
M-LOOP Documentation

Release 3.3.3

Michael R Hush

May 12, 2023

Contents

1	Quick Start	3
2	Contents	5
2.1	Installation	5
2.2	Tutorials	8
2.3	Interfaces	20
2.4	Data	22
2.5	Visualizations	23
2.6	Examples	27
2.7	Contributing	34
2.8	Changelog	35
2.9	M-LOOP API	35
3	Indices	83
	Python Module Index	85
	Index	87

The Machine-Learning Online Optimization Package is designed to automatically and rapidly optimize the parameters of a scientific experiment or computer controller system.

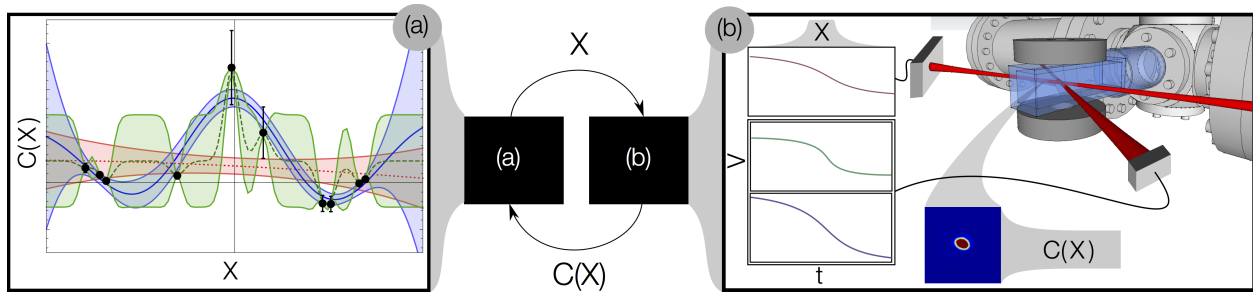


Fig. 1: M-LOOP in control of an ultra-cold atom experiment. M-LOOP was able to find an optimal set of ramps to evaporatively cool a thermal gas and form a Bose-Einstein Condensate.

Using M-LOOP is simple, once the parameters of your experiment are computer controlled, all you need to do is determine a cost function that quantifies the performance of an experiment after a single run. You can then hand over control of the experiment to M-LOOP which will attempt to find a global optimal set of parameters that minimize the cost function, by performing experiments and testing different parameters. M-LOOP uses machine learning to predict how the parameters of the experiment relate to the cost; it uses this model to pick the next best parameters to test to find an optimum as quickly as possible. This approach is known as Bayesian optimization.

M-LOOP not only finds an optimal set of parameters for the experiment it also provides a model of how the parameters are related to the costs which can be used to improve the experiment.

If you use M-LOOP please cite our publication where we first used the package to optimize the production of a Bose-Einstein Condensate:

Fast Machine-Learning Online Optimization of Ultra-Cold-Atom Experiments. *Scientific Reports* **6**, 25890 (2016). DOI: [Link 10.1038/srep25890](https://doi.org/10.1038/srep25890)

<http://www.nature.com/articles/srep25890>

CHAPTER 1

Quick Start

To get M-LOOP running follow the *Installation* instructions and *Tutorials*.

2.1 Installation

M-LOOP is available on PyPI and can be installed with your favorite package manager; simply search for ‘M-LOOP’ and install. However, if you want the latest features and a local copy of the examples you should install M-LOOP using the source code from the [GitHub](#). Detailed installation instructions are provided below.

The installation process involves three steps.

1. Get a Python distribution with the standard scientific packages. We recommend installing *Anaconda*.
2. Install the latest release of *M-LOOP*.
3. (Optional) *Test* your M-LOOP install.

If you are having any trouble with the installation you may need to check your *package dependencies* have been correctly installed. If you are still having trouble, you can [submit an issue](#) to the GitHub.

2.1.1 Anaconda

We recommend installing Anaconda to get a python environment with all the required scientific packages. The Anaconda distribution is available here:

<https://www.anaconda.com/>

Follow the installation instructions they provide.

M-LOOP is targeted at python 3 but also supports 2. Please use python 3 if you do not have a reason to use 2, see *Python 3 vs 2* for details.

2.1.2 M-LOOP

You have two options when installing M-LOOP, you can perform a basic installation of the last release with pip or you can install from source to get the latest features. We recommend installing from source so you can test your installation, see all the examples and get the most recent bug fixes.

Installing from source

Note: If using Anaconda Python, it may be necessary to install Tensorflow using `conda` *before* installing M-LOOP with Pip, otherwise Tensorflow may not install correctly. This can be done with the command `conda install tensorflow`.

M-LOOP can be installed from the latest source code with three commands:

```
git clone git://github.com/michaelhush/M-LOOP.git
cd ./M-LOOP
pip install -e .
```

The first command downloads the latest source code for M-LOOP from GitHub into the current directory, the second moves into the M-LOOP source directory, and the third command builds the package and creates a link from you python package to the source. If you are using linux or MacOS you may need admin privileges to run the installation step.

At any time you can update M-LOOP to the latest version from GitHub by running the command:

```
git pull origin master
```

in the M-LOOP directory.

Installing with pip

Note: If using Anaconda Python, it may be necessary to install Tensorflow using `conda` *before* installing M-LOOP with Pip, otherwise Tensorflow may not install correctly. This can be done with the command `conda install tensorflow`.

M-LOOP can be installed with pip with a single command:

```
pip install M-LOOP
```

If you are using linux or MacOS you may need admin privileges to run the command. To update M-LOOP to the latest version use:

```
pip install M-LOOP --upgrade
```

2.1.3 Testing

If you have installed from source, you can test your installation by running the command:

```
pytest
```

In the M-LOOP source code directory. The tests should take around five minutes to complete. If you find a error please consider [Contributing](#) to the project and report a bug on the [GitHub](#).

If you installed M-LOOP using pip, you will not need to test your installation.

2.1.4 Dependencies

M-LOOP requires the following packages to run correctly.

Package	Version
docutils	>=0.3
matplotlib	>=1.5
numpy	>=1.11
pip	>=7.0
pytest	>=2.9
setuptools	>=26
scikit-learn	>=0.18
scipy	>=0.17
tensorflow	>=2.0.0

These packages should be automatically installed by pip or the script setup.py when you install M-LOOP. The setup script itself requires pytest-runner.

However, if you are using Anaconda some packages that are managed by the conda command may not be correctly updated, even if your installation passes all the tests. In this case, you will have to update these packages manually. You can check what packages you have installed and their version with the command:

```
conda list
```

To install a package that is missing, say for example pytest, use the command:

```
conda install pytest
```

To update a package to the latest version, say for example scikit-learn, use the command:

```
conda update scikit-learn
```

Once you install and update all the required packages with conda M-LOOP should run correctly.

2.1.5 Documentation

The latest documentation will always be available here online. If you would also like a local copy of the documentation, and you have downloaded the source code, enter the docs folder and use the command:

```
make html
```

Which will generate the documentation in docs/_build/html.

2.1.6 Python 3 vs 2

M-LOOP is developed in python 3 and it gets the best performance in this environment. This is primarily because other packages that M-LOOP uses, like numpy, run fastest in python 3. The tests typically take about 20% longer to complete in python 2 than 3.

If you have a specific reason to stay in a python 2 environment (you may use other packages which are not python 3 compatible) then you can still use M-LOOP without upgrading to 3. However, if you do not have a specific reason to stay with python 2, it is highly recommended you use the latest python 3 package.

2.2 Tutorials

Here we provide some tutorials on how to use M-LOOP. M-LOOP is flexible and can be customized with a variety of *options* and *interfaces*. Here we provide some basic tutorials to get you started as quickly as possible.

There are two different approaches to using M-LOOP:

1. You can execute M-LOOP from a command line (or shell) and configure it using a text file.
2. You can use M-LOOP as a *python API*.

If you have a standard experiment that is operated by LabVIEW, Simulink or some other method, then you should use option 1 and follow the *first tutorial*. If your experiment is operated using python, you should consider using option 2 as it will give you more flexibility and control, in which case, look at the *second tutorial*.

2.2.1 Standard experiment

The basic operation of M-LOOP is sketched below.

There are three stages:

1. M-LOOP is started with the command:

```
M-LOOP
```

M-LOOP first looks for the configuration file *exp_config.txt*, which contains options like the number of parameters and their limits, in the folder in which it is executed. Then it starts the optimization process.

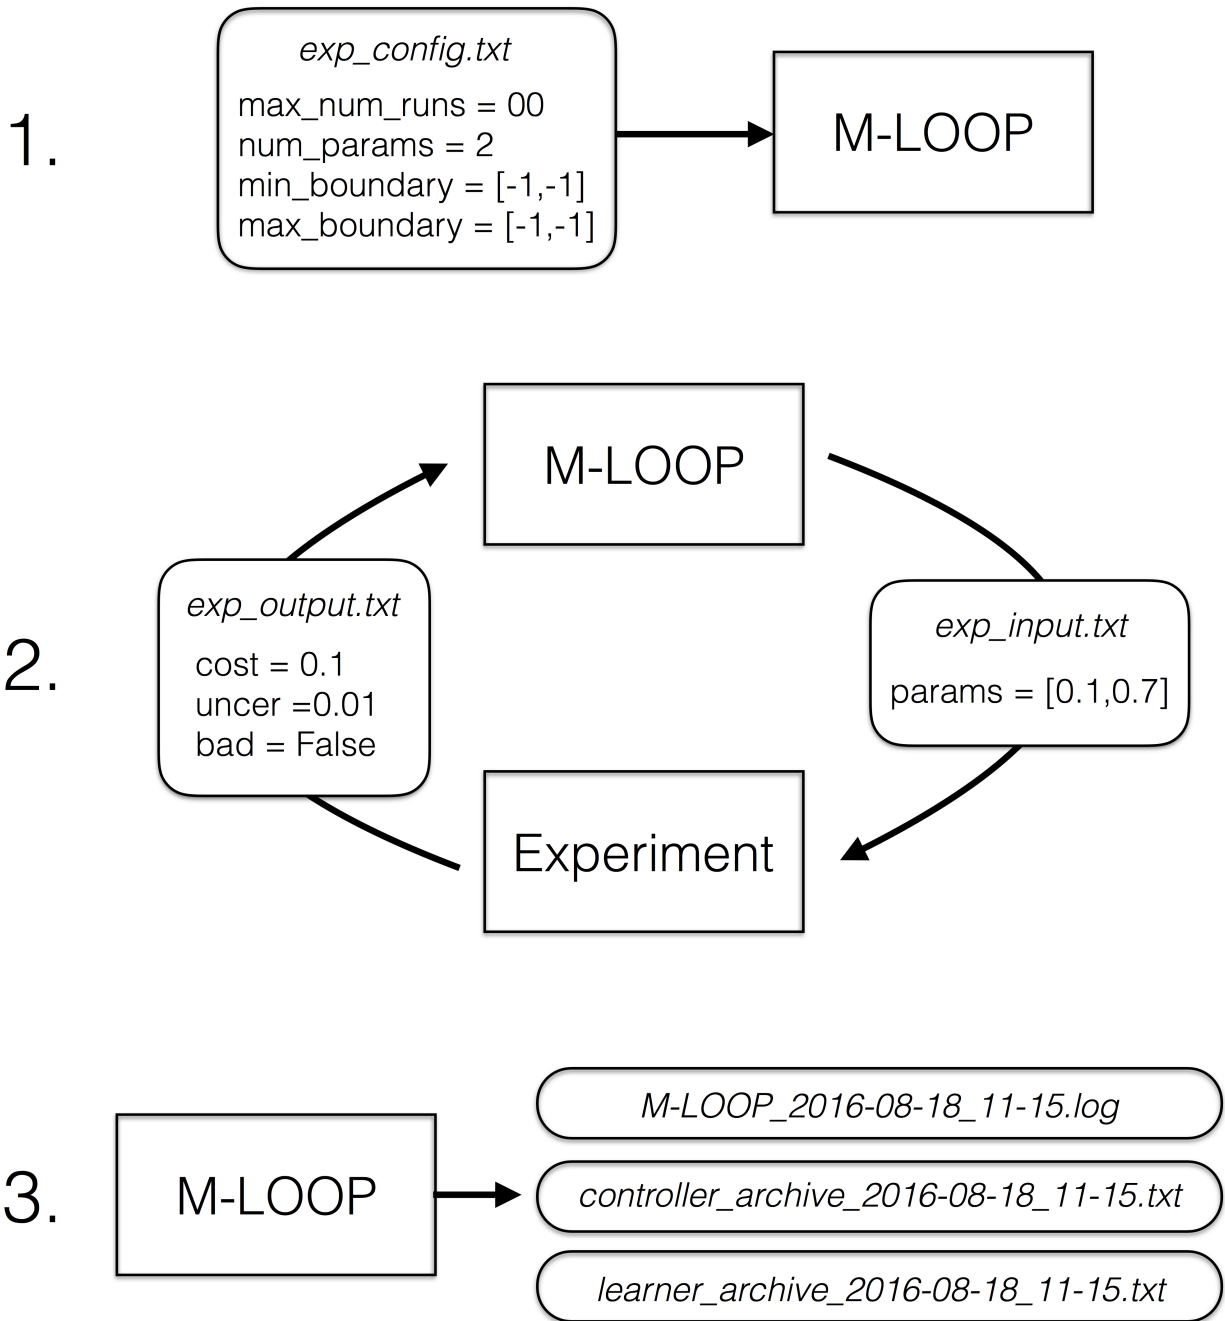
2. M-LOOP controls and optimizes the experiment by exchanging files written to disk. M-LOOP produces a file called *exp_input.txt* which contains a variable params with the next parameters to be run by the experiment. The experiment is expected to run an experiment with these parameters and measure the resultant cost. The experiment should then write the file *exp_output.txt* which contains at least the variable cost which quantifies the performance of that experimental run, and optionally, the variables uncer (for uncertainty) and bad (if the run failed). This process is repeated many times until a halting condition is met.
3. Once the optimization process is complete, M-LOOP prints to the console the parameters and cost of the best run performed during the experiment, and a prediction of what the optimal parameters are (with the corresponding predicted cost and uncertainty). M-LOOP also produces a set of plots that allow the user to visualize the optimization process and cost landscape. During operation and at the end M-LOOP writes these files to disk:
 - *M-LOOP_[datetime].log* a log of the console output and other debugging information during the run.
 - *controller_archive_[datetime].txt* an archive of all the experimental data recorded and the results.
 - *learner_archive_[datetime].txt* an archive of the model created by the machine learner of the experiment.
 - If using the neural net learner, then several *neural_net_archive* files will be saved which store the fitted neural nets.

In what follows we will unpack this process and give details on how to configure and run M-LOOP.

Launching M-LOOP

Launching M-LOOP is performed by executing the command M-LOOP on the console. You can also provide the name of your configuration file if you do not want to use the default with the command:

```
M-LOOP -c [config_filename]
```



Configuration File

The configuration file contains a list of options and settings for the optimization run. Each option must be started on a new line and formatted as:

```
[keyword] = [value]
```

You can add comments to your file using `#`. Everything past the `#` will be ignored. Examples of relevant keywords and syntax for the values are provided in [Examples](#) and a comprehensive list of options are described in [Examples](#). The values should be formatted with python syntax. Strings should be surrounded with single or double quotes and arrays of values can be surrounded with square brackets/parentheses with numbers separated by commas. In this tutorial we will examine the example file `tutorial_config.txt`

```
#Tutorial Config
#-----

#Interface settings
interface_type = 'file'

#Parameter settings
num_params = 2                                #number of parameters
min_boundary = [-1, -1]                       #minimum boundary
max_boundary = [1, 1]                         #maximum boundary
first_params = [0.5, 0.5]                    #first parameters to try
trust_region = 0.4                           #maximum % move distance from best params

#Halting conditions
max_num_runs = 1000                          #maximum number of runs
max_num_runs_without_better_params = 50      #maximum number of runs without finding_
↳better parameters
target_cost = 0.01                           #optimization halts when a cost below this_
↳target is found
max_duration = 36000                          #the optimization will not start a new_
↳iteration after max_duration seconds.

#Learner options
cost_has_noise = True                        #whether the costs are corrupted by noise_
↳or not

#Timing options
no_delay = True                              #wait for learner to make generate new_
↳parameters or use training algorithms

#File format options
interface_file_type = 'txt'                  #file types of *exp_input.mat* and *exp_
↳output.mat*
controller_archive_file_type = 'mat'         #file type of the controller archive
learner_archive_file_type = 'pkl'           #file type of the learner archive

#Visualizations
visualizations = True
```

We will now explain the options in each of their groups. In almost all cases you will only need to adjust the parameters settings and halting conditions, but we have also described a few of the most commonly used extra options.

Parameter settings

The number of parameters and constraints on what parameters can be tried are defined with a few keywords:

```
num_params = 2                #number of parameters
min_boundary = [-1, -1]       #minimum boundary
max_boundary = [1, 1]         #maximum boundary
first_params = [0.5, 0.5]     #first parameters to try
trust_region = 0.4            #maximum % move distance from best params
```

`num_params` defines the number of parameters, `min_boundary` defines the minimum value each of the parameters can take and `max_boundary` defines the maximum value each parameter can take. Here there are two value which each must be between -1 and 1.

`first_params` defines the first parameters the learner will try. You only need to set this if you have a safe set of parameters you want the experiment to start with. Just delete this keyword if any set of parameters in the boundaries will work.

`trust_region` defines the maximum change allowed in the parameters from the best parameters found so far. In the current example the region size is 2 by 2, with a trust region of 40% . Thus the maximum allowed change for the second run will be $[0 \pm 0.8, 0 \pm 0.8]$. Alternatively you can provide a list of values for `trust_region`, which should have one entry for each parameter. In that case each entry specifies the maximum change for the corresponding parameter. When specified as a list, the elements are interpreted as the absolute amplitude of the change, *not* the change as a fraction of the allowed range. Setting `trust_region` to $[0.4, 0.4]$ would make the maximum allowed change for the second run be $[0 \pm 0.4, 0 \pm 0.4]$. Generally, specifying the `trust_region` is only needed if your experiment produces bad results when the parameters are changed significantly between runs. Simply delete this keyword if your experiment works with any set of parameters within the boundaries.

Halting conditions

The halting conditions define when the optimization will stop. We present four options here:

```
max_num_runs = 1000           #maximum number of runs
max_num_runs_without_better_params = 50 #maximum number of runs without finding
    ↪ better parameters
target_cost = 0.01            #optimization halts when a cost below this
    ↪ target is found
max_duration = 36000           #the optimization will not start a new
    ↪ iteration after max_duration seconds.
```

`max_num_runs` is the maximum number of runs that the optimization algorithm is allowed to run. `max_num_runs_without_better_params` is the maximum number of runs allowed before a lower cost and better parameters is found. If a run produces a cost that is less than `target_cost` then the optimization will stop. Finally, if the optimization proceeds for longer than `max_duration` seconds, then the optimization will end once the current iteration finishes (so the actual duration may exceed `max_duration` somewhat as the final iteration finishes).

When multiple halting conditions are set, the optimization process will halt when any one of them is met.

If you do not have any prior knowledge of the problem use only the keyword `max_num_runs` or `max_duration` and set it to the highest value you can wait for. If you have some knowledge about what the minimum attainable cost is or if there is some cost threshold you need to achieve, you might want to set the `target_cost`.

`max_num_runs_without_better_params` is useful if you want to let the optimization algorithm run as long as it needs until there is a good chance the global optimum has been found. If you have some time constraint for how long you would like the optimization to take, set `max_duration`.

If you do not want one of the halting conditions, simply delete it from your file. For example if you just wanted the algorithm to search as long as it can until it found a global minimum you could set:

```
max_num_runs_without_better_params = 50    #maximum number of runs without finding_  
↪better parameters
```

Learner Options

There are many learner specific options (and different learner algorithms) described in [Examples](#). Here we just present a common one:

```
cost_has_noise = True                                #whether the costs are corrupted by noise_  
↪or not
```

If the cost you provide has noise in it, meaning the cost you calculate would fluctuate if you did multiple experiments with the same parameters, then set this flag to `True`. If the costs you provide have no noise then set this flag to `False`. M-LOOP will automatically determine if the costs have noise in them or not, so if you are unsure, just delete this keyword and it will use the default value of `True`.

Timing options

M-LOOP's default optimization algorithm learns how the experiment works by fitting the parameters and costs using a gaussian process. This learning process can take some time. If M-LOOP is asked for new parameters before it has time to generate a new prediction, it will use the training algorithm to provide a new set of parameters to test. This allows for an experiment to be run while the learner is still thinking. The training algorithm by default is differential evolution. This algorithm is also used to do the first initial set of experiments which are then used to train M-LOOP. If you would prefer M-LOOP waits for the learner to come up with its best prediction before running another experiment you can change this behavior with the option:

```
no_delay = True                                       #wait for learner to make generate new_  
↪parameters or use training algorithms
```

Set `no_delay` to `true` to ensure that there are no pauses between experiments and set it to `false` if you want to give M-LOOP the time to come up with its most informed choice. Sometimes doing fewer more intelligent experiments will lead to an optimum quicker than many quick unintelligent experiments. You can delete the keyword if you are unsure and it will default to `True`.

File format options

You can set the file formats for the archives produced at the end and the files exchanged with the experiment with the options:


```

interface_file_type = 'txt'           #file types of *exp_input.mat* and *exp_
↪output.mat*
controller_archive_file_type = 'mat'   #file type of the controller archive
learner_archive_file_type = 'pkl'      #file type of the learner archive

```

`interface_file_type` controls the file format for the files exchanged with the experiment. `controller_archive_file_type` and `learner_archive_file_type` control the format of the respective archives.

