1
>>> fc(x=5)
-1
>>> fc(x=0)
-1
```

We can define any custom constraint that we want. For instance, assume we have a binary function with arguments  $x$  and  $y$ , and we want to make sure that the provided values remain within the unit circle.

```
>>> def f(x, y):
...     return x + y
>>> circle_constraint = lambda x, y: (x ** 2 + y ** 2) <= 1
>>> fc = wrap_constraints(f, default=1234, custom=[circle_constraint])
>>> fc(0.0, 0.0)
0.0
>>> fc(1.0, 0.0)
1.0
>>> fc(0.5, 0.5)
1.0
>>> fc(1, 0.5)
1234
```

`opportunity.make_solver(solver_name, *args, **kwargs)`

Creates a Solver from given parameters.

### Parameters

- **solver\_name** (*string*) – the solver to instantiate
- **args** – positional arguments to solver constructor.
- **kwargs** – keyword arguments to solver constructor.

Use `opportunity.manual()` to get a list of registered solvers. For constructor arguments per solver, please refer to [Solver overview](#).

Raises `KeyError` if

- `solver_name` is not registered
- `*args` and `**kwargs` are invalid to instantiate the solver.

```
opportunity.suggest_solver(num_evals=50, solver_name=None, **kwargs)
```

```
opportunity.cross_validated(x, num_folds=10, y=None, strata=None, folds=None, num_iter=1, regenerate_folds=False, clusters=None, aggregator=<function mean>)
```

Function decorator to perform cross-validation as configured.

### Parameters

- `x` – data to be used for cross-validation
- `num_folds` – number of cross-validation folds (default 10)
- `y` – (optional) labels to be used for cross-validation. If specified, `len(labels)` must equal `len(x)`
- `strata` – (optional) strata to account for when generating folds. Strata signify instances that must be spread across folds. Not every instance must be in a stratum. Specify strata as a list of lists of instance indices.
- `folds` – (optional) prespecified cross-validation folds to be used (list of lists (iterations) of lists (folds)).
- `num_iter` – (optional) number of iterations to use (default 1)
- `regenerate_folds` – (optional) whether or not to regenerate folds on every evaluation (default false)
- `clusters` – (optional) clusters to account for when generating folds. Clusters signify instances that must be assigned to the same fold. Not every instance must be in a cluster. Specify clusters as a list of lists of instance indices.
- `aggregator` – function to aggregate scores of different folds (default: `mean`)

**Returns** a `cross_validated_callable` with the proper configuration.

This resulting decorator must be used on a function with the following signature (+ potential other arguments):

### Parameters

- `x_train` (`iterable`) – training data
- `y_train` (`iterable`) – training labels (optional)
- `x_test` (`iterable`) – testing data
- `y_test` (`iterable`) – testing labels (optional)

`y_train` and `y_test` must be available of the `y` argument to this function is not `None`.

These 4 keyword arguments will be bound upon decoration. Further arguments will remain free (e.g. hyperparameter names).

```
>>> data = list(range(5))
>>> @cross_validated(x=data, num_folds=5, folds=[[i] for i in range(5)]], aggregator=identity)
... def f(x_train, x_test, a):
...     return x_test[0] + a
>>> f(a=1)
[1, 2, 3, 4, 5]
```

```
>>> f(1)
[1, 2, 3, 4, 5]
>>> f(a=2)
[2, 3, 4, 5, 6]
```

The number of folds must be less than or equal to the size of the data.

```
>>> data = list(range(5))
>>> @cross_validated(x=data, num_folds=6)
... def f(x_train, x_test, a):
...     return x_test[0] + a
AssertionError
```

The number of labels (if specified) must match the number of data instances.

```
>>> data = list(range(5))
>>> labels = list(range(3))
>>> @cross_validated(x=data, y=labels, num_folds=2)
... def f(x_train, x_test, a):
...     return x_test[0] + a
AssertionError
```

`opportunity.generate_folds(num_rows, num_folds=10, strata=None, clusters=None)`

Generates folds for a given number of rows.