There are three file formats currently available: ‘mat’ is for MATLAB readable files, ‘pkl’ if for python binary archives created using the [pickle package](#), and ‘txt’ human readable text files. For more details on these formats see [Data](#).

Visualization

By default M-LOOP will display a set of plots that allow the user to visualize the optimization process and the cost landscape. To change this behavior use the option:

```

visualizations = True

```

Set it to false to turn the visualizations off. For more details see [Visualizations](#).

Interface

There are many options for how to connect M-LOOP to your experiment. Here we consider the most generic method, writing and reading files to disk. For other options see [Interfaces](#). If you design a bespoke interface for your experiment please consider [Contributing](#) to the project by sharing your method with other users.

The file interface works under the assumption that your experiment follows the following algorithm.

1. Wait for the file `exp_input.txt` to be made on the disk in the same folder in which M-LOOP is run.
2. Read the parameters for the next experiment from the file (named params).
3. Delete the file `exp_input.txt`.
4. Run the experiment with the parameters provided and calculate a cost, and optionally the uncertainty.
5. Write the cost to the file `exp_output.txt`. Go back to step 1.

It is important you delete the file `exp_input.txt` after reading it, since it is used to as an indicator for the next experiment to run.

When writing the file `exp_output.txt` there are three keywords and values you can include in your file, for example after the first run your experiment may produce the following:

```

cost = 0.5
uncer = 0.01
bad = false

```

`cost` refers to the cost calculated from the experimental data. `uncer`, is optional, and refers to the uncertainty in the cost measurement made. Note, M-LOOP by default assumes there is some noise corrupting costs, which is fitted and compensated for. Hence, if there is some noise in your costs which you are unable to predict from a single measurement, do not worry, you do not have to estimate `uncer`, you can just leave it out. Lastly `bad` can be used to indicate an experiment failed and was not able to produce a cost. If the experiment worked set `bad = false` and if it failed set `bad = true`.

Note you do not have to include all of the keywords, you must provide at least a cost or the bad keyword set to true. For example a successful run can simply be:

```
cost = 0.3
```

and failed experiment can be as simple as:

```
bad = True
```

Once the *exp_output.txt* has been written to disk, M-LOOP will read it and delete it.

Parameters and cost function

Choosing the right parameterization of your experiment and cost function will be an important part of getting great results.

If you have time dependent functions in your experiment you will need to choose a parametrization of these function before interfacing them with M-LOOP. M-LOOP will take more time and experiments to find an optimum, given more parameters. But if you provide too few parameters, you may not be able to achieve your cost target.

Fortunately, the visualizations provided after the optimization will help you determine which parameters contributed the most to the optimization process. Try with whatever parameterization is convenient to start and use the data produced afterwards to guide you on how to better improve the parametrization of your experiment.

Picking the right cost function from experimental observables will also be important. M-LOOP will always find a global optimum as quickly as it can, but if you have a poorly chosen cost function, the global optimum may not be what you really wanted. Make sure you pick a cost function that will uniquely produce the result you want. Again, do not be afraid to experiment and use the data produced by the optimization runs to improve the cost function you are using.

Have a look at our [paper](#) on using M-LOOP to create a Bose-Einstein Condensate for an example of choosing a parametrization and cost function for an experiment.

Results

Once M-LOOP has completed the optimization, it will output results in several ways.

M-LOOP will print results to the console. It will give the parameters of the experimental run that produced the lowest cost. It will also provide a set of parameters which are predicted to produce the lowest average cost. If there is no noise in the costs your experiment produced, then the best parameters and predicted best parameters will be the same. If there was some noise in your costs then it is possible that there will be a difference between the two. This is because the noise might have caused a set of experimental parameters to produce a lower cost than they typically would due to a random fluke. The real optimal parameters that correspond to the minimum average cost are the predicted best parameters. In general, use the predicted best parameters (when provided) as the final result of the experiment.

M-LOOP will produce an archive for the controller and machine learner. The controller archive contains all the data gathered during the experimental run and also other configuration details set by the user. By default it will be a 'txt' file which is human readable. If the meaning of a keyword and its associated data in the file is unclear, just search the documentation with the keyword to find a description. The learner archive contains a model of the experiment produced by the machine learner algorithm, which is currently a gaussian process by default. By default it will also be a 'txt' file. For more detail on these files see [Data](#).

M-LOOP, by default, will produce a set of visualizations. These plots show the optimizations process over time and also predictions made by the learner of the cost landscape. For more details on these visualizations and their interpretation see [Visualizations](#).

2.2.2 Python controlled experiment

If you have an experiment that is already under python control you can use M-LOOP as an API. Below we go over the example python script *python_controlled_experiment.py*. You should also read over the [first tutorial](#) to get a general idea of how M-LOOP works.

When integrating M-LOOP into your laboratory remember that it will be controlling your experiment, not vice versa. Hence, at the top level of your python script you will execute M-LOOP which will then call on your experiment when needed. Your experiment will not be making calls of M-LOOP.

An example script for a python controlled experiment is given in the examples folder called *python_controlled_experiment.py*, which is included below

```

1  #Imports for python 2 compatibility
2  from __future__ import absolute_import, division, print_function
3  __metaclass__ = type
4
5  #Imports for M-LOOP
6  import mloop.interfaces as mli
7  import mloop.controllers as mlc
8  import mloop.visualizations as mlv
9
10 #Other imports
11 import numpy as np
12 import time
13
14 #Declare your custom class that inherits from the Interface class
15 class CustomInterface(mli.Interface):
16
17     #Initialization of the interface, including this method is optional
18     def __init__(self):
19         #You must include the super command to call the parent class, Interface,
20         ↳ constructor
21         super(CustomInterface, self).__init__()
22
23         #Attributes of the interface can be added here
24         #If you want to precalculate any variables etc. this is the place to do it
25         #In this example we will just define the location of the minimum
26         self.minimum_params = np.array([0,0.1,-0.1])
27
28         #You must include the get_next_cost_dict method in your class
29         #this method is called whenever M-LOOP wants to run an experiment
30         def get_next_cost_dict(self,params_dict):
31
32             #Get parameters from the provided dictionary
33             params = params_dict['params']
34
35             #Here you can include the code to run your experiment given a particular set,
36             ↳ of parameters
37             #In this example we will just evaluate a sum of sinc functions
38             cost = -np.sum(np.sinc(params - self.minimum_params))
39             #There is no uncertainty in our result
40             uncer = 0
41             #The evaluation will always be a success
42             bad = False
43             #Add a small time delay to mimic a real experiment
44             time.sleep(1)

```

(continues on next page)

(continued from previous page)

```

44     #The cost, uncertainty and bad boolean must all be returned as a dictionary
45     #You can include other variables you want to record as well if you want
46     cost_dict = {'cost':cost, 'uncer':uncer, 'bad':bad}
47     return cost_dict
48
49 def main():
50     #M-LOOP can be run with three commands
51
52     #First create your interface
53     interface = CustomInterface()
54     #Next create the controller. Provide it with your interface and any options you
55     ↪ want to set
56     controller = mlc.create_controller(interface,
57                                     max_num_runs = 1000,
58                                     target_cost = -2.99,
59                                     num_params = 3,
60                                     min_boundary = [-2,-2,-2],
61                                     max_boundary = [2,2,2])
62     #To run M-LOOP and find the optimal parameters just use the controller method
63     ↪ optimize
64     controller.optimize()
65
66     #The results of the optimization will be saved to files and can also be accessed
67     ↪ as attributes of the controller.
68     print('Best parameters found:')
69     print(controller.best_params)
70
71     #You can also run the default sets of visualizations for the controller with one
72     ↪ command
73     mlv.show_all_default_visualizations(controller)
74
75 #Ensures main is run when this code is run as a script
76 if __name__ == '__main__':
77     main()

```

Each part of the code is explained in the following sections.

Imports

The start of the script imports the libraries that are necessary for M-LOOP to work:

```

#Imports for python 2 compatibility
from __future__ import absolute_import, division, print_function
__metaclass__ = type

#Imports for M-LOOP
import mloop.interfaces as mli
import mloop.controllers as mlc
import mloop.visualizations as mlv

#Other imports
import numpy as np
import time

```

The first group of imports are just for python 2 compatibility. M-LOOP is targeted at python3, but has been designed to be bilingual. These imports ensure backward compatibility.

The second group of imports are the most important modules M-LOOP needs to run. The interfaces and controllers modules are essential, while the visualizations module is only needed if you want to view your data afterwards.

Lastly, you can add any other imports you may need.

Custom Interface

M-LOOP takes an object oriented approach to controlling the experiment. This is different than the functional approach taken by other optimization packages, like scipy. When using M-LOOP you must make your own class that inherits from the Interface class in M-LOOP. This class must implement a method called `get_next_cost_dict` that takes a set of parameters, runs your experiment and then returns the appropriate cost and uncertainty.

An example of the simplest implementation of a custom interface is provided below

```
#Declare your custom class that inherits from the Interface class
class SimpleInterface(mli.Interface):

    #the method that runs the experiment given a set of parameters and returns a
    ↪cost
    def get_next_cost_dict(self, params_dict):

        #The parameters come in a dictionary and are provided in a numpy array
        params = params_dict['params']

        #Here you can include the code to run your experiment given a
    ↪particular set of parameters
        #For this example we just evaluate a simple function
        cost = np.sum(params**2)
        uncer = 0
        bad = False

        #The cost, uncertainty and bad boolean must all be returned as a
    ↪dictionary
        cost_dict = {'cost':cost, 'uncer':uncer, 'bad':bad}
        return cost_dict
```

The code above defines a new class that inherits from the Interface class in M-LOOP. Note that this code is different from the example above; we will consider this later. It is slightly more complicated than just defining a method, however there is a lot more flexibility when taking this approach. You should put the code you use to run your experiment in the `get_next_cost_dict` method. This method is executed by the interface whenever M-LOOP wants a cost corresponding to a set of parameters.

When you actually run M-LOOP you will need to make an instance of your interface. To make an instance of the class above you would use:

```
interface = SimpleInterface()
```

This interface is then provided to the controller, which is discussed in the next section.

Dictionaries are used for both input and output of the method, to give the user flexibility. For example, if you had a bad run, you do not have to return a cost and uncertainty, you can just return a dictionary with bad set to True:

```
cost_dict = {'bad':True}
return cost_dict
```

By taking an object oriented approach, M-LOOP can provide a lot more flexibility when controlling your experiment. For example if you wish to start up your experiment or perform some initial numerical analysis you can add a customized constructor or `__init__` method for the class. We consider this in the main example:

```
#Declare your custom class that inherits from the Interface class
class CustomInterface(mli.Interface):

    #Initialization of the interface, including this method is optional
    def __init__(self):
        #You must include the super command to call the parent class, Interface,
        ↳ constructor
        super(CustomInterface, self).__init__()

        #Attributes of the interface can be added here
        #If you want to precalculate any variables etc. this is the place to do it
        #In this example we will just define the location of the minimum
        self.minimum_params = np.array([0,0.1,-0.1])

    #You must include the get_next_cost_dict method in your class
    #this method is called whenever M-LOOP wants to run an experiment
    def get_next_cost_dict(self,params_dict):

        #Get parameters from the provided dictionary
        params = params_dict['params']

        #Here you can include the code to run your experiment given a particular set,
        ↳ of parameters
        #In this example we will just evaluate a sum of sinc functions
        cost = -np.sum(np.sinc(params - self.minimum_params))
        #There is no uncertainty in our result
        uncer = 0
        #The evaluation will always be a success
        bad = False
        #Add a small time delay to mimic a real experiment
        time.sleep(1)

        #The cost, uncertainty and bad boolean must all be returned as a dictionary
        #You can include other variables you want to record as well if you want
        cost_dict = {'cost':cost, 'uncer':uncer, 'bad':bad}
        return cost_dict
```

In this code snippet we also implement a constructor with the `__init__()` method. Here we just define a numpy array which defines the minimum_parameter values. We can call this variable whenever we need in the `get_next_cost_dict` method. You can also define your own custom methods in your interface or even inherit from other classes.

Once you have implemented your own Interface running M-LOOP can be done in three lines.

Running M-LOOP

Once you have made your interface class, running M-LOOP can be as simple as three lines. In the example script M-LOOP is run in the main method:

```
def main():
    #M-LOOP can be run with three commands
```

(continues on next page)

(continued from previous page)

```

#First create your interface
interface = CustomInterface()
#Next create the controller. Provide it with your interface and any options you_
↪want to set
controller = mlc.create_controller(interface,
                                   max_num_runs = 1000,
                                   target_cost = -2.99,
                                   num_params = 3,
                                   min_boundary = [-2,-2,-2],
                                   max_boundary = [2,2,2])

#To run M-LOOP and find the optimal parameters just use the controller method_
↪optimize
controller.optimize()

```

In the code snippet we first make an instance of our custom interface class called interface. We then create an instance of a controller. The controller will run the experiment and perform the optimization. You must provide the controller with the interface and any of the M-LOOP options you would normally provide in the configuration file. In this case we give five options, which do the following:

1. `max_num_runs = 1000` sets the maximum number of runs to be 1000.
2. `target_cost = -2.99` sets a cost that M-LOOP will halt at once it has been reached.
3. `num_params = 3` sets the number of parameters to be 3.
4. `min_boundary = [-2,-2,-2]` defines the minimum values of each of the parameters.
5. `max_boundary = [2,2,2]` defines the maximum values of each of the parameters.

There are many other options you can use. Have a look at [Configuration File](#) for a detailed introduction into all the important configuration options. Remember you can include any option you would include in a configuration file as keywords for the controller. For more options you should look at all the config files in [Examples](#), or for a comprehensive list look at the [M-LOOP API](#).

Once you have created your interface and controller you can run M-LOOP by calling the optimize method of the controller. So in summary M-LOOP is executed in three lines:

```

interface = CustomInterface()
controller = mlc.create_controller(interface, [options])
controller.optimize()

```

Results

The results will be displayed on the console and also saved in a set of files. Have a read over [Results](#) for more details on the results displayed and saved. Also read [Data](#) for more details on data formats and how it is stored.

Within the python environment you can also access the results as attributes of the controller after it has finished optimization. The example includes a simple demonstration of this:

```

#The results of the optimization will be saved to files and can also be accessed_
↪as attributes of the controller.
print('Best parameters found:')
print(controller.best_params)

```

All of the results saved in the controller archive can be directly accessed as attributes of the controller object. For a comprehensive list of the attributes of the controller generated after an optimization run see the [M-LOOP API](#).

Visualizations

For each controller there is normally a default set of visualizations available. The visualizations for the Gaussian Process, the default optimization algorithm, are described in [Visualizations](#). Visualizations can be called through the visualization module. The example includes a simple demonstration of this:

```
#You can also run the default sets of visualizations for the controller with one_
↪command
mlv.show_all_default_visualizations(controller)
```

This code snippet will display all the visualizations available for that controller. There are many other visualization methods and options available that let you control which plots are displayed and when. See the [M-LOOP API](#) for details.

2.3 Interfaces

Currently M-LOOP supports three ways to interface your experiment

1. File interface where parameters and costs are exchanged between the experiment and M-LOOP through files written to disk. This approach is described in a [tutorial](#).
2. Shell interface where parameters and costs are exchanged between the experiment and M-LOOP through information piped through a shell (or command line). This option should be considered if you can execute your experiment using a command from a shell.
3. Implementing your own interface through the M-LOOP python API.

Each of these options is described below. If you have any suggestions for interfaces please consider [Contributing](#) to the project.

2.3.1 File interface

The simplest method to connect your experiment to M-LOOP is with the file interface where data is exchanged by writing files to disk. To use this interface you can include the option:

```
interface='file'
```

in your configuration file. The file interface happens to be the default, so this is not necessary.

The file interface works under the assumption that your experiment follows the following algorithm.

1. Wait for the file *exp_input.txt* to be made on the disk in the same folder in which M-LOOP is run.
2. Read the parameters for the next experiment from the file (named params).
3. Delete the file *exp_input.txt*.
4. Run the experiment with the parameters provided and calculate a cost, and optionally the uncertainty.
5. Write the cost to the file *exp_output.txt*. Go back to step 1.

It is important you delete the file *exp_input.txt* after reading it, since it is used to as an indicator for the next experiment to run.

When writing the file *exp_output.txt* there are three keywords and values you can include in your file, for example after the first run your experiment may produce the following:

```
cost = 0.5
uncer = 0.01
bad = false
```

`cost` refers to the cost calculated from the experimental data. `uncer`, is optional, and refers to the uncertainty in the cost measurement made. Note, M-LOOP by default assumes there is some noise corrupting costs, which is fitted and compensated for. Hence, if there is some noise in your costs which you are unable to predict from a single measurement, do not worry, you do not have to estimate `uncer`, you can just leave it out. Lastly `bad` can be used to indicate an experiment failed and was not able to produce a cost. If the experiment worked set `bad = false` and if it failed set `bad = true`.

Note you do not have to include all of the keywords, you must provide at least a cost or the `bad` keyword set to true. For example a successful run can simply be:

```
cost = 0.3
```

and failed experiment can be as simple as:

```
bad = True
```

Once the *exp_output.txt* has been written to disk, M-LOOP will read it and delete it.

2.3.2 Shell interface

The shell interface is used when experiments can be run from a command in a shell. M-LOOP will still need to be configured and executed in the same manner described for a file interface as describe in [tutorial](#). The only difference is how M-LOOP starts the experiment and reads data. To use this interface you must include the following options:

```
interface_type='shell'
command='./run_exp'
params_args_type='direct'
```

in the configuration file. The `interface` keyword simply indicates that you want M-LOOP to operate the experiment through the shell. The other two keywords need to be customized to your needs.

The `command` keyword should be provided with the command on the shell that runs the experiment. In the example above the executable would be *run_exp*. Note M-LOOP will try and execute the command in the folder that you run M-LOOP from. If this causes trouble you should just include the absolute address of your executable. Your command can be more complicated than a single word, for example if you want to include some options like `'./run_exp --verbose -U'` this would also be acceptable.

The `params_args_type` keyword controls how M-LOOP delivers the parameters to the executable. If you use the `'direct'` option the parameters will just be fed directly to the experiment as arguments. For example if the command was `./run_exp` and the parameters to test next were 1.3, -23 and 12, M-LOOP would execute the following command:

```
./run_exp 1.3 -23 12
```

The other `params_args_type` option is `'named'`, in which case each parameter is fed to the experiment as a named option. If the optional `param_names` argument is provided in the configuration file, then M-LOOP will use those as the names of the arguments passed to the executable. If `param_names` is not provided, then M-LOOP will

default to using arguments named `param1`, `param2` and so on. Given the same parameters as before, M-LOOP would execute the command:

```
./run_exp --param1 1.3 --param2 -23 --param3 12
```

After the experiment has run and a cost (and uncertainty or bad value) has been found they must be provided back to M-LOOP through the shell. For example if you experiment completed with a cost 1.3, uncertainty 0.1 you need to program your executable to print the following to the shell:

```
M-LOOP_start
cost = 1.3
uncer = 0.1
M-LOOP_end
```

You can also output other information to the shell and split up the information you provide to M-LOOP if you wish. The following output would also valid:

```
Running experiment... Experiment complete.
Checking it was valid... It worked.
M-LOOP_start
bad = False
M-LOOP_end
Calculating cost... Was 3.2.
M-LOOP_start
cost = 3.2
M-LOOP_end
```

2.3.3 Python interfaces

If your experiment is controlled in python you can use M-LOOP as an API in your own custom python script. In this case you must create your own implementation of the abstract interface class to control the experiment. This is explained in detail in the *[tutorial for python controlled experiments](#)*.

2.4 Data

M-LOOP saves all data produced by the experiment in archives which are saved to disk during and after the optimization run. The archives also contain information derived from the data, including the machine learning model for how the experiment works. Here we explain how to interpret the file archives.

2.4.1 File Formats

M-LOOP currently supports three file formats for all file input and output.

- ‘txt’ text files: Human readable text files. This is the default file format for all outputs. The advantage of text files is they are easy to read, and there will be no format compatibility issues in the future. However, there will be some loss of precision in your data. To ensure you keep all significant figure you may want to use ‘pkl’ or ‘mat’.
- ‘mat’ MATLAB files: Matlab files that can be opened and written with MATLAB or `numpy`.
- ‘pkl’ pickle files: a serialization of a python dictionary made with `pickle` [<https://docs.python.org/3/library/pickle.html>](https://docs.python.org/3/library/pickle.html). Your data can be retrieved from this dictionary using the appropriate keywords.

2.4.2 File Keywords

The archives contain a set of keywords/variable names with associated data. The quickest way to understand what the values mean for a particular keyword is to search the documentation for a description.

For a comprehensive list of all the keywords looks at the attributes described in the API.

For the controller archive see *controllers*.

For the learner archive see *learners*. The generic keywords are described in the class `Learner`, with learner specific options described in the derived classes, for example `GaussianProcessLearner`.

2.4.3 Converting files

If for whatever reason you want to convert files between the formats you can do so using the utilities module of M-LOOP. For example the following python code will convert the file `controller_archive_2016-08-18_12-18.pkl` from a 'pkl' file to a 'mat' file:

```
import mloop.utilities as mlu

saved_dict = mlu.get_dict_from_file('./M-LOOP_archives/controller_archive_2016-08-18_
↪12-18.pkl')
mlu.save_dict_to_file(saved_dict, './M-LOOP_archives/controller_archive_2016-08-18_12-
↪18.mat')
```

2.5 Visualizations

At the end of an optimization run a set of visualizations will be produced by default.

The number of visualizations, and what they show, will depend on what controller type you use. Generally there will always be three plots that present the data from the controller. In addition, for most controller types there will be more plots which present data gathered by the learner. The information presented in these plots is explained below. The plots which start with *Controller:* are generated from the controller archive, while plots that start with *Learner:* are generated from the learner archive.

Often optimization runs can involve many parameters, which can make the plots that display values of different parameters too busy and difficult to interpret. To avoid this issue, low-level plotting functions support an optional `parameter_subset` argument. Passing a list of indices for `parameter_subset` instructs those functions to only display the data for the parameters corresponding to those indices. The high-level plotting functions, i.e. the ones that generate all of the plots that are implemented for a given archive, instead support an optional `max_parameters_per_plot` argument. When that argument is provided, plots with many parameters will be broken up into several different plots, each displaying the data for at most `max_parameters_per_plot` arguments.

Occasionally the legend can obscure some of the data in the plots. The positions of legends can be adjusted by calling `mloop.visualizations.set_legend_location()`, which accepts any of the values that can be used for `loc` in `matplotlib's legend()` function. For example, to set the legend outside of the plot, you can use `set_legend_location((1, 0))`. Note that `set_legend_location()` must be called before generating a plot in order for it to have an effect. To move the legend in an existing plot, call `set_legend_location()` then recreate the plot.

2.5.1 Controller Visualizations

Regardless of controller type, there are always three plots produced for the controller. These are as follows:

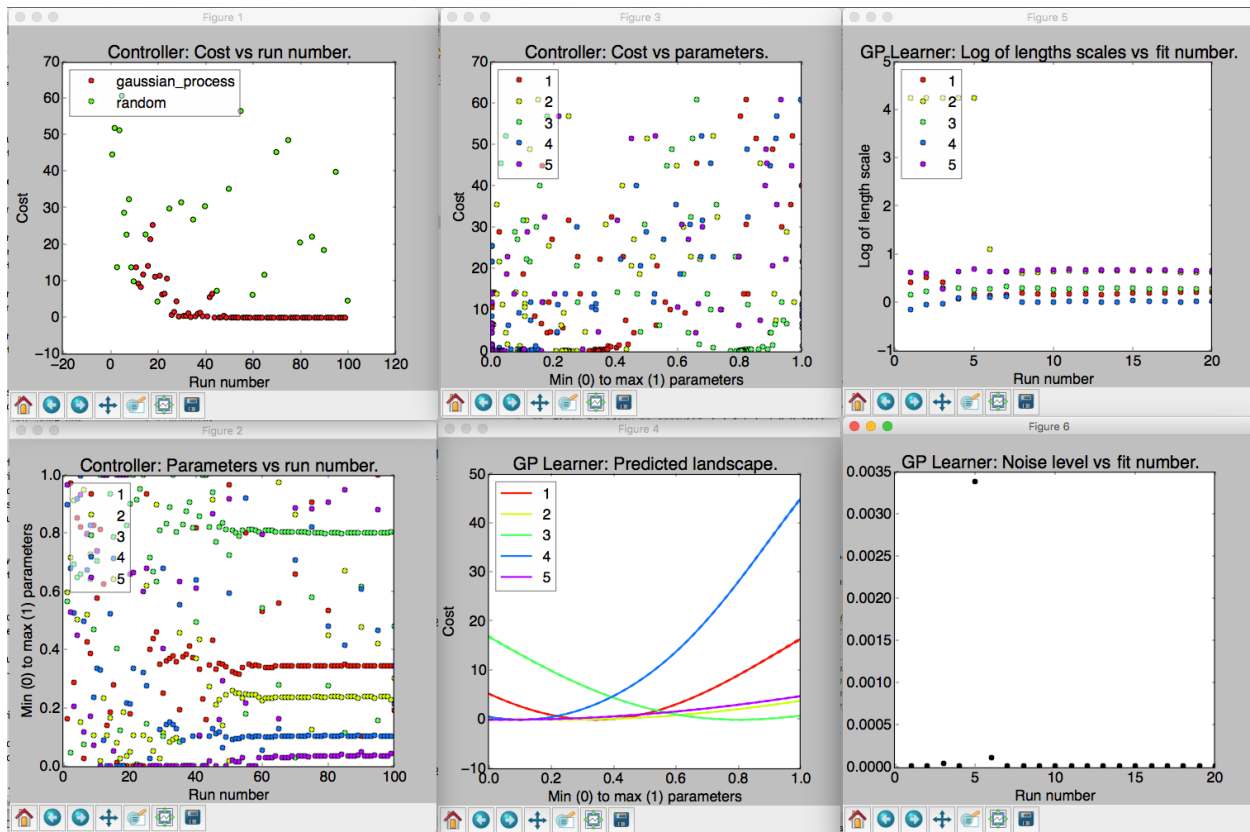


Fig. 1: An example of the six visualizations automatically produced when M-LOOP is run with the default controller, the Gaussian process machine learner.

- **Controller: Cost vs run number.** Here the cost returned by the experiment versus run number is plotted. The legend shows what algorithm was used to generate the parameters tested by the experiment. If you use the Gaussian process or neural net, there will also be another algorithm used throughout the optimization algorithm in order to (a) ensure parameters are generated fast enough and (b) add new prior free data to ensure the Gaussian process converges to the correct model.
- **Controller: Parameters vs run number.** The parameters values are all plotted against the run number. Note the parameters will all be scaled between their minimum and maximum value. The legend indicates which color corresponds to which parameter.
- **Controller: Cost vs parameters.** The cost versus the parameters. Here each of the parameters tested are plotted against the cost they returned as a set. Again the parameter values are all scaled between their minimum and maximum values.

2.5.2 Learner Visualizations

Which visualizations are generated to display results from the learner depend on what learner was used for the optimization. Each section below describes the plots produced by a different learner.

Gaussian Process

- **GP Learner: Predicted landscape.** 1D cross sections of the predicted cost landscape about point with the best recorded cost are plotted for each parameter. The color of the cross section corresponds to the parameter that is varied in the cross section. This predicted landscape is generated by the model fit to the experiment by the Gaussian process. Be sure to check after an optimization run that all parameters contributed. If one parameter produces a flat cross section, it is likely it did not have any influence on the final cost or its length scale was not fit well. You may want to remove it on the next optimization run or set limits on its fitted length scale. If a trust region was specified, then markers indicating the upper and lower limits of the trust region for each parameter will also be displayed in the plot.
- **GP Learner: Log of length scales vs fit number.** The Gaussian process fits a correlation length to each of the parameters in the experiment. Here we see a plot of the correlation lengths versus fit number. The last correlation lengths (highest fit number) are the most reliable values. Correlation lengths indicate how sensitive the cost is to changes in these parameters. If the correlation length is large, the parameter has a very little influence on the cost; if the correlation length is small, the parameter will have a very large influence on the cost. The correlation lengths are not precisely estimated. They should only be trusted accurate to +/- an order of magnitude. If a parameter has an extremely large value at the end of the optimization, say 5 or more, it is unlikely to have much affect on the cost and should be removed on the next optimization run.
- **GP Learner: Noise level vs fit number.** This is the estimated noise in the costs as a function of fit number. The most reliable estimate of the noise level will be the last value (highest fit number). The noise level is useful for quantifying the intrinsic noise and uncertainty in your cost value. Most other optimization algorithms will not provide this estimate. The noise level estimate may be helpful when isolating what part of your system can be optimized and what part is due to random fluctuations. This plot will only be generated if the `cost_has_noise` option was set to `True`.

Neural Net

- **Neural Net Learner: Predicted Landscape.** The neural net learner visualizer produces a plot showing 1D cross sections of the predicted cost about the point with the best recorded cost, just as the Gaussian learner visualizer does. The main difference is that the neural net learner will actually produce multiple instances of this kind of plot. The first few show the results predicted by each of the independent neural nets. The last one shows the average of the costs predicted by each independent neural net, as well as dashed lines showing the maximum and minimum values predicted by any neural net for each point. If a trust region was specified, then

markers indicating the upper and lower limits of the trust region for each parameter will also be displayed in the plot.

- **Neural Net Learner: Cost Surface.** This plot is only generated if the optimization had two parameters. It plots the predicted cost landscape as a surface in a 3D space where the x and y axes are the optimization parameters and the z axis shows the predicted cost.
- **Neural Net Learner: Cost Surface Density.** This plot is also only generated if the optimization had two parameters. It shows the same data as the cost surface plot, except that the predicted cost is plotted using a color scale rather than using a third dimension.
- **Neural Net Learner: Loss vs Epoch.** While fitting the neural nets their loss is calculated, which is a measure of how well the predicted cost fits the measured values. In M-LOOP this is measured as the mean of the square of the deviation between the predicted and measured values, plus a contribution from the regularization loss which is used to reduce overfitting. Each neural net records its loss every ten training epochs. This plot displays those recorded losses. Note that an “epoch” here is not the same as a run of the experiment. One epoch corresponds to one iteration over the full data set while fitting a neural net. Generally the fitting routine will go through many epochs during one fit, and the number of epochs per fit will vary.
- **Neural Net Learner: Regularization History.** The neural nets use L2 regularization to smooth their predicted landscapes in an attempt to avoid overfitting the data. The strength of the regularization is set by the regularization coefficient, which is a hyperparameter that is tuned during the optimization if `update_hyperparameters` is set to `True`. Generally larger regularization coefficient values force the landscape to be smoother while smaller values allow it to vary more quickly. A value too large can lead to underfitting while a value too small can lead to overfitting. The ideal regularization coefficient value will depend on many factors, such as the shape of the actual cost landscape, the SNR of the measured costs, and even the number of measured costs. This method plots the initial regularization coefficient value and the optimal values found for the regularization coefficient when performing the hyperparameter tuning. One curve showing the history of values used for the regularization coefficient is plotted for each neural net. If `update_hyperparameters` was set to `False` during the optimization, then only the initial default value will be plotted.

Differential Evolution

- **Differential Evolution Learner: Parameters vs Generation.** This plot displays the values tried for each of the parameters for each generation. Because there are multiple runs per generation, there are many points for each parameter within each generation.
- **Differential Evolution Learner: Costs vs Generation.** This plot displays the measured costs for each generation. Because there are multiple runs per generation, there are many different values for the cost plotted for each generation.

Nelder-Mead

As of yet there are no visualizations implemented for the Nelder-Mead learner. The controller’s archive may still be plotted though when Nelder-Mead is used.

Random

As of yet there are no visualizations implemented for the random learner. The controller’s archive may still be plotted though when the random controller is used.

2.5.3 Reproducing visualizations

If you have a controller and learner archive and would like to examine the visualizations again, it is best to do so using the *M-LOOP API*. For example the following code will plot the visualizations again from the files *controller_archive_2016-08-23_13-59.mat* and *learner_archive_2016-08-18_12-18.pkl*:

```
import mloop.visualizations as mlv

mlv.configure_plots()
mlv.show_all_default_visualizations_from_archive(
    controller_filename='controller_archive_2016-08-23_13-59.mat',
    learner_filename='learner_archive_2016-08-18_12-18.pkl',
)
```

2.6 Examples

M-LOOP includes a series of example configuration files for each of the controllers and interfaces. The examples can be found in examples folder. For some controllers there are two files, ones ending with *_basic_config* which includes the standard configuration options and *_complete_config* which include a comprehensive list of all the configuration options available.

The options available are also comprehensively documented in the *M-LOOP API* as keywords for each of the classes. However, the quickest and easiest way to learn what options are available, if you are not familiar with python, is to just look at the provided examples.

Each of the example files is used when running tests of M-LOOP. So please copy and modify them elsewhere if you use them as a starting point for your configuration file.

2.6.1 Interfaces

There are currently two interfaces supported: ‘file’ and ‘shell’. You can specify which interface you want with the option:

```
interface_type = [name]
```

The default will be ‘file’. The specific options for each of the interfaces are described below.

File Interface

The file interface exchanges information with the experiment by writing files to disk. You can change the names of the files used for the file interface and their type. The file interface options are described in *file_interface_config.txt*.

```
#File Interface Options
#-----

interface_type = 'file'           #The type of interface
interface_out_filename = 'exp_input' #The filename of the file output by the
↪interface and input into the experiment
interface_in_filename = 'exp_output' #The filename o the file input into the
↪interface and output by the experiment
interface_file_type = 'txt'       #The file_type of both the input and output
↪files, can be 'txt', 'pkl' or 'mat'.
```

Shell Interface

The shell interface is for experiments that can be run through a command executed in a shell. Information is then piped between M-LOOP and the experiment through the shell. You can change the command to run the experiment and the way the parameters are formatted. The shell interface options are described in *shell_interface_config.txt*

```
#Command Line Interface Options
#-----

interface_type = 'shell'          #The type of interface
command = 'python shell_script.py' #The command for the command line to run the
↪experiment to get a cost from the parameters
params_args_type = 'direct'      #The format of the parameters when providing
↪them on the command line. 'direct' simply appends them, e.g. python shell_script.py
↪7 2 1, 'named' names each parameter, e.g. python shell_script.py --param1 7 --
↪param2 2 --param3 1
```

2.6.2 Controllers

There are currently five built-in controller types: 'gaussian_process', 'neural_net', 'differential_evolution', 'nelder_mead', and 'random'. The default controller is 'gaussian_process'. You can set which controller you want to use with the option:

```
controller_type = [name]
```

Each of the controllers and their specific options are described below. There is also a set of common options shared by all controllers which is described in *controller_config.txt*. The common options include the parameter settings and the halting conditions.

```
#General Controller Options
#-----

#Halting conditions
max_num_runs = 1000          #number of planned runs
target_cost = 0.1           #cost to beat
max_num_runs_without_better_params = 100 #max allowed number of runs between finding
↪better parameters
max_duration = 36000         #the optimization will not start a new
↪iteration after max_duration seconds.

#Parameter controls
num_params = 2              #Number of parameters
min_boundary = [0,0]        #Minimum value for each parameter
max_boundary = [2,2]        #Maximum value for each parameter

#Filename related
controller_archive_filename = 'agogo' #filename prefix for controller archive,
↪can include path
controller_archive_file_type = 'mat'  #file_type for controller archive
learner_archive_filename = 'ogoga'    #filename prefix for learner archive, can
↪include path
learner_archive_file_type = 'pkl'     #file_type for learner archive
archive_extra_dict = {'test':'this_is'} #dictionary of any extra data to be put in
↪archive
```

In addition to the built-in controllers, you can also use controllers provided by external Python packages. In this

case, you can set `controller_type` to `'module_name:controller_name'`, where `module_name` is the name of the Python module containing the controller and `controller_name` is the name of the controller class (or the function that creates the controller object). The parameters for such controllers should be documented by the corresponding external packages.

Gaussian process

The Gaussian process controller is the default controller. It uses a [Gaussian process](#) to develop a model for how the parameters relate to the measured cost, effectively creating a model for how the experiment operates. This model is then used when picking new points to test.

There are two example files for the Gaussian-process controller: `gaussian_process_simple_config.txt` which contains the basic options.

```
#Gaussian Process Basic Options
#-----

#General options
max_num_runs = 100           #number of planned runs
target_cost = 0.1

#Gaussian process controller options
controller_type = 'gaussian_process'  #name of controller to use
num_params = 3                       #number of parameters
min_boundary = [-0.8, -0.9, -1.1]    #minimum boundary
max_boundary = [0.8, 0.9, 1.1]       #maximum boundary
trust_region = 0.4                  #maximum % move distance from best params
cost_has_noise = False              #whether cost function has noise
```

`gaussian_process_complete_config.txt` which contains a comprehensive list of options.

```
#Gaussian Process Complete Options
#-----

#General options
max_num_runs = 100           #number of planned runs
target_cost = 0.1            #cost to beat

#Gaussian process options
controller_type = 'gaussian_process'
num_params = 2               #number of parameters
min_boundary = [-10., -10.]  #minimum boundary
max_boundary = [10., 10.]    #maximum boundary
param_names = ['a', 'b']    #names for parameters
length_scale = [1.0]         #initial lengths scales for GP
length_scale_bounds = [1e-5, 1e5] #limits on values fit for length_scale
minimum_uncertainty = 1e-8    #minimum uncertainty of cost, required to
↳avoid fitting errors
cost_has_noise = True        #whether cost function has noise
noise_level = 0.1            #initial noise level estimate, cost's variance
↳(standard deviation squared)
noise_level_bounds = [1e-5, 1e5] #limits on values fit for noise_level
update_hyperparameters = True #whether noise level and lengths scales are
↳updated
trust_region = [5, 5]        #maximum move distance from best params
default_bad_cost = 10        #default cost for bad run
default_bad_uncertainty = 1   #default uncertainty for bad run
```

(continues on next page)

(continued from previous page)

```

learner_archive_filename = 'a_word'      #filename of gp archive, can include path
learner_archive_file_type = 'mat'        #file type of archive
predict_global_minima_at_end = True      #find predicted global minima at end
no_delay = True                          #whether to wait for the GP to make_
↳ predictions or not. Default True (do not wait)

#Training source options
training_type = 'random'                  #training type can be random, differential_
↳ evolution, or nelder_mead
first_params = [1.9, -1.0]                #first parameters to try in initial training
num_training_runs = 20                    #number of training runs before using machine_
↳ learner to pick parameters
training_filename = None                  #filename for training from previous experiment

#if you use nelder_mead for the initial training source see the_
↳ CompleteNelderMeadConfig.txt for options.

```

Note that `noise_level` corresponds to a variance, not a standard deviation. In particular `noise_level` estimates the variance if the cost for a given set of parameters were measured many times. This is in contrast to the cost uncertainty that the user optionally passes to M-LOOP with the cost itself; that should be the standard deviation. In other words the cost uncertainty should estimate the standard deviation if the cost for a given set of parameters were measured many times.

Neural net

The neural net controller also uses a machine-learning-based algorithm. It is similar to the Gaussian process controller in that it constructs a model of how the parameters relate to the cost and then uses that model for the optimization. However instead of modeling with a Gaussian process, it works by modeling with a sampled neural net.

The neural net models aren't always as robust and reliable as the Gaussian process. However, the time required to fit a Gaussian process scales as the cube of the number of data points, while the time to train a neural net only scales linearly. Often the Gaussian process fitting can be prohibitively slow for long optimizations with many parameters, while the neural net training remains relatively fast. That makes the neural net controller a good choice for high-dimensional optimizations.

There are two example files for the neural net controller: `neural_net_simple_config.txt` which contains the basic options.

```

#Neural Net Basic Options
#-----

#General options
max_num_runs = 100                      #number of planned runs
target_cost = 0.1

#Neural net controller options
controller_type = 'neural_net'           #name of controller to use
num_params = 3                           #number of parameters
min_boundary = [-0.8, -0.9, -1.1]        #minimum boundary
max_boundary = [0.8, 0.9, 1.1]           #maximum boundary
trust_region = 0.4                       #maximum move distance from best params, as a_
↳ fraction of the allowed range

```

`neural_net_complete_config.txt` which contains a comprehensive list of options.

```

#Neural Net Complete Options
#-----

#General options
max_num_runs = 100                #number of planned runs
target_cost = 0.1                 #cost to beat

#Neural net controller options
controller_type = 'neural_net'    #name of controller to use
num_params = 2                   #number of parameters
min_boundary = [-10., -10.]      #minimum boundary
max_boundary = [10., 10.]        #maximum boundary
param_names = ['a', 'b']        #names for parameters
minimum_uncertainty = 1e-8       #minimum uncertainty of cost, required to
    ↳avoid fitting errors
trust_region = [5, 5]            #maximum move distance from best params
default_bad_cost = 10            #default cost for bad run
default_bad_uncertainty = 1      #default uncertainty for bad run
learner_archive_filename = 'a_word' #filename of neural net learner archive, can
    ↳include path
learner_archive_file_type = 'txt' #file type of neural net learner archive
predict_global_minima_at_end = True #find predicted global minima at end
no_delay = True                  #whether to wait for the GP to make
    ↳predictions or not. Default True (do not wait)
update_hyperparameters = False   #whether hyperparameters should be tuned to
    ↳avoid overfitting. Default False.

#Training source options
training_type = 'random'         #training type can be random, differential_
    ↳evolution, or nelder_mead
first_params = [1.9, -1.0]       #first parameters to try in initial training
num_training_runs = 20           #number of training runs before using machine_
    ↳learner to pick parameters
training_filename = None         #filename for training from previous experiment

#if you use nelder_mead for the initial training source see the_
    ↳CompleteNelderMeadConfig.txt for options.

```

Differential evolution

The differential evolution (DE) controller uses a [DE algorithm](#) for optimization. DE is a type of evolutionary algorithm, and is historically the most commonly used in automated optimization. DE will eventually find a global solution, however it can take many experiments before it does so.