#### Parameters

- `num_rows` – number of data instances
- `num_folds` – number of folds to use (default 10)
- `strata` – (optional) list of lists to indicate different sampling strata. Not all rows must be in a stratum. The number of rows per stratum must be larger than or equal to `num_folds`.
- `clusters` – (optional) list of lists indicating clustered instances. Clustered instances must be placed in a single fold to avoid information leaks.

**Returns** a list of folds, each fold is a list of instance indices

```
>>> folds = generate_folds(num_rows=6, num_folds=2, clusters=[[1, 2]], strata=[[3, 4]])
>>> len(folds)
2
>>> i1 = [idx for idx, fold in enumerate(folds) if 1 in fold]
>>> i2 = [idx for idx, fold in enumerate(folds) if 2 in fold]
>>> i1 == i2
True
>>> i3 = [idx for idx, fold in enumerate(folds) if 3 in fold]
>>> i4 = [idx for idx, fold in enumerate(folds) if 4 in fold]
>>> i3 == i4
False
```

**Warning:** Instances in strata are not necessarily spread out over all folds. Some folds may already be full due to clusters. This effect should be negligible.

`opportunity.pmap(f, *args, **kwargs)`  
Parallel map using multiprocessing.

#### Parameters

- `f` – the callable
- `args` – arguments to `f`, as iterables

**Returns** a list containing the results

**Warning:** This function will not work in IPython: <https://github.com/claesem/optimality/issues/8>.

**Warning:** Python's multiprocessing library is incompatible with Jython.

`optimality.available_solvers()`

Returns a list of all available solvers.

These can be used in `optimality.make_solver()`.

`optimality.call_log2dataframe(log)`

Converts a call log into a pandas data frame. This function errors if you don't have pandas available.

**Parameters** `log` – call log to be converted, as returned by e.g. `optimality.minimize`

**Returns** a pandas data frame capturing the same information as the call log

`optimality.maximize_structured(f, search_space, num_evals=50, pmap=<built-in function map>)`

Basic function maximization routine. Maximizes `f` within the given box constraints.

### Parameters

- `f` – the function to be maximized
- `search_space` – the search space (see [Structured search spaces](#) for details)
- `num_evals` – number of permitted function evaluations
- `pmap (callable)` – the map function to use

**Returns** retrieved maximum, extra information and solver info

This function will implicitly choose an appropriate solver and its initialization based on `num_evals` and the box constraints.

`optimality.minimize_structured(f, search_space, num_evals=50, pmap=<built-in function map>)`

Basic function minimization routine. Minimizes `f` within the given box constraints.

### Parameters

- `f` – the function to be minimized
- `search_space` – the search space (see [Structured search spaces](#) for details)
- `num_evals` – number of permitted function evaluations
- `pmap (callable)` – the map function to use

**Returns** retrieved minimum, extra information and solver info

This function will implicitly choose an appropriate solver and its initialization based on `num_evals` and the box constraints.

---

## Bibliography

---

- [DEAP2012] Fortin, Félix-Antoine, et al. “DEAP: Evolutionary algorithms made easy.” *Journal of Machine Learning Research* 13.1 (2012): 2171-2175.
- [RAND] Bergstra, James, and Yoshua Bengio. *Random search for hyper-parameter optimization*. The Journal of Machine Learning Research 13.1 (2012): 281-305
- [PSO2010] Kennedy, James. *Particle swarm optimization*. Encyclopedia of Machine Learning. Springer US, 2010. 760-766.
- [PSO2002] Clerc, Maurice, and James Kennedy. *The particle swarm-explosion, stability, and convergence in a multi-dimensional complex space*. Evolutionary Computation, IEEE Transactions on 6.1 (2002): 58-73.
- [NELDERMEAD] Nelder, John A. and Mead, R. *A simplex method for function minimization*. Computer Journal 7: 308–313, 1965.