There are two example files for the differential evolution controller: *differential_evolution_simple_config.txt* which contains the basic options.

```

#Differential Evolution Basic Options
#-----

#General options
max_num_runs = 500                #number of planned runs
target_cost = 0.1                 #cost to beat

#Differential evolution controller options
controller_type = 'differential_evolution'

```

(continues on next page)

(continued from previous page)

```

num_params = 1                #number of parameters
min_boundary = [-4.8]          #minimum boundary
max_boundary = [10.0]          #maximum boundary
trust_region = 0.6             #maximum % move distance from best params
first_params = [5.3]           #first parameters to try

```

differential_evolution_complete_config.txt which contains a comprehensive list of options.

```

#Differential Evolution Complete Options
#-----

#General options
max_num_runs = 500             #number of planned runs
target_cost = 0.1              #cost to beat

#Differential evolution controller options
controller_type = 'differential_evolution'
num_params = 2                 #number of parameters
min_boundary = [-1.2, -2]      #minimum boundary
max_boundary = [10.0, 4]       #maximum boundary
param_names = ['a', 'b']      #names for parameters
trust_region = [3.2, 3.1]      #maximum move distance from best params
first_params = None           #first parameters to try if None a random set
↳ of parameters is chosen
evolution_strategy = 'best2'   #evolution strategy can be 'best1', 'best2',
↳ 'rand1' and 'rand2'. Best uses the best point, rand uses a random one, the number
↳ indicates the number of directions added.
population_size = 10           #a multiplier for the population size of a
↳ generation
mutation_scale = (0.4, 1.1)    #the minimum and maximum value for the
↳ mutation scale factor. Each generation is randomly selected from this. Each value
↳ must be between 0 and 2.
cross_over_probability = 0.8    #the probability a parameter will be resampled
↳ during a mutation in a new generation
restart_tolerance = 0.02        #the fraction the standard deviation in the
↳ costs of the population must reduce from the initial sample, before the search is
↳ restarted.
learner_archive_filename = 'a_word' #filename of the learner archive, can include
↳ path
learner_archive_file_type = 'mat' #file type of the learner archive

```

Nelder–Mead

The Nelder–Mead controller implements the [Nelder–Mead method](#) for optimization. You can control the starting point and size of the initial simplex of the method with the configuration file.

There are two example files for the Nelder–Mead controller: *nelder_mead_simple_config.txt* which contains the basic options.

```

#Nelder-Mead Basic Options
#-----

#General options
max_num_runs = 100             #number of planned runs
target_cost = 0.1              #cost to beat

```

(continues on next page)

(continued from previous page)

```
#Specific options
controller_type = 'nelder_mead'
num_params = 3           #number of parameters
min_boundary = [-1, -1, -1] #minimum boundary
max_boundary = [1, 1, 1]  #maximum boundary
initial_simplex_scale = 0.4 #initial size of simplex relative to the boundary_
↪size.
```

nelder_mead_complete_config.txt which contains a comprehensive list of options.

```
#Nelder-Mead Complete Options
#-----

#General options
max_num_runs = 100           #number of planned runs
target_cost = 0.1           #cost to beat

#Specific options
controller_type = 'nelder_mead'
num_params = 5               #number of parameters
min_boundary = [-1.1, -1.2, -1.3, -1.4, -1.5] #minimum boundary
max_boundary = [1.1, 1.1, 1.1, 1.1, 1.1]      #maximum boundary
param_names = ['a', 'b', 'c', 'd', 'e']       #names for parameters
initial_simplex_corner = [-0.21, -0.23, -0.24, -0.23, -0.25] #initial corner of the_
↪simplex
initial_simplex_displacements = [1, 1, 1, 1, 1] #initial displacements_
↪for the N+1 (i this case 6) points of the simplex
```

Random

The random optimization algorithm picks parameters randomly from a uniform distribution from within the parameter bounds or trust region.

There are two example files for the random controller: *random_simple_config.txt* which contains the basic options.

```
#Random Basic Options
#-----

#General options
max_num_runs = 10           #number of planned runs

#Random controller options
controller_type = 'random'
num_params = 1               #number of parameters
min_boundary = [1.2]         #minimum boundary
max_boundary = [10.0]        #maximum boundary
trust_region = 0.1           #maximum % move distance from best params
first_params = [5.3]         #first parameters to try
```

random_complete_config.txt which contains a comprehensive list of options.

```
#Random Complete Options
#-----
```

(continues on next page)

(continued from previous page)

```

#General options
max_num_runs = 20                                #number of planned runs

#Random controller options
controller_type = 'random'
num_params = 2                                    #number of parameters
min_boundary = [1.2, -2]                         #minimum boundary
max_boundary = [10.0, 4]                         #maximum boundary
param_names = ['a', 'b']                        #names for parameters
trust_region = [0.2, 0.5]                       #maximum move distance from best params
first_params = [5.1, -1.0]                      #first parameters to try

```

2.6.3 Logging

You can control the filename of the logs and also the level which is reported to the file and the console. For more information see [logging levels](#). The logging options are described in *logging_config.txt*.

```

#Logging Options
#-----

log_filename = 'cl_run'                          #Prefix for logging filename, can include path
file_log_level=logging.DEBUG                     #Logging level saved in file
console_log_level=logging.WARNING                #Logging level presented to console, normally INFO

```

2.6.4 Extras

Extras refers to options related to post processing your data once the optimization is complete. Currently the only extra option is for visualizations. The extra options are described in *extras_config.txt*.

```

#Extra Options
#-----

visualizations=False                            #whether plots should be presented after run

```

2.7 Contributing

If you use M-LOOP please consider contributing to the project. There are many quick and easy ways to help out.

- If you use M-LOOP be sure to cite the paper where it was first used: ‘Fast machine-learning online optimization of ultra-cold-atom experiments’, Sci Rep 6, 25890 (2016).
- Star and watch the [M-LOOP GitHub](#).
- Make a suggestion on what features you would like added, or report an issue, on the [GitHub](#) or by [email](#).
- Contribute your own code to the [M-LOOP GitHub](#). This could be the interface you designed, more options, or a completely new solver.

Finally spread the word! Let others know the success you have had with M-LOOP and recommend they try it too.

2.7.1 Contributors

M-LOOP is written and maintained by [Michael R Hush <MichaelRHush@gmail.com>](mailto:MichaelRHush@gmail.com)

Other contributors, listed alphabetically, are:

- John W. Bastian - design, first demonstration
- Patrick J. Everitt - testing, design, first demonstration
- Kyle S. Hardman - design, first demonstration
- Anton van den Hengel - design, first demonstration
- Joe J. Hope - design, first demonstration
- Carlos C. N. Kuhn - first demonstration
- Andre N. Luiten - first demonstration
- Gordon D. McDonald - first demonstration
- Manju Perumbil - first demonstration
- Ian R. Petersen - first demonstration
- Ciaran D. Quinlivan - first demonstration
- Alex Ratcliff - testing
- Nick P. Robins - first demonstration
- Mahasen A. Sooriyabandara - first demonstration
- Richard Taylor - testing
- Zak Vendeiro - ease of use enhancements
- Paul B. Wigley - testing, design, first demonstration

2.8 Changelog

M-LOOP's changelog is available on its [Github releases page](#).

2.9 M-LOOP API

M-LOOP can also be used as a library in python. This is particularly useful if the experiment you are optimizing can be controlled by python.

2.9.1 mloop

M-LOOP: Machine-Learning Online Optimization Package

Python package for performing automated, online optimization of scientific experiments or anything that can be computer controlled. The package employs machine learning algorithms to rapidly find optimal parameters for systems under control.

If you use this package please cite the article <http://www.nature.com/articles/srep25890>.

To contribute to the project or report a bug visit the project's github <https://github.com/michaelhush/M-LOOP>.

2.9.2 controllers

Module of all the controllers used in M-LOOP. The controllers, as the name suggests, control the interface to the experiment and all the learners employed to find optimal parameters.

```
class mloop.controllers.Controller(interface,          max_num_runs=inf,          target_cost=-
                                inf,          max_num_runs_without_better_params=inf,
                                max_duration=inf, controller_archive_filename='controller_archive',
                                controller_archive_file_type='txt',
                                archive_extra_dict=None,          start_datetime=None,
                                **kwargs)
```

Bases: object

Abstract class for controllers.

The controller controls the entire M-LOOP process. The controllers for each algorithm all inherit from this class. The class stores a variety of data which all algorithms use and also includes all of the archiving and saving features.

In order to implement your own controller class the minimum requirement is to add a learner to the learner variable and implement the *next_parameters()* method where you provide the appropriate information to the learner and get the next parameters. See the *RandomController* for a simple implementation of a controller. Note the first three keywords are all possible halting conditions for the controller. If any of them are satisfied the controller will halt (meaning an OR condition is used). This base class also creates an empty attribute *self.learner*. The simplest way to make a working controller is to assign a learner of some kind to this variable, and add appropriate queues and events from it.

Parameters *interface* (*interface*) – The interface process. It is run by the controller.

Keyword Arguments

- **max_num_runs** (*Optional [float]*) – The number of runs before the controller stops. If set to *float('+inf')* the controller will run forever assuming no other halting conditions are met. Default *float('+inf')*, meaning the controller will run until another halting condition is met.
- **target_cost** (*Optional [float]*) – The target cost for the run. If a run achieves a cost lower than the target, the controller is stopped. Default *float('-inf')*, meaning the controller will run until another halting condition is met.
- **max_num_runs_without_better_params** (*Optional [float]*) – The optimization will halt if the number of consecutive runs without improving over the best measured value thus far exceeds this number. Default *float('+inf')*, meaning the controller will run until another halting condition is met.
- **max_duration** (*Optional [float]*) – The maximum duration of the optimization, in seconds of wall time. The actual duration may exceed this value slightly, but no new iterations will start after *max_duration* seconds have elapsed since *start_datetime*. Default is *float('+inf')*, meaning the controller will run until another halting condition is met.
- **controller_archive_filename** (*Optional [string]*) – Filename for archive. The archive contains costs, parameter history and other details depending on the controller type. Default *'controller_archive'*.
- **controller_archive_file_type** (*Optional [string]*) – File type for archive. Can be either *'txt'* for a human readable text file, *'pkl'* for a python pickle file, *'mat'* for a matlab file, or *None* to forgo saving a controller archive. Default *'txt'*.
- **archive_extra_dict** (*Optional [dict]*) – A dictionary with any extra variables that are to be saved to the archive. If *None*, nothing is added. Default *None*.

- **start_datetime** (*Optional datetime*) – Datetime for when the controller was started.

params_out_queue

Queue for parameters to next be run by the experiment.

Type queue

costs_in_queue

Queue for costs (and other details) that have been returned by experiment.

Type queue

interface_error_queue

Queue for returning errors encountered by the interface.

Type queue

end_interface

Event used to trigger the end of the interface.

Type event

learner

The placeholder for the learner. Creating this variable is the minimum requirement to make a working controller class.

Type None

learner_params_queue

The parameters queue for the learner.

Type queue

learner_costs_queue

The costs queue for the learner.

Type queue

end_learner

Event used to trigger the end of the learner.

Type event

num_in_costs

Counter for the number of costs received.

Type int

num_out_params

Counter for the number of parameters received.

Type int

out_params

List of all parameters sent out by controller.

Type list

out_extras

Any extras associated with the output parameters.

Type list

in_costs

List of costs received by controller.

Type list

in_uncers

List of uncertainties received by controller.

Type list

best_cost

The lowest, and best, cost received by the learner.

Type float

best_uncer

The uncertainty associated with the best cost.

Type float

best_params

The best parameters received by the learner.

Type array

best_index

The run number that produced the best cost.

Type float

`_enforce_boundaries` (*params*)

Enforce the minimum and maximum parameter boundaries.

If the values in *params* extend outside of the allowed boundaries set by *self.min_boundary* and *self.max_boundary* by a nontrivial amount, then this method will raise an *RuntimeError*.

To avoid numerical precision problems, this method actually gently clips values which barely exceed the boundaries. This is because variables are internally scaled and thus may very slightly violate the boundaries after being unscaled. If a parameter's value only slightly exceeds the boundaries, then its value will be set equal to the boundary. If a value exceeds the boundary by a nontrivial amount, then a *RuntimeError* will be raised.

Note that although this method is forgiving about input parameter values very slightly exceeding the boundaries, it is completely strict about returning parameter values which obey the boundaries. Thus it is safe to assume that the returned values are within the range set by *self.min_boundary* and *self.max_boundary* (inclusively).

Parameters *params* (*array*) – Array of values to be experimentally tested.

Raises *RuntimeError* – A *RuntimeError* is raised if any value in *params* exceeds the parameter value boundaries by a nontrivial amount.

Returns

The input *params*, except that any values which slightly exceed the boundaries will have been clipped to stay within the boundaries exactly.

Return type array

`_first_params` ()

Checks queue to get the first parameters.

Returns Parameters for first experiment

`_get_cost_and_in_dict` ()

Get cost, uncertainty, parameters, bad and extra data from experiment.

This method stores results in lists and also puts data in the appropriate ‘current’ variables. This method doesn’t return anything and instead stores all of its results in the internal storage arrays and the ‘current’ variables.

If the interface encounters an error, it will pass the error to the controller here so that the error can be re-raised in the controller’s thread (note that the interface runs in a separate thread).

`_next_params()`

Send latest cost info and get next parameters from the learner.

`_optimization_routine()`

Runs controller main loop. Gives parameters to experiment and saves costs returned.

`_put_params_and_out_dict(params, param_type=None, **kwargs)`

Send parameters to queue with optional additional keyword arguments.

This method also saves sent variables in appropriate storage arrays.

Parameters

- **`params`** (*array*) – Array of values to be experimentally tested.
- **`param_type`** (*Optional, str*) – The learner type which generated the parameter values. Because some learners use other learners as trainers, the parameter type can be different for different iterations during a given optimization. This value will be stored in *self.out_type* and in the *out_type* list in the controller archive. If *None*, then it will be set to *self.learner.OUT_TYPE*. Default *None*.

Keyword Arguments `kwargs`** – Any additional keyword arguments will be stored in *self.out_extras* and in the *out_extras* list in the controller archive.

`_send_to_learner()`

Send the latest cost info the the learner.

`_shut_down()`

Shutdown and clean up resources of the controller. end the learners, queue_listener and make one last save of archive.

`_start_up()`

Start the learner and interface threads/processes.

`_update_controller_with_learner_attributes()`

Update the controller with properties from the learner.

`check_end_conditions()`

Check whether any of the end contions have been met.

In particular this method check for any of the following conditions:

- If the number of iterations has reached *max_num_runs*.
- If the *target_cost* has been reached.
- If *max_num_runs_without_better_params* iterations in a row have occurred without any improvement.
- If *max_duration* seconds or more has elapsed since *start_datetime*.

Returns

***True*, if the controller should continue or *False* if one or more halting conditions have been met and the controller should end.**

Return type bool

optimize()

Optimize the experiment. This code learner and interface processes/threads are launched and appropriately ended. Starts both threads and catches kill signals and shuts down appropriately.

print_results()

Print results from optimization run to the logs

save_archive()

Save the archive associated with the controller class. Only occurs if the filename for the archive is not None. Saves with the format previously set.

exception `mloop.controllers.ControllerInterrupt`

Bases: `Exception`

Exception that is raised when the controlled is ended with the end flag or event.

class `mloop.controllers.DifferentialEvolutionController` (*interface*, ***kwargs*)

Bases: `mloop.controllers.Controller`

Controller for the differential evolution learner. :param `params_out_queue`: Queue for parameters to next be run by experiment. :type `params_out_queue`: queue :param `costs_in_queue`: Queue for costs (and other details) that have been returned by experiment. :type `costs_in_queue`: queue :param ***kwargs*: Dictionary of options to be passed to Controller parent class and differential evolution learner. :type ***kwargs*: Optional [dict]

class `mloop.controllers.GaussianProcessController` (*interface*, *num_params=None*,
min_boundary=None,
max_boundary=None,
trust_region=None,
learner_archive_filename='learner_archive',
learner_archive_file_type='txt',
param_names=None, ***kwargs*)

Bases: `mloop.controllers.MachineLearnerController`

Controller for the Gaussian Process solver. Primarily suggests new points from the Gaussian Process learner. However, during the initial few runs it must rely on a different optimization algorithm to get some points to seed the learner. :param `interface`: The interface to the experiment under optimization. :type `interface`: Interface :param ***kwargs*: Dictionary of options to be passed to MachineLearnerController parent class and Gaussian Process learner. :type ***kwargs*: Optional [dict]

Keyword Args:

class `mloop.controllers.MachineLearnerController` (*interface*, *training_type='differential_evolution'*,
num_training_runs=None,
no_delay=True, *num_params=None*,
min_boundary=None,
max_boundary=None,
trust_region=None,
learner_archive_filename='learner_archive',
learner_archive_file_type='txt',
param_names=None, ***kwargs*)

Bases: `mloop.controllers.Controller`

Abstract Controller class for the machine learning based solvers. :param `interface`: The interface to the experiment under optimization. :type `interface`: Interface :param ***kwargs*: Dictionary of options to be passed to Controller parent class and initial training learner. :type ***kwargs*: Optional [dict]

Keyword Arguments

- **training_type** (*Optional [string]*) – The type for the initial training source can be ‘random’ for the random learner, ‘nelder_mead’ for the Nelder–Mead learner or

‘differential_evolution’ for the Differential Evolution learner. This learner is also called if the machine learning learner is too slow and a new point is needed. Default ‘differential_evolution’.

- **num_training_runs** (*Optional [int]*) – The number of training runs to before starting the learner. If None, will be ten or double the number of parameters, whatever is larger.
- **no_delay** (*Optional [bool]*) – If True, there is never any delay between a returned cost and the next parameters to run for the experiment. In practice, this means if the machine learning learner has not prepared the next parameters in time the learner defined by the initial training source is used instead. If false, the controller will wait for the machine learning learner to predict the next parameters and there may be a delay between runs.

_get_cost_and_in_dict()

Get cost, uncertainty, parameters, bad, and extra data from experiment.

This method calls `_get_cost_and_in_dict()` of the parent *Controller* class and additionally sends the results to machine learning learner.

_optimization_routine()

Overrides `_optimization_routine`. Uses the parent routine for the training runs. Implements a customized `_optimization_routine` when running the machine learning learner.

_shut_down()

Shutdown and clean up resources of the machine learning controller.

_start_up()

Runs parent method and also starts `training_learner`.

print_results()

Adds some additional output to the results specific to controller.

class `mloop.controllers.NelderMeadController` (*interface*, ***kwargs*)

Bases: `mloop.controllers.Controller`

Controller for the Nelder–Mead solver. Suggests new parameters based on the Nelder–Mead algorithm. Can take no boundaries or hard boundaries. More details for the Nelder–Mead options are in the learners section.
 :param `params_out_queue`: Queue for parameters to next be run by experiment. :type `params_out_queue`: queue
 :param `costs_in_queue`: Queue for costs (and other details) that have been returned by experiment. :type `costs_in_queue`: queue
 :param `**kwargs`: Dictionary of options to be passed to Controller parent class and Nelder–Mead learner. :type `**kwargs`: Optional [dict]

class `mloop.controllers.NeuralNetController` (*interface*, *num_params=None*, *min_boundary=None*, *max_boundary=None*, *trust_region=None*, *learner_archive_filename='learner_archive'*, *learner_archive_file_type='txt'*, *param_names=None*, ***kwargs*)

Bases: `mloop.controllers.MachineLearnerController`

Controller for the Neural Net solver. Primarily suggests new points from the Neural Net learner. However, during the initial few runs it must rely on a different optimization algorithm to get some points to seed the learner. :param `interface`: The interface to the experiment under optimization. :type `interface`: Interface :param `**kwargs`: Dictionary of options to be passed to MachineLearnerController parent class and Neural Net learner. :type `**kwargs`: Optional [dict]

Keyword Args:

class `mloop.controllers.RandomController` (*interface*, ***kwargs*)

Bases: `mloop.controllers.Controller`

Controller that simply returns random variables for the next parameters. Costs are stored but do not influence future points picked. :param params_out_queue: Queue for parameters to next be run by experiment. :type params_out_queue: queue :param costs_in_queue: Queue for costs (and other details) that have been returned by experiment. :type costs_in_queue: queue :param **kwargs: Dictionary of options to be passed to Controller and Random Learner. :type **kwargs: Optional [dict]

```
mloop.controllers.create_controller(interface, controller_type='gaussian_process', **con-
                                   troller_config_dict)
```

Start the controller with the options provided. :param interface: Interface with queues and events to be passed to controller :type interface: interface

Keyword Arguments

- **controller_type** (*Optional [str]*) – Defines the type of controller can be 'random', 'nelder', 'gaussian_process' or 'neural_net'. Alternatively, the controller can belong to an external module, in which case this parameter should be 'module_name:controller_name'. Defaults to 'gaussian_process'.
- ****controller_config_dict** – Options to be passed to controller.

Returns threadable object which must be started with start() to get the controller running.

Return type *Controller*

Raises ValueError – if controller_type is an unrecognized string

2.9.3 interfaces

Module of the interfaces used to connect the controller to the experiment.

```
class mloop.interfaces.FileInterface(interface_out_filename='exp_input',           in-
                                     interface_in_filename='exp_output',         inter-
                                     face_file_type='txt', **kwargs)
```

Bases: *mloop.interfaces.Interface*

Interfaces between the files produced by the experiment and the queues accessed by the controllers.

Parameters

- **params_out_queue** (*queue*) – Queue for parameters to next be run by experiment.
- **costs_in_queue** (*queue*) – Queue for costs (and other details) that have been returned by experiment.

Keyword Arguments

- **interface_out_filename** (*Optional [string]*) – filename for file written with parameters.
- **interface_in_filename** (*Optional [string]*) – filename for file written with parameters.
- **interface_file_type** (*Optional [string]*) – file type to be written either 'mat' for matlab or 'txt' for readable text file. Defaults to 'txt'.

```
get_next_cost_dict(params_dict)
```

Implementation of file read in and out. Put parameters into a file and wait for a cost file to be returned.

```
class mloop.interfaces.Interface(interface_wait=1, **kwargs)
```

Bases: *threading.Thread*

A abstract class for interfaces which populate the costs_in_queue and read from the params_out_queue. Inherits from Thread

Parameters

- **interface_wait** (*Optional [float]*) – Time between polling when needed in interface.
- **params_out_queue** (*queue*) – Queue for parameters to next be run by experiment.
- **costs_in_queue** (*queue*) – Queue for costs (and other details) that have been returned by experiment.
- **end_event** (*event*) – Event which triggers the end of the interface.

Keyword Arguments **interface_wait** (*float*) – Wait time when polling for files or queues is needed.

get_next_cost_dict (*params_dict*)

Abstract method.

This is the only method that needs to be implemented to make a working interface.

Given the parameters the interface must then produce a new cost. This may occur by running an experiment or program. If an error is raised by this method, the optimization will halt.

Parameters **params_dict** (*dictionary*) – A dictionary containing the parameters. Use *params_dict['params']* to access them.

Returns

The cost and other properties derived from the experiment when it was run with the parameters. If just a cost was produced provide *{'cost': [float]}*, if you also have an uncertainty provide *{'cost': [float], 'uncer': [float]}*. If the run was bad you can simply provide *{'bad': True}*. For completeness you can always provide all three using *{'cost': [float], 'uncer': [float], 'bad': [bool]}*. Any extra keys provided will also be saved by the controller.

Return type *cost_dict* (dictionary)

run ()

The run sequence for the interface.

This method does NOT need to be overloaded create a working interface.

exception *mloop.interfaces.InterfaceInterrupt*

Bases: *Exception*

The *InterfaceInterrupt* is now deprecated.

The *Interface* class can now handle arbitrary errors, so there is no need for *Interface.get_next_cost_dict()* to raise an *InterfaceInterrupt* in particular. Instead raise an appropriate error given the situation.

class *mloop.interfaces.ShellInterface* (*command='./run_exp', params_args_type='direct', param_names=None, **kwargs*)

Bases: *mloop.interfaces.Interface*

Interface for running programs from the shell.

Parameters

- **params_out_queue** (*queue*) – Queue for parameters to next be run by experiment.
- **costs_in_queue** (*queue*) – Queue for costs (and other details) that have been returned by experiment.

Keyword Arguments

- **command** (*Optional [string]*) – The command used to run the experiment. Default `./run_exp`.
- **params_args_type** (*Optional [string]*) – The style used to pass parameters. Can be `direct` or `named`. If `direct` it is assumed the parameters are fed directly to the program. For example if I wanted to run the parameters `[7,5,9]` with the command `./run_exp` I would use the syntax:

```
./run_exp 7 5 9
```

`named` on the other hand requires an option for each parameter. The options should be name `-param1`, `-param2` etc (unless `param_names` is specified, see below)). The same example as before would be

```
./run_exp --param1 7 --param2 5 --param3 9
```

Default `direct`.

- **param_names** (*Optional [string]*) – List of names for parameters to be passed as options to the shell command, replacing `-param1`, `-param2`, etc. If set to `None` then the default parameter names will be used. Default `None`.

get_next_cost_dict (*params_dict*)

Implementation of running a command with parameters on the command line and reading the result.

class `mloop.interfaces.TestInterface` (*test_landscape=None, **kwargs*)

Bases: `mloop.interfaces.Interface`

Interface for testing. Returns fake landscape data directly to learner.

Parameters

- **params_out_queue** (*queue*) – Parameters to be used to evaluate fake landscape.
- **costs_in_queue** (*queue*) – Queue for costs (and other details) that have been calculated from fake landscape.

Keyword Arguments

- **test_landscape** (*Optional [TestLandscape]*) – Landscape that can be given a set of parameters and a cost and other values. If `None` creates a the default landscape. Default `None`
- **out_queue_wait** (*Optional [float]*) – Time in seconds to wait for queue before checking end flag.

get_next_cost_dict (*params_dict*)

Test implementation. Gets the next cost from the `test_landscape`.

`mloop.interfaces.create_interface` (*interface_type='file', **interface_config_dict*)

Start a new interface with the options provided.

Parameters

- **interface_type** (*Optional [str]*) – Defines the type of interface, can be `file`, `shell` or `test`. Default `file`.
- ****interface_config_dict** – Options to be passed to interface.

Returns An interface as defined by the keywords

Return type interface

2.9.4 launchers

Modules of launchers used to start M-LOOP.

`mloop.launchers._pop_extras_kwargs` (*kwargs*)

Remove the keywords used in the extras section (if present), and return them.

Returns tuple made of (*extras_kwargs*, *kwargs*), where *extras_kwargs* are keywords for the extras and *kwargs* are the others that were provided.

`mloop.launchers.launch_extras` (*controller*, *visualizations=True*, ***kwargs*)

Launch post optimization extras. Including visualizations.

Keyword Arguments *visualizations* (*Optional [bool]*) – If true run default visualizations for the controller. Default false.

`mloop.launchers.launch_from_file` (*config_filename*, ***kwargs*)

Launch M-LOOP using a configuration file. See configuration file documentation.

Parameters

- **config_filename** (*str*) – Filename of configuration file
- ****kwargs** – keywords that override the keywords in the file.

Returns Controller for optimization.

Return type controller (*Controller*)

2.9.5 learners

Module of learners used to determine what parameters to try next given previous cost evaluations.

Each learner is created and controlled by a controller.

```
class mloop.learners.DifferentialEvolutionLearner (first_params=None,
                                                    trust_region=None,          evolu-
                                                    tion_strategy='best1',      popula-
                                                    tion_size=15, mutation_scale=(0.5,
                                                    1),      cross_over_probability=0.7,
                                                    restart_tolerance=0.01, **kwargs)
```

Bases: `mloop.learners.Learner`, `threading.Thread`

Adaption of the differential evolution algorithm in scipy.

Parameters

- **params_out_queue** (*queue*) – Queue for parameters sent to controller.
- **costs_in_queue** (*queue*) – Queue for costs for gaussian process. This must be tuple
- **end_event** (*event*) – Event to trigger end of learner.

Keyword Arguments

- **first_params** (*Optional [array]*) – The first parameters to test. If None will just randomly sample the initial condition. Default None.
- **trust_region** (*Optional [float or array]*) – The trust region defines the maximum distance the learner will travel from the current best set of parameters. If None, the learner will search everywhere. If a float, this number must be between 0 and 1 and defines maximum distance the learner will venture as a percentage of the boundaries. If it is

an array, it must have the same size as the number of parameters and the numbers define the maximum absolute distance that can be moved along each direction.

- **evolution_strategy** (*Optional [string]*) – the differential evolution strategy to use, options are 'best1', 'best2', 'rand1' and 'rand2'. The default is 'best1'.
- **population_size** (*Optional [int]*) – multiplier proportional to the number of parameters in a generation. The generation population is set to `population_size * parameter_num`. Default 15.
- **mutation_scale** (*Optional [tuple]*) – The mutation scale when picking new points. Otherwise known as differential weight. When provided as a tuple (min,max) a mutation constant is picked randomly in the interval. Default (0.5,1.0).
- **cross_over_probability** (*Optional [float]*) – The recombination constant or crossover probability, the probability a new points will be added to the population.
- **restart_tolerance** (*Optional [float]*) – when the current population have a spread less than the initial tolerance, namely `stdev(curr_pop) < restart_tolerance stdev(init_pop)`, it is likely the population is now in a minima, and so the search is started again.

has_trust_region

Whether the learner has a trust region.

Type bool

num_population_members

The number of parameters in a generation.

Type int

params_generations

History of the parameters generations. A list of all the parameters in the population, for each generation created.

Type list

costs_generations

History of the costs generations. A list of all the costs in the population, for each generation created.

Type list

init_std

The initial standard deviation in costs of the population. Calculated after sampling (or resampling) the initial population.

Type float

curr_std

The current standard deviation in costs of the population. Calculated after sampling each generation.

Type float

OUT_TYPE = 'differential_evolution'**_best1** (*index*)

Use best parameters and two others to generate mutation.

Parameters **index** (*int*) – Index of member to mutate.

_best2 (*index*)

Use best parameters and four others to generate mutation.

Parameters **index** (*int*) – Index of member to mutate.

_rand1 (*index*)

Use three random parameters to generate mutation.

Parameters **index** (*int*) – Index of member to mutate.

_rand2 (*index*)

Use five random parameters to generate mutation.

Parameters **index** (*int*) – Index of member to mutate.

generate_population ()

Sample a new random set of variables

mutate (*index*)

Mutate the parameters at index.

Parameters **index** (*int*) – Index of the point to be mutated.

next_generation ()

Evolve the population by a single generation

random_index_sample (*index*, *num_picks*)

Randomly select a num_picks of indexes, without index.

Parameters

- **index** (*int*) – The index that is not included
- **num_picks** (*int*) – The number of picks.

run ()

Runs the Differential Evolution Learner.

save_generation ()

Save history of generations.

update_archive ()

Update the archive.

```
class mloop.learners.GaussianProcessLearner (length_scale=None,
                                             length_scale_bounds=None,          up-
                                             date_hyperparameters=True,
                                             cost_has_noise=True,   noise_level=None,
                                             noise_level_bounds=None, **kwargs)
```

Bases: [mloop.learners.MachineLearner](#), `multiprocessing.context.Process`

Gaussian process learner.

Generates new parameters based on a gaussian process fitted to all previous data.

Parameters

- **params_out_queue** (*queue*) – Queue for parameters sent to controller.
- **costs_in_queue** (*queue*) – Queue for costs for gaussian process. This must be tuple.
- **end_event** (*event*) – Event to trigger end of learner.

Keyword Arguments

- **length_scale** (*Optional [array]*) – The initial guess for length scale(s) of the gaussian process. The array can either of size one or the number of parameters or *None*. If it is size one, it is assumed that all of the correlation lengths are the same. If it is an array with length equal to the number of the parameters then all the parameters have their own independent length scale. If it is set to *None* and a learner archive from a Gaussian process optimization is provided for *training_filename*, then it will be set to the value recorded for

length_scale in that learner archive. If set to *None* but *training_filename* does not specify a learner archive from a Gaussian process optimization, then it is assumed that all of the length scales should be independent and they are all given an initial value of equal to one tenth of their allowed range. Default *None*.

- **length_scale_bounds** (*Optional [array]*) – The limits on the fitted length scale values, specified as a single pair of numbers e.g. *[min, max]*, or a list of pairs of numbers, e.g. *[[min_0, max_0], ..., [min_N, max_N]]*. This only has an effect if *update_hyperparameters* is set to *True*. If one pair is provided, the same limits will be used for all length scales. Alternatively one pair of *[min, max]* can be provided for each length scale. For example, possible valid values include *[1e-5, 1e5]* and *[[1e-2, 1e2], [5, 5], [1.6e-4, 1e3]]* for optimizations with three parameters. If set to *None*, then the length scale will be bounded to be between 0.001 and 10 times the allowed range for each parameter.
- **update_hyperparameters** (*Optional [bool]*) – Whether the length scales and noise estimate should be updated when new data is provided. Default *True*.
- **cost_has_noise** (*Optional [bool]*) – If *True* the learner assumes there is common additive white noise that corrupts the costs provided. This noise is assumed to be on top of the uncertainty in the costs (if it is provided). If *False*, it is assumed that there is no noise in the cost (or if uncertainties are provided no extra noise beyond the uncertainty). Default *True*.
- **noise_level** (*Optional [float]*) – The initial guess for the noise level (variance, not standard deviation) in the costs. This is only used if *cost_has_noise* is *True*. If it is set to *None* and a learner archive from a Gaussian process optimization is provided for *training_filename*, then it will be set to the value recorded for *noise_level* in that learner archive. If set to *None* but *training_filename* does not specify a learner archive from a Gaussian process optimization, then it will automatically be set to the variance of the training data costs.
- **noise_level_bounds** (*Optional [array]*) – The limits on the fitted *noise_level* values, specified as a single pair of numbers *[min, max]*. This only has an effect if *update_hyperparameters* and *cost_has_noise* are both set to *True*. If set to *None*, the value *[1e-5 * var, 1e5 * var]* will be used where *var* is the variance of the training data costs. Default *None*.
- **training_filename** (*Optional [str]*) – The name of a learner archive from a previous optimization from which to extract past results for use in the current optimization. If *None*, no past results will be used. Default *None*.
- **trust_region** (*Optional [float or array]*) – The trust region defines the maximum distance the learner will travel from the current best set of parameters. If *None*, the learner will search everywhere. If a float, this number must be between 0 and 1 and defines maximum distance the learner will venture as a percentage of the boundaries. If it is an array, it must have the same size as the number of parameters and the numbers define the maximum absolute distance that can be moved along each direction.
- **default_bad_cost** (*Optional [float]*) – If a run is reported as bad and *default_bad_cost* is provided, the cost for the bad run is set to this default value. If *default_bad_cost* is *None*, then the worst cost received is set to all the bad runs. Default *None*.
- **default_bad_uncertainty** (*Optional [float]*) – If a run is reported as bad and *default_bad_uncertainty* is provided, the uncertainty for the bad run is set to this default value. If *default_bad_uncertainty* is *None*, then the uncertainty is set to a tenth of the best to worst cost range. Default *None*.

- **minimum_uncertainty** (*Optional [float]*) – The minimum uncertainty associated with provided costs. Must be above zero to avoid fitting errors. Default *1e-8*.
- **predict_global_minima_at_end** (*Optional [bool]*) – If *True* attempts to find the global minima when the learner is ended. Does not if *False*. Default *True*.

all_params

Array containing all parameters sent to learner.

Type array

all_costs

Array containing all costs sent to learner.

Type array

all_uncers

Array containing all uncertainties sent to learner.

Type array

scaled_costs

Array containing all the costs scaled to have zero mean and a standard deviation of 1. Needed for training the gaussian process.

Type array

bad_run_indexes

list of indexes to all runs that were marked as bad.

Type list

best_cost

Minimum received cost, updated during execution.

Type float

best_params

Parameters of best run. (reference to element in params array).

Type array

best_index

index of the best cost and params.

Type int

worst_cost

Maximum received cost, updated during execution.

Type float

worst_index

index to run with worst cost.

Type int

cost_range

Difference between *worst_cost* and *best_cost*.

Type float

generation_num

Number of sets of parameters to generate each generation. Set to 4.

Type int

length_scale_history

List of length scales found after each fit.

Type list

noise_level_history

List of noise levels found after each fit.

Type list

fit_count

Counter for the number of times the gaussian process has been fit.

Type int

cost_count

Counter for the number of costs, parameters and uncertainties added to learner.

Type int

params_count

Counter for the number of parameters asked to be evaluated by the learner.

Type int

gaussian_process

Gaussian process that is fitted to data and used to make predictions

Type GaussianProcessRegressor

cost_scaler

Scaler used to normalize the provided costs.

Type StandardScaler

params_scaler

Scaler used to normalize the provided parameters.

Type StandardScaler

has_trust_region

Whether the learner has a trust region.

Type bool

OUT_TYPE = 'gaussian_process'

_check_length_scale_bounds()

Ensure self.length_scale_bounds has a valid value, otherwise raise a ValueError.

_check_noise_level_bounds()

Ensure self.noise_level has a valid value, otherwise raise a ValueError.

_transform_length_scale_bounds(length_scale_bounds, inverse=False)

Transform length scale bounds to or from scaled units.

This method functions similarly to *self.transform_length_scales()*, except that it transforms the bounds for the length scales. The same scalings used for the length scales themselves are applied here to the lower and upper bounds. To transform from real/unscaled units to scaled units, call this method with *inverse* set to *False*. To perform the inverse transformation, namely to transform length scale bounds from scaled units to real/unscaled units, call this method with *inverse* set to *True*.

The output array will have a separate scaled min/max value pair for each parameter length scale. In other words, the output will be an array with two columns (one for min values and one for max values) and one

row for each parameter length scale. This will be the case even if *length_scale_bounds* consists of a single min/max value pair because the scalings are generally different for different parameters.

Note that although *length_scale_bounds* can be a 1D array with only two entries (a single min/max pair shared by all parameters), this method always returns a 2D array with a separate min/max pair for each parameter because the scaling factors aren't generally the same for all of the parameters. This implies that transforming a 1D array then performing the inverse transformation will yield a 2D array of identical min/max pairs rather than the original 1D array.

Parameters

- **length_scale_bounds** (*array*) – The bounds for the Gaussian process's length scales which should be transformed to or from scaled units. This can either be (a) a 1D array with two entries of the form *[min, max]* or (b) a 2D array with two columns (min and max values respectively) and one row for each parameter length scale.
- **inverse** (*bool*) – This argument controls whether the forward or inverse transformation is applied. If *False*, then the forward transformation is applied, which takes *length_scale_bounds* in real/unscaled units and transforms them to scaled units. If *True* then this method assumes that *length_scale_bounds* are in scaled units and transforms them into real/unscaled units. Default *False*.

Raises `ValueError` – A `ValueError` is raised if *length_scale_bounds* does not have an acceptable shape. The allowed shapes are (2,) (a single min/max pair shared by all parameters) or (*self.num_params*, 2) (a separate min/max pair for each parameter).

Returns

The transformed length scale bounds. These will be in scaled units if *inverse* is *False* or in real/unscaled units if *inverse* is *True*. Note that *transformed_length_scale_bounds* will always be a 2D array of shape (*self.num_params*, 2) even if *length_scale_bounds* was a single pair of min/max values.

Return type *transformed_length_scale_bounds* (*array*)

`_transform_length_scales` (*length_scales*, *inverse=False*)

Transform length scales to or from scaled units.

This method uses *self.params_scaler* to transform length scales to/from scaled units. To transform from real/unscaled units to scaled units, call this method with *inverse* set to *False*. To perform the inverse transformation, namely to transform length scales from scaled units to real/unscaled units, call this method with *inverse* set to *True*.

Notably length scales should be scaled, but not offset, when they are transformed. For this reason, they should not simply be passed through *self.params_scaler.transform()* and instead should be passed through this method.

Although *length_scales* can be a single float, this method always returns a 1D array because the scaling factors aren't generally the same for all of the parameters. This implies that transforming a float then performing the inverse transformation will yield a 1D array of identical entries rather than a single float.

Parameters

- **length_scales** (*float or array*) – Length scale(s) for the Gaussian process which should be transformed to or from scaled units. Can be either a single float or a 1D array of length *self.num_params*.
- **inverse** (*bool*) – This argument controls whether the forward or inverse transformation is applied. If *False*, then the forward transformation is applied, which takes *length_scales* in real/unscaled units and transforms them to scaled units. If *True* then this method as-

sumes that *length_scales* are in scaled units and transforms them into real/unscaled units. Default *False*.

Returns

The transformed length scales. These will be in scaled units if *inverse* is *False* or in real/unscaled units if *inverse* is *True*. Note that *transformed_length_scales* will be a 1D array even if *length_scales* was a single float.

Return type *transformed_length_scales* (array)

create_gaussian_process ()

Create a Gaussian process.

find_global_minima ()

Search for the global minima predicted by the Gaussian process.

This method will attempt to find the global minima predicted by the Gaussian process, but it is possible for it to become stuck in local minima of the predicted cost landscape.

This method does not return any values, but creates the attributes listed below.

predicted_best_parameters

The parameter values which are predicted to yield the best results, as a 1D array.

Type array

predicted_best_cost

The predicted cost at the *predicted_best_parameters*, in a 1D 1-element array.

Type array

predicted_best_uncertainty

The uncertainty of the predicted cost at *predicted_best_parameters*, in a 1D 1-element array.

Type array

find_next_parameters ()

Get the next parameters to test.

This method searches for the parameters expected to give the minimum biased cost, as predicted by the Gaussian process. The biased cost is not just the predicted cost, but a weighted sum of the predicted cost and the uncertainty in the predicted cost. See *self.predict_biased_cost()* for more information.

This method additionally increments *self.params_count* appropriately.

Returns

The next parameter values to try, stored in a 1D array.

Return type *next_params* (array)

fit_gaussian_process ()

Fit the Gaussian process to the current data

predict_biased_cost (*params*, *perform_scaling=True*)

Predict the biased cost at the given parameters.

The biased cost is a weighted sum of the predicted cost and the uncertainty of the predicted cost. In particular, the bias function is:

$$biased_cost = cost_bias * pred_cost - uncer_bias * pred_uncer$$

Parameters

- **params** (*array*) – A 1D array containing the values for each parameter. These should be in real/unscaled units if *perform_scaling* is *True* or they should be in scaled units if *perform_scaling* is *False*.
- **perform_scaling** (*bool, optional*) – Whether or not the parameters and biased costs should be scaled. If *True* then this method takes in parameter values in real/unscaled units then returns a biased predicted cost in real/unscaled units. If *False*, then this method takes parameter values in scaled units and returns a biased predicted cost in scaled units. Note that this method cannot determine on its own if the values in *params* are in real/unscaled units or scaled units; it is up to the caller to pass the correct values. Defaults to *True*.

Returns

Biased cost predicted for the given parameters. This will be in real/unscaled units if *perform_scaling* is *True* or it will be in scaled units if *perform_scaling* is *False*.

Return type pred_bias_cost (float)

predict_cost (*params, perform_scaling=True, return_uncertainty=False*)

Predict the cost for *params* using *self.gaussian_process*.

This method also optionally returns the uncertainty of the predicted cost.

By default (with *perform_scaling=True*) this method will use *self.params_scaler* to scale the input values and then use *self.cost_scaler* to scale the cost back to real/unscaled units. If *perform_scaling* is *False*, then this scaling will NOT be done. In that case, *params* should consist of already-scaled parameter values and the returned cost (and optional uncertainty) will be in scaled units.