- [HANSEN2001] Nikolaus Hansen and Andreas Ostermeier. *Completely derandomized self-adaptation in evolution strategies*. Evolutionary computation, 9(2):159-195, 2001.
- [DEAP2012] Felix-Antoine Fortin, Francois-Michel De Rainville, Marc-Andre Gardner, Marc Parizeau and Christian Gagne, *DEAP: Evolutionary Algorithms Made Easy*, *Journal of Machine Learning Research*, pp. 2171-2175, no 13, jul 2012.
- [TPE2011] Bergstra, James S., et al. “Algorithms for hyper-parameter optimization.” Advances in Neural Information Processing Systems. 2011.
- [TPE2013] Bergstra, James, Daniel Yamins, and David Cox. “Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures.” Proceedings of The 30th International Conference on Machine Learning. 2013.
- [Hyperopt] <http://jaberg.github.io/hyperopt/>
- [SOBOL] Ilya Sobol, USSR Computational Mathematics and Mathematical Physics, Volume 16, pages 236-242, 1977.
- [SOBOL2] Ilya Sobol, Levitan, The Production of Points Uniformly Distributed in a Multidimensional Cube (in Russian), Preprint IPM Akad. Nauk SSSR, Number 40, Moscow 1976.
- [ANTONOV] Antonov, Saleev, USSR Computational Mathematics and Mathematical Physics, Volume 19, 1980, pages 252 - 256.
- [BRATLEY] Paul Bratley, Bennett Fox, Algorithm 659: Implementing Sobol’s Quasirandom Sequence Generator, ACM Transactions on Mathematical Software, Volume 14, Number 1, pages 88-100, 1988.

- [FOX] Bennett Fox, Algorithm 647: Implementation and Relative Efficiency of Quasirandom Sequence Generators, ACM Transactions on Mathematical Software, Volume 12, Number 4, pages 362-376, 1986.
- [DEAP2012] Fortin, Félix-Antoine, et al. “DEAP: Evolutionary algorithms made easy.” Journal of Machine Learning Research 13.1 (2012): 2171-2175.
- [RAND] Bergstra, James, and Yoshua Bengio. *Random search for hyper-parameter optimization*. The Journal of Machine Learning Research 13.1 (2012): 281-305
- [PSO2010] Kennedy, James. *Particle swarm optimization*. Encyclopedia of Machine Learning. Springer US, 2010. 760-766.
- [PSO2002] Clerc, Maurice, and James Kennedy. *The particle swarm-explosion, stability, and convergence in a multi-dimensional complex space*. Evolutionary Computation, IEEE Transactions on 6.1 (2002): 58-73.
- [NELDERMEAD] Nelder, John A. and Mead, R. *A simplex method for function minimization*. Computer Journal 7: 308–313, 1965.
- [HANSEN2001] Nikolaus Hansen and Andreas Ostermeier. *Completely derandomized self-adaptation in evolution strategies*. Evolutionary computation, 9(2):159-195, 2001.
- [DEAP2012] Felix-Antoine Fortin, Francois-Michel De Rainville, Marc-Andre Gardner, Marc Parizeau and Christian Gagne, *DEAP: Evolutionary Algorithms Made Easy*, Journal of Machine Learning Research, pp. 2171-2175, no 13, jul 2012.
- [TPE2011] Bergstra, James S., et al. “Algorithms for hyper-parameter optimization.” Advances in Neural Information Processing Systems. 2011.
- [TPE2013] Bergstra, James, Daniel Yamins, and David Cox. “Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures.” Proceedings of The 30th International Conference on Machine Learning. 2013.
- [Hyperopt] <http://jaberg.github.io/hyperopt/>
- [SOBOL] Ilya Sobol, USSR Computational Mathematics and Mathematical Physics, Volume 16, pages 236-242, 1977.
- [SOBOL2] Ilya Sobol, Levitan, The Production of Points Uniformly Distributed in a Multidimensional Cube (in Russian), Preprint IPM Akad. Nauk SSSR, Number 40, Moscow 1976.