Parameters

- **params** (*array*) – A 1D array containing the values for each parameter. These should be in real/unscaled units if *perform_scaling* is *True* or they should be in scaled units if *perform_scaling* is *False*.
- **perform_scaling** (*bool, optional*) – Whether or not the parameters and costs should be scaled. If *True* then this method takes in parameter values in real/unscaled units then returns a predicted cost (and optionally the predicted cost uncertainty) in real/unscaled units. If *False*, then this method takes parameter values in scaled units and returns a cost (and optionally the predicted cost uncertainty) in scaled units. Note that this method cannot determine on its own if the values in *params* are in real/unscaled units or scaled units; it is up to the caller to pass the correct values. Defaults to *True*.
- **return_uncertainty** (*bool, optional*) – This optional argument controls whether or not the predicted cost uncertainty is returned with the predicted cost. The predicted cost uncertainty will be in real/unscaled units if *perform_scaling* is *True* and will be in scaled units if *perform_scaling* is *False*. Defaults to *False*.

Returns

Predicted cost at *params*. The cost will be in real/unscaled units if *perform_scaling* is *True* and will be in scaled units if *perform_scaling* is *False*.

uncertainty (float, optional): The uncertainty of the predicted cost. This will be in the same units (either real/unscaled or scaled) as the returned *cost*. The *cost_uncertainty* will only be returned if *return_uncertainty* is *True*.

Return type cost (float)

run ()

Starts running the Gaussian process learner. When the new parameters event is triggered, reads the cost

information provided and updates the Gaussian process with the information. Then searches the Gaussian process for new optimal parameters to test based on the biased cost. Parameters to test next are put on the output parameters queue.

update_archive()

Update the archive.

update_bias_function()

Set the constants for the cost bias function.

```
class mloop.learners.Learner(num_params=None, min_boundary=None, max_boundary=None,
                             learner_archive_filename='learner_archive',
                             learner_archive_file_type='txt',           start_datetime=None,
                             param_names=None, **kwargs)
```

Bases: object

Base class for all learners. Contains default boundaries and some useful functions that all learners use.

The class that inherits from this class should also inherit from `threading.Thread` or `multiprocessing.Process`, depending if you need the learner to be a genuine parallel process or not.

Keyword Arguments

- **num_params** (*Optional [int]*) – The number of parameters to be optimized. If None defaults to 1. Default None.
- **min_boundary** (*Optional [array]*) – Array with minimum values allowed for each parameter. Note if certain values have no minimum value you can set them to -inf for example `[-1, 2, float('-inf')]` is a valid `min_boundary`. If None sets all the boundaries to '-1'. Default None.
- **max_boundary** (*Optional [array]*) – Array with maximum values allowed for each parameter. Note if certain values have no maximum value you can set them to +inf for example `[0, float('inf'), 3, -12]` is a valid `max_boundary`. If None sets all the boundaries to '1'. Default None.
- **learner_archive_filename** (*Optional [string]*) – Name for python archive of the learners current state. If None, no archive is saved. Default None. But this is typically overloaded by the child class.
- **learner_archive_file_type** (*Optional [string]*) – File type for archive. Can be either 'txt' a human readable text file, 'pkl' a python dill file, 'mat' a matlab file or None if there is no archive. Default 'mat'.
- **log_level** (*Optional [int]*) – Level for the learners logger. If None, set to warning. Default None.
- **start_datetime** (*Optional [datetime]*) – Start date time, if None, is automatically generated.
- **param_names** (*Optional [list of str]*) – A list of names of the parameters for use e.g. in plot legends. Number of elements must equal `num_params`. If None, each name will be set to an empty string. Default None.

params_out_queue

Queue for parameters created by learner.

Type queue

costs_in_queue

Queue for costs to be used by learner.

Type queue

end_event

Event to trigger end of learner.

Type event

all_params

Array containing all parameters sent to learner.

Type array

all_costs

Array containing all costs sent to learner.

Type array

all_uncers

Array containing all uncertainties sent to learner.

Type array

bad_run_indexes

list of indexes to all runs that were marked as bad.

Type list

OUT_TYPE = ''**_parse_cost_message** (*message*)

Parse a message sent from the controller via *self.costs_in_queue*.

Parameters *message* (*tuple*) – A tuple put in *self.costs_in_queue* by the controller. It should be of the form (*params*, *cost*, *uncer*, *bad*) where *params* is an array specifying the parameter values used, *cost* is the measured cost for those parameter values, *uncer* is the uncertainty measured for those parameter values, and *bad* is a boolean indicating whether the run was bad.

Raises `ValueError` – A `ValueError` is raised if the number of parameters in the provided *params* doesn't match *self.num_params*.

Returns

A tuple of the form (*params*, *cost*, *uncer*, *bad*). For more information on the meaning of those parameters, see the entry for the *message* argument above.

Return type tuple

_prepare_logger ()

Prepare the logger.

If *self.log* already exists, then this method silently returns without changing anything.

_set_trust_region (*trust_region*)

Sets trust region properties for learner that have this. Common function for learners with trust regions.

Parameters *trust_region* (*float* or *array*) – Property defines the trust region.

_shut_down ()

Shut down and perform one final save of learner.

_update_run_data_attributes (*params*, *cost*, *uncer*, *bad*)

Update attributes that store the results returned by the controller.

Parameters

- **params** (*array*) – Array of control parameter values.

- **cost** (*float*) – The cost measured for *params*.
- **uncer** (*float*) – The uncertainty measured for *params*.
- **bad** (*bool*) – Whether or not the run was bad.

check_in_boundary (*param*)

Check given parameters are within stored boundaries.

Parameters **param** (*array*) – array of parameters

Returns True if the parameters are within boundaries, False otherwise.

Return type bool

check_in_diff_boundary (*param*)

Check given distances are less than the boundaries.

Parameters **param** (*array*) – array of distances

Returns True if the distances are smaller or equal to boundaries, False otherwise.

Return type bool

check_num_params (*param*)

Check the number of parameters is right.

put_params_and_get_cost (*params*, ***kwargs*)

Send parameters to queue and whatever additional keywords.

Also saves sent and received variables in appropriate storage arrays.

Parameters **params** (*array*) – array of values to be sent to file

Returns cost from the cost queue

save_archive ()

Save the archive associated with the learner class. Only occurs if the filename for the archive is not None.

Saves with the format previously set.

update_archive ()

Update the dictionary of parameters and values to save to the archive.

Child classes should call this method and also updated *self.archive_dict* with any other parameters and values that need to be saved to the learner archive.

exception `mloop.learners.LearnerInterrupt`

Bases: `Exception`

Exception that is raised when the learner is ended with the end flag or event.

class `mloop.learners.MachineLearner` (*trust_region=None*, *default_bad_cost=None*, *default_bad_uncertainty=None*, *minimum_uncertainty=1e-08*, *predict_global_minima_at_end=True*, *training_filename=None*, ***kwargs*)

Bases: `mloop.learners.Learner`

A parent class for more specific machine learner classes.

This class is not intended to be used directly.

Keyword Arguments

- **trust_region** (*Optional [float or array]*) – The trust region defines the maximum distance the learner will travel from the current best set of parameters. If None, the learner will search everywhere. If a float, this number must be between 0 and 1 and defines maximum distance the learner will venture as a percentage of the boundaries. If it is

an array, it must have the same size as the number of parameters and the numbers define the maximum absolute distance that can be moved along each direction.

- **default_bad_cost** (*Optional [float]*) – If a run is reported as bad and *default_bad_cost* is provided, the cost for the bad run is set to this default value. If *default_bad_cost* is *None*, then the worst cost received is set to all the bad runs. Default *None*.
- **default_bad_uncertainty** (*Optional [float]*) – If a run is reported as bad and *default_bad_uncertainty* is provided, the uncertainty for the bad run is set to this default value. If *default_bad_uncertainty* is *None*, then the uncertainty is set to a tenth of the best to worst cost range. Default *None*.
- **minimum_uncertainty** (*Optional [float]*) – The minimum uncertainty associated with provided costs. Must be above zero to avoid fitting errors. Default *1e-8*.
- **predict_global_minima_at_end** (*Optional [bool]*) – If *True* finds the global minima when the learner is ended. Does not if *False*. Default *True*.
- **training_filename** (*Optional [str]*) – The name of a learner archive from a previous optimization from which to extract past results for use in the current optimization. If *None*, no past results will be used. Default *None*.

all_params

Array containing all parameters sent to learner.

Type array

all_costs

Array containing all costs sent to learner.

Type array

all_uncers

Array containing all uncertainties sent to learner.

Type array

scaled_costs

Array containing all the costs scaled to have zero mean and a standard deviation of 1. Needed for training the learner.

Type array

bad_run_indexes

list of indexes to all runs that were marked as bad.

Type list

best_cost

Minimum received cost, updated during execution.

Type float

best_params

Parameters of best run. (reference to element in params array).

Type array

best_index

index of the best cost and params.

Type int

worst_cost

Maximum received cost, updated during execution.

Type float

worst_index

index to run with worst cost.

Type int

cost_range

Difference between `worst_cost` and `best_cost`

Type float

params_count

Counter for the number of parameters asked to be evaluated by the learner.

Type int

has_trust_region

Whether the learner has a trust region.

Type bool

`_find_predicted_minimum` (*scaled_figure_of_merit_function*, *scaled_search_region*,
params_scaler, *scaled_jacobian_function*=None)

Find the predicted minimum of *scaled_figure_of_merit_function*().

The search for the minimum is constrained to be within *scaled_search_region*.

The *scaled_figure_of_merit_function*() should take inputs in scaled units and generate outputs in scaled units. This is necessary because *scipy.optimize.minimize*() (which is used internally here) can struggle if the numbers are too small or too large. Using scaled parameters and figures of merit brings the numbers closer to ~1, which can improve the behavior of *scipy.optimize.minimize*().

Parameters

- **`scaled_figure_of_merit_function`** (*function*) – This should be a function which accepts an array of scaled parameter values and returns a predicted figure of merit. Importantly, both the input parameter values and the returned value should be in scaled units.
- **`scaled_search_region`** (*array*) – The scaled parameter-space bounds for the search. The returned minimum position will be constrained to be within this region. The *scaled_search_region* should be a 2D array of shape (*self.num_params*, 2) where the first column specifies lower bounds and the second column specifies upper bounds for each parameter (in scaled units).
- **`params_scaler`** (*mloop.utilities.ParameterScaler*) – A *ParameterScaler* instance for converting parameters to scaled units.
- **`scaled_jacobian_function`** (*function*, *optional*) – An optional function giving the Jacobian of *scaled_figure_of_merit_function*() which will be used by *scipy.optimize.minimize*() if provided. As with *scaled_figure_of_merit_function*(), the *scaled_jacobian_function*() should accept and return values in scaled units. If *None* then no Jacobian will be provided to *scipy.optimize.minimize*(). Defaults to *None*.

Returns

The scaled parameter values which minimize *scaled_figure_of_merit_function*() within *scaled_search_region*. They are provided as a 1D array of values in scaled units.

Return type `best_scaled_params` (array)

`_reconcile_kwarg_and_training_val` (*kwargs_, name, training_value*)

Utility function for comparing values from kwargs to training values.

When a training archive is specified there can be two values specified for some parameters; one from user's config/kwargs and one from the training archive. This function compares the values. If the values are the same then the value is returned, and if they are different a *ValueError* is raised. Care is taken not to raise that error though if one of the values is *None* since that can mean that a value wasn't specified. In that case the other value is returned, or *None* is returned if they are both *None*.

Parameters

- **`kwargs_`** (*[dict]*) – The dictionary of keyword arguments passed to `__init__()`.
- **`name`** (*[str]*) – The name of the parameter.
- **`training_value`** (*[any]*) – The value for the parameter in the training archive.

Raises *ValueError* – A *ValueError* is raised if the value of the parameter in the keyword arguments doesn't match the value from the training archive.

Returns

The value for the parameter, taken from either *kwargs_* or *training_value*, or both if they are the same.

Return type *[any]*

`get_params_and_costs` ()

Get the parameters and costs from the queue and place in their appropriate all_*[type]* arrays.

Also updates bad costs, best parameters, and search boundaries given trust region.

`update_archive` ()

Update the archive.

`update_bads` ()

Best and/or worst costs have changed, update the values associated with bad runs accordingly.

`update_search_params` ()

Update the list of parameters to use for the next search.

`update_search_region` ()

If trust boundaries is not none, updates the search boundaries based on the defined trust region.

`wait_for_new_params_event` ()

Waits for a new parameters event and starts a new parameter generation cycle.

Also checks end event and will break if it is triggered.

```
class mloop.learners.NelderMeadLearner (initial_simplex_corner=None,          ini-
                                     tial_simplex_displacements=None,      ini-
                                     tial_simplex_scale=None, **kwargs)
```

Bases: *mloop.learners.Learner*, *threading.Thread*

Nelder–Mead learner. Executes the Nelder–Mead learner algorithm and stores the needed simplex to estimate the next points.

Parameters

- **`params_out_queue`** (*queue*) – Queue for parameters from controller.
- **`costs_in_queue`** (*queue*) – Queue for costs for nelder learner. The queue should be populated with cost (float) corresponding to the last parameter sent from the Nelder–Mead Learner. Can be a float('inf') if it was a bad run.
- **`end_event`** (*event*) – Event to trigger end of learner.

Keyword Arguments

- **initial_simplex_corner** (*Optional [array]*) – Array for the initial set of parameters, which is the lowest corner of the initial simplex. If None the initial parameters are randomly sampled if the boundary conditions are provided, or all are set to 0 if boundary conditions are not provided.
- **initial_simplex_displacements** (*Optional [array]*) – Array used to construct the initial simplex. Each array is the positive displacement of the parameters above the init_params. If None and there are no boundary conditions, all are set to 1. If None and there are boundary conditions assumes the initial conditions are scaled. Default None.
- **initial_simplex_scale** (*Optional [float]*) – Creates a simplex using a the boundary conditions and the scaling factor provided. If None uses the init_simplex if provided. If None and init_simplex is not provided, but boundary conditions are is set to 0.5. Default None.

init_simplex_corner

Parameters for the corner of the initial simple used.

Type array**init_simplex_disp**

Parameters for the displacements about the simplex corner used to create the initial simple.

Type array**simplex_params**

Parameters of the current simplex

Type array**simplex_costs**

Costs associated with the parameters of the current simplex

Type array**OUT_TYPE = 'nelder_mead'****run()**

Runs Nelder–Mead algorithm to produce new parameters given costs, until end signal is given.

update_archive()

Update the archive.

class mloop.learners.NeuralNetLearner (*update_hyperparameters=False, **kwargs*)Bases: *mloop.learners.MachineLearner*, *multiprocessing.context.Process*

Learner that uses a neural network for function approximation.

Parameters

- **params_out_queue** (*queue*) – Queue for parameters sent to controller.
- **costs_in_queue** (*queue*) – Queue for costs.
- **end_event** (*event*) – Event to trigger end of learner.

Keyword Arguments

- **update_hyperparameters** (*Optional [bool]*) – Whether the hyperparameters used to prevent overfitting should be tuned by trying out different values. Setting to *True* can reduce overfitting of the model, but can slow down the fitting due to the computational cost of trying different values. Default *False*.

- **training_filename** (*Optional [str]*) – The name of a learner archive from a previous optimization from which to extract past results for use in the current optimization. If *None*, no past results will be used. Default *None*.
- **trust_region** (*Optional [float or array]*) – The trust region defines the maximum distance the learner will travel from the current best set of parameters. If *None*, the learner will search everywhere. If a float, this number must be between 0 and 1 and defines maximum distance the learner will venture as a fraction of the boundaries. If it is an array, it must have the same size as the number of parameters and the numbers define the maximum absolute distance that can be moved along each direction.
- **default_bad_cost** (*Optional [float]*) – If a run is reported as bad and *default_bad_cost* is provided, the cost for the bad run is set to this default value. If *default_bad_cost* is *None*, then the worst cost received is set to all the bad runs. Default *None*.
- **default_bad_uncertainty** (*Optional [float]*) – If a run is reported as bad and *default_bad_uncertainty* is provided, the uncertainty for the bad run is set to this default value. If *default_bad_uncertainty* is *None*, then the uncertainty is set to one tenth of the best to worst cost range. Default *None*.
- **minimum_uncertainty** (*Optional [float]*) – The minimum uncertainty associated with provided costs. Must be above zero to avoid fitting errors. Default *1e-8*.
- **predict_global_minima_at_end** (*Optional [bool]*) – If *True*, find the global minima when the learner is ended. Does not if *False*. Default *True*.

all_params

Array containing all parameters sent to learner.

Type array

all_costs

Array containing all costs sent to learner.

Type array

all_uncers

Array containing all uncertainties sent to learner.

Type array

scaled_costs

Array containing all the costs scaled to have zero mean and a standard deviation of 1.

Type array

bad_run_indexes

list of indexes to all runs that were marked as bad.

Type list

best_cost

Minimum received cost, updated during execution.

Type float

best_params

Parameters of best run. (reference to element in params array).

Type array

best_index

Index of the best cost and params.

Type int

worst_cost
Maximum received cost, updated during execution.

Type float

worst_index
Index to run with worst cost.

Type int

cost_range
Difference between *worst_cost* and *best_cost*

Type float

generation_num
Number of sets of parameters to generate each generation. Set to 3.

Type int

cost_count
Counter for the number of costs, parameters and uncertainties added to learner.

Type int

params_count
Counter for the number of parameters asked to be evaluated by the learner.

Type int

neural_net
Neural net that is fitted to data and used to make predictions.

Type NeuralNet

cost_scaler
Scaler used to normalize the provided costs.

Type StandardScaler

cost_scaler_init_index
The number of params to use to initialise *cost_scaler*.

Type int

has_trust_region
Whether the learner has a trust region.

Type bool

OUT_TYPE = 'neural_net'

_fit_neural_net (*index*)
Fits a neural net to the data.

cost_scaler must have been fitted before calling this method.

_init_cost_scaler ()
Initialises the cost scaler. cost_scaler_init_index must be set.

create_neural_net ()
Creates the neural net. Must be called from the same process as fit_neural_net, predict_cost and predict_costs_from_param_array.

find_global_minima (*net_index=None*)

Search for the global minima predicted by the neural net.

This method will attempt to find the global minima predicted by the neural net, but it is possible for it to become stuck in local minima of the predicted cost landscape.

This method does not return any values, but creates the attributes listed below.

Parameters **net_index** (*int, optional*) – The index of the neural net to use to predict the cost. If *None* then a net will be randomly chosen. Defaults to *None*.

predicted_best_parameters

The parameter values which are predicted to yield the best results, as a 1D array.

Type array

predicted_best_cost

The predicted cost at the *predicted_best_parameters*, in a 1D 1-element array.

Type array

find_next_parameters (*net_index=None*)

Get the next parameters to test.

This method searches for the parameters expected to give the minimum cost, as predicted by a neural net.

This method additionally increments *self.params_count* appropriately.

Parameters **net_index** (*int, optional*) – The index of the neural net to use to predict the cost. If *None* then a net will be randomly chosen. Defaults to *None*.

Returns

The next parameter values to try, stored in a 1D array.

Return type next_params (array)

get_losses ()

get_regularization_histories ()

Get the regularization coefficient values used by the nets.

Returns

The values used by the neural nets for the regularization coefficient. There is one list per net, which includes all of the regularization coefficient values used by that net during the optimization. If the optimization was run with *update_hyperparameters* set to *False*, then each net's list will only have one entry, namely the initial default value for the regularization coefficient. If the optimization was run with *updated_hyperparameters* set to *True* then the list will also include the optimal values for the regularization coefficient determined during each hyperparameter fitting.

Return type list of list of float

import_neural_net ()

Imports neural net parameters from the training dictionary provided at construction. Must be called from the same process as *fit_neural_net*, *predict_cost* and *predict_costs_from_param_array*. You must call exactly one of this and *create_neural_net* before calling other methods.

predict_cost (*params, net_index=None, perform_scaling=True*)

Predict the cost from the neural net for *params*.

This method is a wrapper around *mloop.neuralnet.NeuralNet.predict_cost()*.

Parameters

- **params** (*array*) – A 1D array containing the values for each parameter. These should be in real/unscaled units if *perform_scaling* is *True* or they should be in scaled units if *perform_scaling* is *False*.
- **net_index** (*int*, *optional*) – The index of the neural net to use to predict the cost. If *None* then a net will be randomly chosen. Defaults to *None*.
- **perform_scaling** (*bool*, *optional*) – Whether or not the parameters and costs should be scaled. If *True* then this method takes in parameter values in real/unscaled units then returns a predicted cost in real/unscaled units. If *False*, then this method takes parameter values in scaled units and returns a cost in scaled units. Note that this method cannot determine on its own if the values in *params* are in real/unscaled units or scaled units; it is up to the caller to pass the correct values. Defaults to *True*.

Returns

Predicted cost for *params*. This will be in real/unscaled units if *perform_scaling* is *True* or it will be in scaled units if *perform_scaling* is *False*.

Return type cost (float)

predict_cost_gradient (*params*, *net_index=None*, *perform_scaling=True*)

Predict the gradient of the cost function at *params*.

This method is a wrapper around *mloop.neuralnet.NeuralNet.predict_cost_gradient()*.

Parameters

- **params** (*array*) – A 1D array containing the values for each parameter. These should be in real/unscaled units if *perform_scaling* is *True* or they should be in scaled units if *perform_scaling* is *False*.
- **net_index** (*int*, *optional*) – The index of the neural net to use to predict the cost gradient. If *None* then a net will be randomly chosen. Defaults to *None*.
- **perform_scaling** (*bool*, *optional*) – Whether or not the parameters and costs should be scaled. If *True* then this method takes in parameter values in real/unscaled units then returns a predicted cost gradient in real/unscaled units. If *False*, then this method takes parameter values in scaled units and returns a cost gradient in scaled units. Note that this method cannot determine on its own if the values in *params* are in real/unscaled units or scaled units; it is up to the caller to pass the correct values. Defaults to *True*.