- [ANTONOV] Antonov, Saleev, USSR Computational Mathematics and Mathematical Physics, Volume 19, 1980, pages 252 - 256.
- [BRATLEY] Paul Bratley, Bennett Fox, Algorithm 659: Implementing Sobol’s Quasirandom Sequence Generator, ACM Transactions on Mathematical Software, Volume 14, Number 1, pages 88-100, 1988.
- [FOX] Bennett Fox, Algorithm 647: Implementation and Relative Efficiency of Quasirandom Sequence Generators, ACM Transactions on Mathematical Software, Volume 12, Number 4, pages 362-376, 1986.
- [TPE2011] Bergstra, James S., et al. “Algorithms for hyper-parameter optimization.” Advances in Neural Information Processing Systems. 2011
- [LEE2003] Wee Sun Lee and Bing Liu. Learning with positive and unlabeled examples using weighted logistic regression. In Proceedings of the Twentieth International Conference on Machine Learning (ICML), 2003.

·  
optunity.communication, 122  
optunity.constraints, 122  
optunity.cross\_validation, 122  
optunity.functions, 122  
optunity.metrics, 122  
optunity.search\_spaces, 122  
optunity.solvers, 122  
optunity.solvers.CMAES, 122  
optunity.solvers.GridSearch, 122  
optunity.solvers.NelderMead, 122  
optunity.solvers.ParticleSwarm, 122  
optunity.solvers.RandomSearch, 122  
optunity.solvers.Sobol, 122  
optunity.solvers.solver\_registry, 122  
optunity.solvers.TPE, 122  
optunity.solvers.util, 122  
optunity.standalone, 122

## 0

optunity, 123



**A**

absolute\_error() (in module `optunity.metrics`), 122  
accept\_server\_connection() (in module `optunity.communication`), 122  
accuracy() (in module `optunity.metrics`), 122  
add\_to\_queue() (`optunity.communication.EvalManager` method), 122  
append() (`optunity.solvers.util.ThreadSafeQueue` method), 122  
`Args` (class in `optunity.functions`), 122  
`args` (`optunity.constraints.ConstraintViolation` attribute), 122  
`args` (`optunity.functions.MaximumEvaluationsException` attribute), 122  
`assign_grid_points()` (`optunity.solvers.GridSearch.GridSearch` method), 122  
`auc()` (in module `optunity.metrics`), 122  
`available_solvers()` (in module `optunity`), 128

**B**

`bitwise_xor()` (`optunity.solvers.Sobol.Sobol` static method), 122  
`bounds` (`optunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
`bounds` (`optunity.solvers.RandomSearch.RandomSearch` attribute), 122  
`bounds` (`optunity.solvers.Sobol.Sobol` attribute), 122  
`bounds` (`optunity.solvers.TPE.TPE` attribute), 122  
box constraints, 26  
`brier()` (in module `optunity.metrics`), 122

**C**

`call_log2dataframe()` (in module `optunity`), 128  
`call_log2dataframe()` (in module `optunity.functions`), 122  
`CallLog` (class in `optunity.functions`), 122  
`cases` (`optunity.search_spaces.Options` attribute), 122  
`choice` (`optunity.search_spaces.Node` attribute), 122  
`clone()` (`optunity.solvers.ParticleSwarm.ParticleSwarm`.`Particle` method), 122

`CMA_ES` (class in `optunity.solvers.CMAES`), 122  
`compute_curve()` (in module `optunity.metrics`), 122  
`constr_lb_c()` (in module `optunity.constraints`), 122  
`constr_lb_o()` (in module `optunity.constraints`), 122  
`constr_range_cc()` (in module `optunity.constraints`), 122  
`constr_range_co()` (in module `optunity.constraints`), 122  
`constr_range_oc()` (in module `optunity.constraints`), 122  
`constr_range_oo()` (in module `optunity.constraints`), 122  
`constr_ub_c()` (in module `optunity.constraints`), 122  
`constr_ub_o()` (in module `optunity.constraints`), 122  
`constrained()` (in module `optunity.constraints`), 122  
constraint (`optunity.constraints.ConstraintViolation` attribute), 122  
`ConstraintViolation`, 122  
`content` (`optunity.search_spaces.SearchTree` attribute), 122  
`content` (`optunity.solvers.util.ThreadSafeQueue` attribute), 122  
`contingency_table()` (in module `optunity.metrics`), 122  
`contingency_tables()` (in module `optunity.metrics`), 122  
`copy()` (`optunity.solvers.util.ThreadSafeQueue` method), 122  
`cross_validated()` (in module `optunity`), 126  
`cross_validated()` (in module `optunity.cross_validation`), 122  
`cv` (`optunity.communication.EvalManager` attribute), 122

**D**

`data` (`optunity.functions.CallLog` attribute), 122  
`decode()` (`optunity.search_spaces.SearchTree` method), 122  
`delete()` (`optunity.functions.CallLog` method), 122  
`desc_brief` (`optunity.solvers.GridSearch.GridSearch` attribute), 122  
`desc_brief` (`optunity.solvers.NelderMead.NelderMead` attribute), 122  
`desc_brief` (`optunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
`desc_brief` (`optunity.solvers.RandomSearch.RandomSearch` attribute), 122  
`desc_brief` (`optunity.solvers.Sobol.Sobol` attribute), 122

desc\_full (opportunity.solvers.GridSearch.GridSearch attribute), 122  
desc\_full (opportunity.solvers.NelderMead.NelderMead attribute), 122  
desc\_full (opportunity.solvers.ParticleSwarm.ParticleSwarm attribute), 122  
desc\_full (opportunity.solvers.RandomSearch.RandomSearch attribute), 122  
desc\_full (opportunity.solvers.Sobol.Sobol attribute), 122

**E**

error\_rate() (in module opportunity.metrics), 122  
EvalManager (class in opportunity.communication), 122

**F**

fbeta() (in module opportunity.metrics), 122  
flush\_queue() (opportunity.communication.EvalManager method), 122  
fold\_request() (in module opportunity.standalone), 122  
from\_dict() (opportunity.functions.CallLog static method), 122  
ftol (opportunity.solvers.NelderMead.NelderMead attribute), 122

**G**

generate() (opportunity.solvers.ParticleSwarm.ParticleSwarm method), 122  
generate\_folds() (in module opportunity), 127  
generate\_folds() (in module opportunity.cross\_validation), 122  
get() (in module opportunity.solvers.solver\_registry), 122  
get() (opportunity.communication.EvalManager method), 122  
get() (opportunity.functions.CallLog method), 122  
GridSearch (class in opportunity.solvers.GridSearch), 122

**H**

hyperparameters, 25

**I**

i4\_bit\_hi1() (opportunity.solvers.Sobol.Sobol static method), 122  
i4\_bit\_lo0() (opportunity.solvers.Sobol.Sobol static method), 122  
i4\_sobol() (opportunity.solvers.Sobol.Sobol static method), 122  
i4\_sobol\_generate() (opportunity.solvers.Sobol.Sobol static method), 122  
identity() (in module opportunity.cross\_validation), 122  
insert() (opportunity.functions.CallLog method), 122  
items() (opportunity.functions.CallLog method), 122

**J**

json\_decode() (in module opportunity.communication), 122

json\_encode() (in module opportunity.communication), 122

**K**

key (opportunity.search\_spaces.Node attribute), 122  
keys() (opportunity.functions.Args method), 122  
keys() (opportunity.functions.CallLog method), 122  
kwargs (opportunity.constraints.ConstraintViolation attribute), 122

**L**

lambda\_ (opportunity.solvers.CMAES.CMA\_ES attribute), 122  
list\_mean() (in module opportunity.cross\_validation), 122  
lock (opportunity.functions.CallLog attribute), 122  
lock (opportunity.solvers.util.ThreadSafeQueue attribute), 122  
logged() (in module opportunity.functions), 122  
logloss() (in module opportunity.metrics), 122  