Returns

The predicted gradient at *params*. This will be in real/unscaled units if *perform_scaling* is *True* or it will be in scaled units if *perform_scaling* is *False*.

Return type cost_gradient (np.float64)

predict_costs_from_param_array (*params*, *net_index=None*, *perform_scaling=True*)

Produces an array of predicted costs from an array of *params*.

This method is similar to *self.predict_cost()* except that it accepts many different sets of parameter values simultaneously and returns a predicted cost for each set of parameter values.

Parameters

- **params** (*array*) – A 2D array containing the values for each parameter. This should essentially be a list of sets of parameters values, where each set specifies one value for each parameter. In other words, each row of the array should be one set of parameter values which could be passed to *self.predict_cost()*. These should be in real/unscaled units if *perform_scaling* is *True* or they should be in scaled units if *perform_scaling* is *False*.

- **net_index**(*int*, *optional*) – The index of the neural net to use to predict the cost. If *None* then a net will be randomly chosen. Defaults to *None*.
- **perform_scaling**(*bool*, *optional*) – Whether or not the parameters and costs should be scaled. If *True* then this method takes in parameter values in real/unscaled units then returns a predicted cost in real/unscaled units. If *False*, then this method takes parameter values in scaled units and returns a cost in scaled units. Note that this method cannot determine on its own if the values in *params* are in real/unscaled units or scaled units; it is up to the caller to pass the correct values. Defaults to *True*.

Returns

A list of floats giving the predicted cost for each set of parameter values.

Return type list

run()

Starts running the neural network learner. When the new parameters event is triggered, reads the cost information provided and updates the neural network with the information. Then searches the neural network for new optimal parameters to test based on the biased cost. Parameters to test next are put on the output parameters queue.

update_archive()

Update the archive.

class `mloop.learners.RandomLearner`(*trust_region=None*, *first_params=None*, ***kwargs*)

Bases: `mloop.learners.Learner`, `threading.Thread`

Random learner. Simply generates new parameters randomly with a uniform distribution over the boundaries. Learner is perhaps a misnomer for this class.

Parameters ***kwargs* (*Optional dict*) – Other values to be passed to Learner.

Keyword Arguments

- **min_boundary** (*Optional [array]*) – If set to *None*, overrides default learner values and sets it to a set of value 0. Default *None*.
- **max_boundary** (*Optional [array]*) – If set to *None* overrides default learner values and sets it to an array of value 1. Default *None*.
- **first_params** (*Optional [array]*) – The first parameters to test. If *None* will just randomly sample the initial condition.
- **trust_region** (*Optional [float or array]*) – The trust region defines the maximum distance the learner will travel from the current best set of parameters. If *None*, the learner will search everywhere. If a float, this number must be between 0 and 1 and defines maximum distance the learner will venture as a percentage of the boundaries. If it is an array, it must have the same size as the number of parameters and the numbers define the maximum absolute distance that can be moved along each direction.

OUT_TYPE = 'random'

run()

Puts the next parameters on the queue which are randomly picked from a uniform distribution between the minimum and maximum boundaries when a cost is added to the cost queue.

2.9.6 testing

Module of classes used to test M-LOOP.

```
class mloop.testing.FakeExperiment (test_landscape=None,      experiment_file_type='txt',
                                   exp_wait=0, poll_wait=1, **kwargs)
```

Bases: threading.Thread

Pretends to be an experiment and reads files and prints files based on the costs provided by a TestLandscape. Executes as a thread.

Keyword Arguments

- **test_landscape** (*Optional TestLandscape*) – landscape to generate costs from.
- **experiment_file_type** (*Optional [string]*) – currently supports: 'txt' where the output is a text file with the parameters as a list of numbers, and 'mat' a matlab file with variable parameters with the next_parameters. Default is 'txt'.

Attributes self.end_event (Event): Used to trigger end of experiment.

run()

Implementation of file read in and out. Put parameters into a file and wait for a cost file to be returned.

set_landscape (test_landscape)

Set new test landscape.

Parameters test_landscape (*TestLandscape*) – Landscape to generate costs from.

```
class mloop.testing.TestLandscape (num_params=1)
```

Bases: object

Produces fake landscape data for testing, default functions are set for each of the methods which can then be over ridden.

Keyword Arguments num_parameters (*Optional [int]*) – Number of parameters for landscape, defaults to 1.

get_cost_dict (params)

Return cost from fake landscape given parameters.

Parameters params (*array*) – Parameters to evaluate cost.

set_bad_region (min_boundary, max_boundary, bad_cost=None, bad_uncer=None)

Adds a region to landscape that is reported as bad.

Parameters

- **min_boundary** (*array*) – minimum boundary for bad region
- **max_boundary** (*array*) – maximum boundary for bad region

set_default_landscape()

Set landscape functions to their defaults

set_noise_function (proportional=0.0, absolute=0.0)

Adds noise to the function.

with the formula:

$$c'(c, x) = c (1 + s_p p) + s_a a$$

where s_i are gaussian random variables, p is the proportional noise factor and a is the absolute noise factor, and c is the cost before noise is added

the uncertainty is then:

```
u = sqrt((cp)^2 + a^2)
```

Keyword Arguments

- **proportional** (*Optional [float]*) – the proportional factor. Defaults to 0
- **absolute** (*Optional [float]*) – the absolute factor. Defaults to 0

set_quadratic_landscape (*minimum=None, scale=None*)

Set deterministic part of landscape to be a quadratic.

with the formula:

```
c(x) = \sum_i a_i * (x_i - x_0,i)^2
```

where x_i are the parameters, $x_{0,i}$ is the location of the minimum and a_i are the scaling factors.

Keyword Arguments

- **minimum** (*Optional [array]*) – Location of the minimum. If set to None is at the origin. Default None.
- **scales** (*Optional [array]*) – scaling of quadratic along the dimension specified. If set to None the scaling is one.

set_random_quadratic_landscape (*min_region, max_region, random_scale=True, min_scale=0, max_scale=3*)

Make a quadratic function with a minimum randomly placed in a region with random scales

Parameters

- **min_region** (*array*) – minimum region boundary
- **max_region** (*array*) – maximum region boundary

Keyword Arguments

- **random_scale** (*Optional bool*) – If true randomize the scales of the parameters. Default True.
- **min_scale** (*float*) – Natural log of minimum scale factor. Default 0.
- **max_scale** (*float*) – Natural log of maximum scale factor. Default 3.

2.9.7 utilities

Module of common utility methods and attributes used by all the modules.

class mloop.utilities.NullQueueListener

Bases: object

Shell class with start and stop functions that do nothing. Queue listener is not implemented in python 2. Current fix is to simply use the multiprocessing class to pipe straight to the cmd line if running on python 2. This is class is just a placeholder.

start ()

Does nothing

stop ()

Does nothing

```
class mloop.utilities.ParameterScaler (min_boundary, max_boundary, **kwargs)
```

Bases: sklearn.preprocessing._data.MinMaxScaler

Class for scaling parameters based on their min/max value constraints.

All parameters are mapped (by default) between [0, 1] by linearly rescaling them. In particular the minimum value for a parameter is mapped to 0 and the maximum value is mapped to 1.

This class inherits from scikit-learn's *MinMaxScaler*. The primary difference is that values are scaled by the minimum and maximum values set by the user, rather than the minimum and maximum values actually used in a dataset.

Parameters

- **min_boundary** (*np.array*) – The minimum values allowed for each parameter.
- **max_boundary** (*np.array*) – The maximum values allowed for each parameter.
- ***args** – Additional arguments are passed to *MinMaxScaler.__init__()*.
- ****kwargs** – Arbitrary keyword arguments are passed to *MinMaxScaler.__init__()*.

```
partial_fit (X=None, *args, **kwargs)
```

Teach the scaler that we want to scale things based on the minimum and maximum boundaries

```
mloop.utilities.__config_logger (log_filename='M-LOOP', file_log_level=10, console_log_level=20, start_datetime=None, **kwargs)
```

Configure and the root logger.

Keyword Arguments

- **log_filename** (*Optional [string]*) – Filename prefix for log. Default M-LOOP run . If None, no file handler is created
- **file_log_level** (*Optional[int]*) – Level of log output for file, default is logging.DEBUG = 10
- **console_log_level** (*Optional[int]*) – Level of log output for console, default is logging.INFO = 20
- **start_datetime** (*Optional datetime.datetime*) – The date and time to use in the filename suffix, represented as an instance of the datetime class defined in the datetime module. If set to None, then this function will use the result returned by datetime.datetime.now(). Default None.

Returns Dict with extra keywords not used by the logging configuration.

Return type dictionary

```
mloop.utilities.__generate_legend_labels (param_indices, all_param_names)
```

Generate a list of labels for the legend of a plot.

This is a helper function for visualization methods, used to generate the labels in legends for plots that show the values for optimization parameters. The label has the parameter's index and, if available, a colon followed by the parameter's name e.g. '3: some_name'. If no name is available, then the label will simply be a string representation of the parameter's index, e.g. '3'.

Parameters

- **param_indices** (*list-like of int*) – The indices of the parameters for which labels should be generated. Generally this should be the same as the list of indices of parameters included in the plot.

- **all_param_names** (*list-like of str*) – The names of all parameters from the optimization. Note this this argument should be *all* of the names for all of the parameters, not just the ones to be included in the plot legend.

Returns The labels generated for use in a plot legend.

Return type labels (list of str)

`mloop.utilities._param_names_from_file_dict` (*file_dict*)

Extract the value for 'param_names' from a training dictionary.

Versions of M-LOOP <= 2.2.0 didn't support the param_names option, so archives generated by those versions do not have an entry for param_names. This helper function takes the dict generated when `get_dict_from_file()` is called on an archive and returns the value for param_names if present. If there is no entry for param_names, it returns a list of empty strings. This makes it possible to use the param_names data in archives from newer versions of M-LOOP while retaining the ability to plot data from archives generated by older versions of M-LOOP.

Parameters **file_dict** (*dict*) – A dict containing data from an archive, such as those returned by `get_dict_from_file()`.

Returns

List of the names of the optimization parameters if present in `file_dict`. If not present, then param_names will be set to None.

Return type param_names (list of str)

`mloop.utilities.check_file_type_supported` (*file_type*)

Checks whether the file type is supported

Returns True if file_type is supported, False otherwise.

Return type bool

`mloop.utilities.chunk_list` (*list_, chunk_size*)

Divide a list into sublists of length chunk_size.

All elements in `list_` will be included in exactly one of the sublists and will be in the same order as in `list_`. If the length of `list_` is not divisible by chunk_size, then the final sublist returned will have fewer than chunk_size elements.

Examples

```
>>> chunk_list([1, 2, 3, 4, 5], 2)
[[1, 2], [3, 4], [5]]
>>> chunk_list([1, 2, 3, 4, 5], None)
[[1, 2, 3, 4, 5]]
>>> chunk_list([1, 2, 3, 4, 5], float('inf'))
[[1, 2, 3, 4, 5]]
```

Parameters

- **list** (*list-like*) – A list (or similar) to divide up into smaller lists.
- **chunk_size** (*int*) – The number of elements to have in each sublist. The last sublist will have fewer elements than this if the length of `list_` is not divisible by chunk_size. If set to float('inf') or None, then all elements will be put into one sublist.

Returns

List of sublists, each of which contains elements from the input **list_**. Each sublist has length `chunk_size` except for the last one which may have fewer elements.

Return type (list)

`mloop.utilities.config_logger(**kwargs)`

Wrapper for `_config_logger`.

`mloop.utilities.datetime_to_string(datetime)`

Method for changing a datetime into a standard string format used by all packages.

`mloop.utilities.dict_to_txt_file(tdict, filename)`

Method for writing a dict to a file with syntax similar to how files are input.

Parameters

- **tdict** (*dict*) – Dictionary to be written to file.
- **filename** (*string*) – Filename for file.

`mloop.utilities.generate_filename_suffix(file_type, file_datetime=None, random_bytes=False)`

Method for generating a string with date and extension for end of file names.

This method returns a string such as ‘_2020-06-13_04-20.txt’ where the date and time specify when this function was called.

Parameters

- **file_type** (*string*) – The extension to use at the end of the filename, e.g. ‘txt’. Note that the period should NOT be included.
- **file_datetime** (*Optional datetime.datetime*) – The date and time to use in the filename suffix, represented as an instance of the datetime class defined in the datetime module. If set to None, then this function will use the result returned by `datetime.datetime.now()`. Default None.
- **random_bytes** (*Optional bool*) – If set to True, six random bytes will be added to the filename suffix. This can be useful avoid duplication if multiple filenames are created with the same datetime.

Returns

A string giving the suffix that can be appended to a filename prefix to give a full filename with timestamp and extension, such as ‘_2020-06-13_04-20.txt’. The date and time specify when this function was called.

Return type string

`mloop.utilities.get_controller_type_from_learner_archive(learner_filename)`

Determine the controller_type used in an optimization.

This function returns the value used for `controller_type` during an optimization run, determined by examining the optimization’s learner archive.

Parameters **learner_filename** (*String*) – The file name including extension, and optionally including path, of a learner archive.

Returns

A string specifying the value for `controller_type` used during the optimization run that produced the provided learner archive.

Return type `controller_type` (String)

`mloop.utilities.get_dict_from_file(filename, file_type=None)`

Method for getting a dictionary from a file, of a given format.

Parameters `filename` (*str*) – The filename for the file.

Keyword Arguments `file_type` (*Optional str*) – The file_type for the file. Can be ‘mat’ for matlab, ‘txt’ for text, or ‘pkl’ for pickle. If set to None, then file_type will be automatically determined from the file extension. Default None.

Returns Dictionary of values in file.

Return type dict

`mloop.utilities.get_file_type(filename)`

Get the file type of a file from the extension in its filename.

Parameters `filename` (*String*) – The filename including extension, and optionally including path, from which to extract the file type.

Returns

The file’s type, inferred from its extension. The type does NOT include a leading period.

Return type file_type (String)

`mloop.utilities.safe_cast_to_array(in_array)`

Attempts to safely cast the input to an array. Takes care of border cases

Parameters `in_array` (*array or equivalent*) – The array (or otherwise) to be converted to a list.

Returns array that has been squeezed and 0-D cases change to 1-D cases

Return type array

`mloop.utilities.safe_cast_to_list(in_array)`

Attempts to safely cast a numpy array to a list, if not a numpy array just casts to list on the object.

Parameters `in_array` (*array or equivalent*) – The array (or otherwise) to be converted to a list.

Returns List of elements from in_array

Return type list

`mloop.utilities.save_dict_to_file(dictionary, filename, file_type=None)`

Method for saving a dictionary to a file, of a given format.

Parameters

- **dictionary** – The dictionary to be saved in the file.
- **filename** – The filename for the saved file.

Keyword Arguments `file_type` (*Optional str*) – The file_type for the file. Can be ‘mat’ for matlab, ‘txt’ for text, or ‘pkl’ for pickle. If set to None, then file_type will be automatically determined from the file extension. Default None.

`mloop.utilities.txt_file_to_dict(filename)`

Method for taking a file and changing it to a dict. Every line in file is a new entry for the dictionary and each element should be written as:

```
[key] = [value]
```

White space does not matter.

Parameters `filename` (*string*) – Filename of file.

Returns Dictionary of values in file.

Return type dict

2.9.8 visualizations

Module of classes used to create visualizations of data produced by the experiment and learners.

class `mloop.visualizations.ControllerVisualizer` (*filename*, *file_type=None*, ***kwargs*)
Bases: object

ControllerVisualizer creates figures from a Controller Archive.

Note that the data from the training archive, if one was provided to the learner at the beginning of the optimization, is NOT included in the controller archive generated during the optimization. Therefore any data from the training archive is not included in the plots generated by this class. This is in contrast to some of the learner visualizer classes.

Parameters `filename` (*String*) – Filename of the controller archive.

Keyword Arguments `file_type` (*String*) – Can be ‘mat’ for matlab, ‘pkl’ for pickle or ‘txt’ for text. If set to None, then the type will be determined from the extension in filename. Default None.

create_visualizations (*plot_cost_vs_run=True*, *plot_parameters_vs_run=True*,
plot_parameters_vs_cost=True, *max_parameters_per_plot=None*)
Runs the plots for a controller file.

Keyword Arguments

- **plot_cost_vs_run** (*Optional [bool]*) – If True plot cost versus run number, else do not. Default True.
- **plot_parameters_vs_run** (*Optional [bool]*) – If True plot parameters versus run number, else do not. Default True.
- **plot_parameters_vs_cost** (*Optional [bool]*) – If True plot parameters versus cost number, else do not. Default True.
- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than `parameters_per_plot`, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to None, then all parameters will be included in the same plot regardless of how many there are. Default None.

plot_cost_vs_run ()

Create a plot of the costs versus run number.

Note that the data from the training archive, if one was provided to the learner at the beginning of the optimization, will NOT be plotted here.

plot_parameters_vs_cost (*parameter_subset=None*)

Create a plot of the parameters versus run number.

Note that the data from the training archive, if one was provided to the learner at the beginning of the optimization, will NOT be plotted here.

Parameters `parameter_subset` (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the

values of the indices in `parameter_subset` should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default *None*.

plot_parameters_vs_run (*parameter_subset=None*)

Create a plot of the parameters versus run number.

Note that the data from the training archive, if one was provided to the learner at the beginning of the optimization, will NOT be plotted here.

Parameters `parameter_subset` (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in `parameter_subset` should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default *None*.

```
class mloop.visualizations.DifferentialEvolutionVisualizer (filename,
                                                         file_type=None,
                                                         **kwargs)
```

Bases: object

DifferentialEvolutionVisualizer creates figures from a differential evolution archive.

Parameters `filename` (*String*) – Filename of the DifferentialEvolutionVisualizer archive.

Keyword Arguments `file_type` (*String*) – Can be ‘mat’ for matlab, ‘pkl’ for pickle or ‘txt’ for text. If set to *None*, then the type will be determined from the extension in filename. Default *None*.

```
create_visualizations (plot_params_vs_generations=True, plot_costs_vs_generations=True,
                        max_parameters_per_plot=None)
```

Runs the plots from a differential evolution learner file.

Keyword Arguments

- **plot_params_generations** (*Optional [bool]*) – If True plot parameters vs generations, else do not. Default True.
- **plot_costs_generations** (*Optional [bool]*) – If True plot costs vs generations, else do not. Default True.
- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than `parameters_per_plot`, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to *None*, then all parameters will be included in the same plot regardless of how many there are. Default *None*.

```
plot_costs_vs_generations ()
```

Create a plot of the costs versus run number.

```
plot_params_vs_generations (parameter_subset=None)
```

Create a plot of the parameters versus run number.

Parameters `parameter_subset` (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in `parameter_subset` should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default *None*.

```
class mloop.visualizations.GaussianProcessVisualizer (filename, **kwargs)
```

Bases: `mloop.learners.GaussianProcessLearner`

A class for visualizing results from a Gaussian process optimization.

GaussianProcessVisualizer extends of *GaussianProcessLearner*, but is designed not to be used as a learner, but to instead post process a *GaussianProcessLearner* archive file and produce visualizations. This class fixes the Gaussian process hyperparameters to what was last found during the run.

If a training archive was provided at the start of the optimization using *training_filename* and that training archive was generated by a Gaussian process optimization, then some of its data is saved in the new learner archive generated during the optimization. That implies that some of the data, such as fitted hyperparameter values, from the training archive will be included in the plots generated by this class.

Parameters *filename* (*String*) – Filename of the *GaussianProcessLearner* archive.

create_visualizations (*plot_cross_sections=True*, *plot_hyperparameters_vs_fit=True*,
plot_noise_level_vs_fit=True, *max_parameters_per_plot=None*,
***kwargs*)

Runs the plots from a gaussian process learner file.

Keyword Arguments

- **plot_cross_sections** (*Optional [bool]*) – If *True* plot predicted landscape cross sections, else do not. Default *True*.
- **plot_hyperparameters_vs_fit** (*Optional [bool]*) – If *True* plot fitted hyperparameters as a function of fit number, else do not. Default *True*.
- **plot_noise_level_vs_fit** (*Optional [bool]*) – If *True* plot the fitted noise level as a function of fit number, else do not. If there is no fitted noise level (i.e. *cost_has_noise* was set to *False*), then this plot will not be made regardless of the value passed for *plot_noise_level_vs_fit*. Default *True*.
- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than *parameters_per_plot*, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to *None*, then all parameters will be included in the same plot regardless of how many there are. Default *None*.

plot_cross_sections (*parameter_subset=None*)

Produce a figure of the cross section about best cost and parameters.

This method will produce plots showing cross sections of the predicted cost landscape along each parameter axis through the point in parameter space which gave the best measured cost. In other words, one parameter will be varied from its minimum allowed value to its maximum allowed value with the other parameters fixed at the values that they had in the set of parameters that gave the best measured cost. At each point the predicted cost will be plotted. That process will be repeated for each parameter in *parameter_subset*. The x axes will be scaled to the range 0 to 1, corresponding to the minimum and maximum bound respectively for each parameter, so that curves from different cross sections can be overlaid nicely.

The expected value of the cost will be plotted as a solid line. Additionally, dashed lines representing the 1-sigma uncertainty in the predicted cost will be plotted as well. This uncertainty includes contributions from the uncertainty in the model due to taking only a finite number of observations. Additionally, if *cost_has_noise* was set to *True* then the fitted noise level will be added in quadrature with the model uncertainty. Note that as more data points are taken the uncertainty in the model generally decreases, but the predicted noise level will typically converge to a nonzero value. That implies that the predicted cost uncertainty generally won't tend towards zero if *cost_has_noise* is set to *True*, even if many observations are made.