lower (opportunity.solvers.RandomSearch.RandomSearch attribute), 122

**M**

main() (in module opportunity.standalone), 122  
make\_piped\_function() (in module opportunity.communication), 122  
make\_solver() (in module opportunity), 125  
make\_solver() (in module opportunity.standalone), 122  
manual() (in module opportunity), 123  
manual() (in module opportunity.solvers.solver\_registry), 122  
manual\_request() (in module opportunity.standalone), 122  
max\_evals (opportunity.functions.MaximumEvaluationsException attribute), 122  
max\_evals() (in module opportunity.functions), 122  
max\_iter (opportunity.solvers.NelderMead.NelderMead attribute), 122  
max\_or\_min() (in module opportunity.standalone), 122  
max\_speed (opportunity.solvers.ParticleSwarm.ParticleSwarm attribute), 122  
max\_vectorized (opportunity.communication.EvalManager attribute), 122  
maxdim() (opportunity.solvers.Sobol.Sobol static method), 122  
maximize() (in module opportunity), 123  
maximize() (opportunity.solvers.CMAES.CMA\_ES method), 122  
maximize() (opportunity.solvers.GridSearch.GridSearch method), 122  
maximize() (opportunity.solvers.NelderMead.NelderMead method), 122  
maximize() (opportunity.solvers.ParticleSwarm.ParticleSwarm method), 122  
maximize() (opportunity.solvers.RandomSearch.RandomSearch method), 122

maximize() (opportunity.solvers.Sobol.Sobol method), 122  
 maximize() (opportunity.solvers.TPE.TPE method), 122  
 maximize() (opportunity.solvers.util.Solver method), 122  
 maximize\_structured() (in module opportunity), 128  
 MaximumEvaluationsException, 122  
 mean() (in module opportunity.cross\_validation), 122  
 mean\_and\_list() (in module opportunity.cross\_validation), 122  
 message (opportunity.constraints.ConstraintViolation attribute), 122  
 message (opportunity.functions.MaximumEvaluationsException attribute), 122  
 minimize() (in module opportunity), 123  
 minimize() (opportunity.solvers.CMAES.CMA\_ES method), 122  
 minimize() (opportunity.solvers.GridSearch.GridSearch method), 122  
 minimize() (opportunity.solvers.NelderMead.NelderMead method), 122  
 minimize() (opportunity.solvers.ParticleSwarm.ParticleSwarm method), 122  
 minimize() (opportunity.solvers.RandomSearch.RandomSearch method), 122  
 minimize() (opportunity.solvers.Sobol.Sobol method), 122  
 minimize() (opportunity.solvers.TPE.TPE method), 122  
 minimize() (opportunity.solvers.util.Solver method), 122  
 minimize\_structured() (in module opportunity), 128  
 mse() (in module opportunity.metrics), 122

**N**

name (opportunity.solvers.GridSearch.GridSearch attribute), 122  
 name (opportunity.solvers.NelderMead.NelderMead attribute), 122  
 name (opportunity.solvers.ParticleSwarm.ParticleSwarm attribute), 122  
 name (opportunity.solvers.RandomSearch.RandomSearch attribute), 122  
 name (opportunity.solvers.Sobol.Sobol attribute), 122  
 negated() (in module opportunity.functions), 122  
 NelderMead (class in opportunity.solvers.NelderMead), 122  
 Node (class in opportunity.search\_spaces), 122  
 npv() (in module opportunity.metrics), 122  
 nth\_root() (in module opportunity.solvers.GridSearch), 122  
 num\_evals (opportunity.solvers.RandomSearch.RandomSearch attribute), 122  
 num\_evals (opportunity.solvers.Sobol.Sobol attribute), 122  
 num\_evals (opportunity.solvers.TPE.TPE attribute), 122  
 num\_generations (opportunity.solvers.CMAES.CMA\_ES attribute), 122  
 num\_generations (opportunity.solvers.ParticleSwarm.ParticleSwarm attribute), 122

**O**

objective function, 25  
 open\_server\_socket() (in module opportunity.communication), 122  
 open\_socket() (in module opportunity.communication), 122  
 optimize() (in module opportunity), 124  
 optimize() (in module opportunity.standalone), 122  
 optimize() (opportunity.solvers.CMAES.CMA\_ES method), 122  
 optimize() (opportunity.solvers.GridSearch.GridSearch method), 122  
 optimize() (opportunity.solvers.NelderMead.NelderMead method), 122  
 optimize() (opportunity.solvers.ParticleSwarm.ParticleSwarm method), 122  
 optimize() (opportunity.solvers.RandomSearch.RandomSearch method), 122  
 optimize() (opportunity.solvers.Sobol.Sobol method), 122  
 optimize() (opportunity.solvers.TPE.TPE method), 122  
 optimize() (opportunity.solvers.util.Solver method), 122  
 Options (class in opportunity.search\_spaces), 122  
 opportunity (module), 123  
 opportunity.communication (module), 122  
 opportunity.constraints (module), 122  
 opportunity.cross\_validation (module), 122  
 opportunity.functions (module), 122  
 opportunity.metrics (module), 122  
 opportunity.search\_spaces (module), 122  
 opportunity.solvers (module), 122  
 opportunity.solvers.CMAES (module), 122  
 opportunity.solvers.GridSearch (module), 122  
 opportunity.solvers.NelderMead (module), 122  
 opportunity.solvers.ParticleSwarm (module), 122  
 opportunity.solvers.RandomSearch (module), 122  
 opportunity.solvers.Sobol (module), 122  
 opportunity.solvers.solver\_registry (module), 122  
 opportunity.solvers.TPE (module), 122  
 opportunity.solvers.util (module), 122  
 opportunity.standalone (module), 122

**P**

parameter\_tuples (opportunity.solvers.GridSearch.GridSearch attribute), 122  
 parameters (opportunity.functions.Args attribute), 122  
 particle2dict() (opportunity.solvers.ParticleSwarm.ParticleSwarm method), 122  
 ParticleSwarm (class in opportunity.solvers.ParticleSwarm), 122  
 ParticleSwarm.Particle (class in opportunity.solvers.ParticleSwarm), 122

**phi1** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
**phi2** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
**pipe\_eval()** (`opportunity.communication.EvalManager` method), 122  
**pmap()** (in module `opportunity`), 127  
**pmap()** (`opportunity.communication.EvalManager` method), 122  
**pr\_auc()** (in module `opportunity.metrics`), 122  
**precision()** (in module `opportunity.metrics`), 122  
**prepare\_fun()** (in module `opportunity.standalone`), 122  
**processed\_semaphore** (`opportunity.communication.EvalManager` attribute), 122  
**pu\_score()** (in module `opportunity.metrics`), 122

**Q**

**queue** (`opportunity.communication.EvalManager` attribute), 122  
**queue\_lock** (`opportunity.communication.EvalManager` attribute), 122

**R**

**r\_squared()** (in module `opportunity.metrics`), 122  
**random\_permutation()** (in module `opportunity.cross_validation`), 122  
**RandomSearch** (class in `opportunity.solvers.RandomSearch`), 122  
**recall()** (in module `opportunity.metrics`), 122  
**receive()** (in module `opportunity.communication`), 122  
**reflect()** (`opportunity.solvers.NelderMead.NelderMead` static method), 122  
**register\_solver()** (in module `opportunity.solvers.solver_registry`), 122  
**replacements** (`opportunity.communication.EvalManager` attribute), 122  
**roc\_auc()** (in module `opportunity.metrics`), 122

**S**

**scale()** (`opportunity.solvers.NelderMead.NelderMead` static method), 122  
**scale\_unit\_to\_bounds()** (in module `opportunity.solvers.util`), 122  
**score**, 25  
**score()** (in module `opportunity.solvers.util`), 122  
**SearchTree** (class in `opportunity.search_spaces`), 122  
**seed** (`opportunity.solvers.TPE.TPE` attribute), 122  
**select()** (in module `opportunity.cross_validation`), 122  
**semaphore** (`opportunity.communication.EvalManager` attribute), 122  
**send()** (in module `opportunity.communication`), 122  
**shrink\_bounds()** (in module `opportunity.solvers.util`), 122  
**sigma** (`opportunity.solvers.CMAES.CMA_ES` attribute), 122

**simplex\_center()** (`opportunity.solvers.NelderMead.NelderMead` static method), 122  
**skip** (`opportunity.solvers.Sobol.Sobol` attribute), 122  
**smax** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
**smin** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
**Sobol** (class in `opportunity.solvers.Sobol`), 122  
**sobolseed** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` attribute), 122  
**solver**, 26  
**Solver** (class in `opportunity.solvers.util`), 122  
**solver\_names()** (in module `opportunity.solvers.solver_registry`), 122  
**sort\_vertices()** (`opportunity.solvers.NelderMead.NelderMead` static method), 122  
**start** (`opportunity.solvers.CMAES.CMA_ES` attribute), 122  
**start** (`opportunity.solvers.NelderMead.NelderMead` attribute), 122  
**static\_key\_order()** (in module `opportunity.functions`), 122  
**strata\_by\_labels()** (in module `opportunity.cross_validation`), 122  
**suggest\_from\_box()** (`opportunity.solvers.GridSearch.GridSearch` static method), 122  
**suggest\_from\_box()** (`opportunity.solvers.ParticleSwarm.ParticleSwarm` static method), 122  
**suggest\_from\_box()** (`opportunity.solvers.RandomSearch.RandomSearch` static method), 122  
**suggest\_from\_box()** (`opportunity.solvers.Sobol.Sobol` static method), 122  
**suggest\_from\_box()** (`opportunity.solvers.TPE.TPE` static method), 122  
**suggest\_from\_seed()** (`opportunity.solvers.CMAES.CMA_ES` static method), 122  
**suggest\_from\_seed()** (`opportunity.solvers.NelderMead.NelderMead` static method), 122  
**suggest\_solver()** (in module `opportunity`), 126

**T**

**terminal** (`opportunity.search_spaces.Node` attribute), 122  
**ThreadSafeQueue** (class in `opportunity.solvers.util`), 122  
**to\_box()** (`opportunity.search_spaces.SearchTree` method), 122  
**to\_dict()** (`opportunity.functions.CallLog` method), 122  
**TPE** (class in `opportunity.solvers.TPE`), 122  
**train-predict-score (TPS)** chain, 26

## U

uniform\_in\_bounds() (in module optunity.solvers.util),  
[122](#)  
update() (optunity.functions.CallLog method), [122](#)  
updateParticle() (optunity.solvers.ParticleSwarm.ParticleSwarm  
method), [122](#)  
upper (optunity.solvers.RandomSearch.RandomSearch  
attribute), [122](#)

## V

value (optunity.search\_spaces.Node attribute), [122](#)  
values() (optunity.functions.Args method), [122](#)  
values() (optunity.functions.CallLog method), [122](#)  
vectorcontent (optunity.search\_spaces.SearchTree attribute), [122](#)  
vectordict (optunity.search\_spaces.SearchTree attribute),  
[122](#)  
vectorized (optunity.communication.EvalManager  
attribute), [122](#)  
violations\_defaulted() (in module optunity.constraints),  
[122](#)

## W

wrap\_call\_log() (in module optunity), [124](#)  
wrap\_constraints() (in module optunity), [124](#)  
wrap\_constraints() (in module optunity.constraints), [122](#)  
wrap\_decoder() (optunity.search\_spaces.SearchTree  
method), [122](#)  
wraps() (in module optunity.functions), [122](#)