Parameters *parameter_subset* (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in *parameter_subset* should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default *None*.

plot_hyperparameters_vs_fit (*parameter_subset=None*)

Produce a figure of the hyperparameters as a function of fit number.

Only one fit is performed per generation, and multiple parameter sets are run each generation. Therefore the number of fits is generally less than the number of runs.

The plot generated will include the data from the training archive if one was provided with *training_filename* and the training archive was generated by a Gaussian process optimization.

Parameters *parameter_subset* (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in *parameter_subset* should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default *None*.

plot_hyperparameters_vs_run (**args, **kwargs*)

Deprecated. Use *plot_hyperparameters_vs_fit()* instead.

plot_noise_level_vs_fit ()

This method plots the fitted noise level as a function of fit number.

The *noise_level* approximates the variance of values that would be measured if the cost were repeatedly measured for the same set of parameters. Note that this is the variance in those costs; not the standard deviation.

This plot is only relevant to optimizations for which *cost_has_noise* is *True*. If *cost_has_noise* is *False* then this method does nothing and silently returns.

Only one fit is performed per generation, and multiple parameter sets are run each generation. Therefore the number of fits is generally less than the number of runs.

The plot generated will include the data from the training archive if one was provided with *training_filename* and the training archive was generated by a Gaussian process optimization.

plot_noise_level_vs_run (**args, **kwargs*)

Deprecated. Use *plot_noise_level_vs_fit()* instead.

return_cross_sections (*points=100, cross_section_center=None*)

Generate 1D cross sections along each parameter axis.

The cross sections are returned as a list of vectors of parameters values, costs, and uncertainties, corresponding to the 1D cross sections along each parameter axis. The cross sections pass through *cross_section_center*, which will default to the parameters that gave the best measured cost.

Keyword Arguments

- **points** (*int*) – The number of points to sample along each cross section. Defaults to 100.
- **cross_section_center** (*array*) – Parameter array where the center of the cross section should be taken. If *None*, the parameters for the best measured cost are used. Default *None*.

Returns

A tuple (cross_arrays, cost_arrays, uncer_arrays) cross_parameter_arrays (list): A list of arrays for each cross

section, with the values of the varied parameter going from the minimum to maximum allowed value.

cost_arrays (list): A list of arrays for the costs evaluated along each cross section through *cross_section_center*.

uncertainty_arrays (list): A list of uncertainties corresponding to the costs in *cost_arrays*.

run ()

Overrides the GaussianProcessLearner multiprocessing run routine. Does nothing but makes a warning.

class `mloop.visualizations.NelderMeadVisualizer` (*filename*, *file_type=None*, ***kwargs*)

Bases: `object`

Visualization class for the Nelder-Mead learner.

Currently no learner plots are generated for the Nelder-Mead learner. However, controller visualizations may still be generated.

create_visualizations (**args*, ***kwargs*)

class `mloop.visualizations.NeuralNetVisualizer` (*filename*, *file_type=None*, ***kwargs*)

Bases: `mloop.learners.NeuralNetLearner`

A class for visualizing results from a neural net optimization.

NeuralNetVisualizer extends of *NeuralNetLearner*, but is designed not to be used as a learner, but to instead post process a *NeuralNetLearner* archive file and produce visualizations.

Parameters **filename** (*String*) – Filename of the *NeuralNetLearner* archive.

create_visualizations (*plot_cross_sections=True*, *max_parameters_per_plot=None*)

Creates plots from a neural net's learner file.

Keyword Arguments

- **plot_cross_sections** (*Optional [bool]*) – If True plot predicted landscape cross sections, else do not. Default True.
- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than *parameters_per_plot*, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to None, then all parameters will be included in the same plot regardless of how many there are. Default None.

do_cross_sections (*parameter_subset=None*, *plot_individual_cross_sections=True*)

Produce figures of the cross section about best cost and parameters.

This method will produce plots showing cross sections of the predicted cost landscape along each parameter axis through the point in parameter space which gave the best measured cost. In other words, one parameter will be varied from its minimum allowed value to its maximum allowed value with the other parameters fixed at the values that they had in the set of parameters that gave the best measured cost. At each point the predicted cost will be plotted. That process will be repeated for each parameter in *parameter_subset*. The x axes will be scaled to the range 0 to 1, corresponding to the minimum and maximum bound respectively for each parameter, so that curves from different cross sections can be overlaid nicely.

One plot will be created which includes a solid line that shows the mean of the cost landscapes predicted by each net, as well as two dashed lines showing the minimum and maximum of the costs predicted by the nets for those parameter values. If *plot_individual_cross_sections* is set to *True* then additional plots will be created, one for each net, which show each net's predicted cost landscape cross sections.

Parameters

- **parameter_subset** (*list-like*) – The indices of parameters to plot. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in *parameter_subset* should be between 0 and the number of parameters minus one, inclusively. If set to *None*, then all parameters will be plotted. Default None.

- **plot_individual_cross_sections** (*bool*) – Whether or not to create extra plots to show each net’s predicted cross sections in its own figure. The plot of the average/min/max of the different nets’ predicted cross sections in one figure will be generated regardless of this setting. Default *True*.

plot_density_surface()

Produce a density plot of the cost surface (only works when there are 2 parameters)

plot_losses()

Produce a figure of the loss as a function of epoch for each net.

The loss is the mean-squared fitting error of the neural net plus the regularization loss, which is the regularization coefficient times the mean L2 norm of the neural net weight arrays (without the square root). Note that the fitting error is calculated after normalizing the data, so it is in arbitrary units.

As the neural nets are fit, the loss is recorded every 10 epochs. The number of epochs per fit varies, and may be different for different nets. The loss will generally increase at the beginning of each fit as new data points will have been added.

Also note that a lower loss isn’t always better; a loss that is too low can be a sign of overfitting.

plot_regularization_history()

Produces a plot of the regularization coefficient values used.

The neural nets use L2 regularization to smooth their predicted landscapes in an attempt to avoid overfitting the data. The strength of the regularization is set by the regularization coefficient, which is a hyperparameter that is tuned during the optimization if *update_hyperparameters* is set to *True*. Generally larger regularization coefficient values force the landscape to be smoother while smaller values allow it to vary more quickly. A value too large can lead to underfitting while a value too small can lead to overfitting. The ideal regularization coefficient value will depend on many factors, such as the shape of the actual cost landscape, the SNR of the measured costs, and even the number of measured costs.

This method plots the initial regularization coefficient value and the optimal values found for the regularization coefficient when performing the hyperparameter tuning. One curve showing the history of values used for the regularization coefficient is plotted for each neural net. If *update_hyperparameters* was set to *False* during the optimization, then only the initial default value will be plotted.

plot_surface()

Produce a figure of the cost surface (only works when there are 2 parameters)

return_cross_sections (*points=100, cross_section_center=None*)

Generate 1D cross sections along each parameter axis.

The cross sections are returned as a list of vectors of parameters values and costs, corresponding to the 1D cross sections along each parameter axis. The cross sections pass through *cross_section_center*, which will default to the parameters that gave the best measured cost. One such pair of list of parameter vectors and corresponding predicted costs are returned for each net.

Keyword Arguments

- **points** (*int*) – the number of points to sample along each cross section. Default value is 100.
- **cross_section_center** (*array*) – parameter array where the centre of the cross section should be taken. If *None*, the parameters for the best measured cost are used. Default *None*.

Returns

a list of tuple (cross_arrays, cost_arrays), one tuple for each net. *cross_parameter_arrays* (list): a list of arrays for each cross section, with the values of the varied parameter go-

ing from the minimum to maximum value. `cost_arrays` (list): a list of arrays for the costs evaluated along

each cross section through `cross_section_center`.

run()

Overrides the GaussianProcessLearner multiprocessing run routine. Does nothing but makes a warning.

class `mloop.visualizations.RandomVisualizer` (*filename*, *file_type=None*, ***kwargs*)

Bases: `object`

Visualization class for the random learner.

Currently no learner plots are generated for the random learner. However, controller visualizations may still be generated.

create_visualizations (**args*, ***kwargs*)

`mloop.visualizations._color_list_from_num_options` (*num_of_params*)

Gives a list of colors based on a number of options.

A distinct color will be generated for each option.

`mloop.visualizations._ensure_parameter_subset_valid` (*visualizer*, *parameter_subset*)

Make sure indices in `parameter_subset` are acceptable.

Parameters

- **visualizer** (*ControllerVisualizer-like*) – An instance of one of the visualization classes defined in this module, which should have the attributes `param_numbers` and `log`.
- **parameter_subset** (*list-like*) – The indices corresponding to a subset of the optimization parameters. The indices should be 0-based, i.e. the first parameter is identified with index 0. Generally the values of the indices in `parameter_subset` should be integers between 0 and the number of parameters minus one, inclusively.

`mloop.visualizations.configure_plots()`

Configure the setting for the plots.

`mloop.visualizations.create_controller_visualizations` (*filename*, *file_type=None*, ***kwargs*)

Runs the plots for a controller file.

Parameters `filename` (*String*) – Filename of the controller archive.

Keyword Arguments

- **file_type** (*String*) – Can be ‘mat’ for matlab, ‘pkl’ for pickle or ‘txt’ for text. If set to None, then the type will be determined from the extension in filename. Default None.
- ****kwargs** – Additional keyword arguments are passed to the visualizer’s `create_visualizations()` method.

`mloop.visualizations.create_differential_evolution_learner_visualizations` (*filename*, *file_type=None*, ***kwargs*)

Runs the plots from a differential evolution learner file.

Parameters `filename` (*String*) – Filename for the differential evolution learner archive.

Keyword Arguments

- **file_type** (*String*) – Can be ‘mat’ for matlab, ‘pkl’ for pickle or ‘txt’ for text. If set to None, then the type will be determined from the extension in filename. Default None.

- ****kwargs** – Additional keyword arguments are passed to the visualizer’s `create_visualizations()` method.

```
mloop.visualizations.create_gaussian_process_learner_visualizations(filename,
                                                                    file_type=None,
                                                                    **kwargs)
```

Runs the plots from a gaussian process learner file.

Parameters `filename` (*String*) – Filename for the gaussian process learner archive.

Keyword Arguments

- **file_type** (*String*) – Can be ‘mat’ for matlab, ‘pkl’ for pickle or ‘txt’ for text. If set to None, then the type will be determined from the extension in filename. Default None.
- ****kwargs** – Additional keyword arguments are passed to the visualizer’s `create_visualizations()` method.

```
mloop.visualizations.create_learner_visualizations(filename,
                                                    max_parameters_per_plot=None,
                                                    learner_visualization_kwargs=None,
                                                    learner_visualizer_init_kwargs=None)
```

Runs the plots for a learner archive file.

Parameters `filename` (*str*) – Filename for the learner archive.

Keyword Arguments

- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than `parameters_per_plot`, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to None, then all parameters will be included in the same plot regardless of how many there are. If a value for `max_parameters_per_plot` is included in `learner_visualization_kwargs`, then the value in that dictionary will override this setting. Default None.
- **learner_visualization_kwargs** (*dict*) – Keyword arguments to pass to the learner visualizer’s `create_visualizations()` method. If set to None, no additional keyword arguments will be passed. Default None.
- **learner_visualizer_init_kwargs** (*dict*) – Keyword arguments to pass to the learner visualizer’s `__init__()` method. If set to None, no additional keyword arguments will be passed. Default None.

```
mloop.visualizations.create_learner_visualizer_from_archive(filename,
                                                            controller_type=None,
                                                            **kwargs)
```

Create an instance of the appropriate visualizer class for a learner archive.

Parameters `filename` (*String*) – Filename of the learner archive.

Keyword Arguments

- **controller_type** (*NoneType*) – This argument is now deprecated and has no effect. Do not provide a value for `controller_type`; it will be removed in a future version of M-LOOP. If set to anything other than None, a warning will be issued. Default None.
- ****kwargs** – Additional keyword arguments are passed to the visualizer’s `__init__()` method.

Returns

An instance of the appropriate visualizer class for plotting data from filename.

Return type visualizer

```
mloop.visualizations.create_neural_net_learner_visualizations(filename,  
                                                             file_type=None,  
                                                             **kwargs)
```

Creates plots from a neural net's learner file.

Parameters **filename** (*String*) – Filename for the neural net learner archive.

Keyword Arguments

- **file_type** (*String*) – Can be 'mat' for matlab, 'pkl' for pickle or 'txt' for text. If set to None, then the type will be determined from the extension in filename. Default None.
- ****kwargs** – Additional keyword arguments are passed to the visualizer's `create_visualizations()` method.

```
mloop.visualizations.set_legend_location(loc=None)
```

Set the location of the legend in future figures.

Note that this function doesn't change the location of legends in existing figures. It only changes where legends will appear in figures generated after the call to this function. If called without arguments, the legend location for future figures will revert to its default value.

Keyword Arguments **loc** (*Optional str, int, or pair of floats*) – The value to use for loc in the calls to matplotlib's `legend()`. Can be e.g. 2, 'upper right', (1, 0). See matplotlib's documentation for more options and additional information. If set to None then the legend location will be set back to its default value. Default None.

```
mloop.visualizations.show_all_default_visualizations(controller, show_plots=True,  
                                                         max_parameters_per_plot=None)
```

Plots all visualizations available for a controller, and it's internal learners.

Parameters **controller** (*Controller*) – The controller to extract plots from

Keyword Arguments

- **show_plots** (*Optional, bool*) – Determine whether to run `plt.show()` at the end or not. For debugging. Default True.
- **max_parameters_per_plot** (*Optional, int*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than `parameters_per_plot`, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to None, then all parameters will be included in the same plot regardless of how many there are. Default None.

```
mloop.visualizations.show_all_default_visualizations_from_archive(controller_filename,  
                                                                    learner_filename,  
                                                                    con-  
                                                                    troller_type=None,  
                                                                    show_plots=True,  
                                                                    max_parameters_per_plot=None,  
                                                                    con-  
                                                                    troller_visualization_kwargs=None,  
                                                                    learner_visualization_kwargs=None,  
                                                                    learner_visualizer_init_kwargs=None)
```

Plots all visualizations available for a controller and its learner from their archives.

Parameters

- **controller_filename** (*str*) – The filename, including path, of the controller archive.
- **learner_filename** (*str*) – The filename, including path, of the learner archive.

Keyword Arguments

- **controller_type** (*NoneType*) – This argument is now deprecated and has no effect. Do not provide a value for `controller_type`; it will be removed in a future version of M-LOOP. If set to anything other than `None`, a warning will be issued. Default `None`.
- **show_plots** (*bool*) – Determine whether to run `plt.show()` at the end or not. For debugging. Default `True`.
- **max_parameters_per_plot** (*Optional [int]*) – The maximum number of parameters to include in plots that display the values of parameters. If the number of parameters is larger than `parameters_per_plot`, then the parameters will be divided into groups and each group will be plotted in its own figure. If set to `None`, then all parameters will be included in the same plot regardless of how many there are. If a value for `max_parameters_per_plot` is included in `controller_visualization_kwargs`, then the value in that dictionary will override this setting. The same applies to `learner_visualization_kwargs`. Default `None`.
- **controller_visualization_kwargs** (*dict*) – Keyword arguments to pass to the controller visualizer's `create_visualizations()` method. If set to `None`, no additional keyword arguments will be passed. Default `None`.
- **learner_visualization_kwargs** (*dict*) – Keyword arguments to pass to the learner visualizer's `create_visualizations()` method. If set to `None`, no additional keyword arguments will be passed. Default `None`.
- **learner_visualizer_init_kwargs** (*dict*) – Keyword arguments to pass to the learner visualizer's `__init__()` method. If set to `None`, no additional keyword arguments will be passed. Default `None`.

CHAPTER 3

Indices

- `genindex`
- `modindex`
- `search`

m

- `mloop`, [35](#)
- `mloop.controllers`, [36](#)
- `mloop.interfaces`, [42](#)
- `mloop.launchers`, [45](#)
- `mloop.learners`, [45](#)
- `mloop.testing`, [65](#)
- `mloop.utilities`, [67](#)
- `mloop.visualizations`, [72](#)

Symbols

`_best1()` (*mloop.learners.DifferentialEvolutionLearner* method), 46
`_best2()` (*mloop.learners.DifferentialEvolutionLearner* method), 46
`_check_length_scale_bounds()` (*mloop.learners.GaussianProcessLearner* method), 50
`_check_noise_level_bounds()` (*mloop.learners.GaussianProcessLearner* method), 50
`_color_list_from_num_options()` (in module *mloop.visualizations*), 78
`_config_logger()` (in module *mloop.utilities*), 68
`_enforce_boundaries()` (*mloop.controllers.Controller* method), 38
`_ensure_parameter_subset_valid()` (in module *mloop.visualizations*), 78
`_find_predicted_minimum()` (*mloop.learners.MachineLearner* method), 58
`_first_params()` (*mloop.controllers.Controller* method), 38
`_fit_neural_net()` (*mloop.learners.NeuralNetLearner* method), 62
`_generate_legend_labels()` (in module *mloop.utilities*), 68
`_get_cost_and_in_dict()` (*mloop.controllers.Controller* method), 38
`_get_cost_and_in_dict()` (*mloop.controllers.MachineLearnerController* method), 41
`_init_cost_scaler()` (*mloop.learners.NeuralNetLearner* method), 62
`_next_params()` (*mloop.controllers.Controller* method), 39
`_optimization_routine()` (*mloop.controllers.Controller* method), 39
`_optimization_routine()` (*mloop.controllers.MachineLearnerController* method), 41
`_param_names_from_file_dict()` (in module *mloop.utilities*), 69
`_parse_cost_message()` (*mloop.learners.Learner* method), 55
`_pop_extras_kwargs()` (in module *mloop.launchers*), 45
`_prepare_logger()` (*mloop.learners.Learner* method), 55
`_put_params_and_out_dict()` (*mloop.controllers.Controller* method), 39
`_rand1()` (*mloop.learners.DifferentialEvolutionLearner* method), 46
`_rand2()` (*mloop.learners.DifferentialEvolutionLearner* method), 47
`_reconcile_kwarg_and_training_val()` (*mloop.learners.MachineLearner* method), 58
`_send_to_learner()` (*mloop.controllers.Controller* method), 39
`_set_trust_region()` (*mloop.learners.Learner* method), 55
`_shut_down()` (*mloop.controllers.Controller* method), 39
`_shut_down()` (*mloop.controllers.MachineLearnerController* method), 41
`_shut_down()` (*mloop.learners.Learner* method), 55
`_start_up()` (*mloop.controllers.Controller* method), 39
`_start_up()` (*mloop.controllers.MachineLearnerController* method), 41
`_transform_length_scale_bounds()` (*mloop.learners.GaussianProcessLearner* method), 50
`_transform_length_scales()` (*mloop.learners.GaussianProcessLearner* method), 51
`_update_controller_with_learner_attributes()`

(mloop.controllers.Controller method), 39
_update_run_data_attributes()
(mloop.learners.Learner method), 55

A

all_costs (mloop.learners.GaussianProcessLearner attribute), 49
all_costs (mloop.learners.Learner attribute), 55
all_costs (mloop.learners.MachineLearner attribute), 57
all_costs (mloop.learners.NeuralNetLearner attribute), 61
all_params (mloop.learners.GaussianProcessLearner attribute), 49
all_params (mloop.learners.Learner attribute), 55
all_params (mloop.learners.MachineLearner attribute), 57
all_params (mloop.learners.NeuralNetLearner attribute), 61
all_uncers (mloop.learners.GaussianProcessLearner attribute), 49
all_uncers (mloop.learners.Learner attribute), 55
all_uncers (mloop.learners.MachineLearner attribute), 57
all_uncers (mloop.learners.NeuralNetLearner attribute), 61

B

bad_run_indexs (mloop.learners.GaussianProcessLearner attribute), 49
bad_run_indexs (mloop.learners.Learner attribute), 55
bad_run_indexs (mloop.learners.MachineLearner attribute), 57
bad_run_indexs (mloop.learners.NeuralNetLearner attribute), 61
best_cost (mloop.controllers.Controller attribute), 38
best_cost (mloop.learners.GaussianProcessLearner attribute), 49
best_cost (mloop.learners.MachineLearner attribute), 57
best_cost (mloop.learners.NeuralNetLearner attribute), 61
best_index (mloop.controllers.Controller attribute), 38
best_index (mloop.learners.GaussianProcessLearner attribute), 49
best_index (mloop.learners.MachineLearner attribute), 57
best_index (mloop.learners.NeuralNetLearner attribute), 61
best_params (mloop.controllers.Controller attribute), 38

best_params (mloop.learners.GaussianProcessLearner attribute), 49
best_params (mloop.learners.MachineLearner attribute), 57
best_params (mloop.learners.NeuralNetLearner attribute), 61
best_uncer (mloop.controllers.Controller attribute), 38

C

check_end_conditions()
(mloop.controllers.Controller method), 39
check_file_type_supported() (in module *mloop.utilities*), 69
check_in_boundary() (*mloop.learners.Learner method*), 56
check_in_diff_boundary()
(mloop.learners.Learner method), 56
check_num_params() (*mloop.learners.Learner method*), 56
chunk_list() (in module *mloop.utilities*), 69
config_logger() (in module *mloop.utilities*), 70
configure_plots() (in module *mloop.visualizations*), 78
Controller (class in *mloop.controllers*), 36
ControllerInterrupt, 40
ControllerVisualizer (class in *mloop.visualizations*), 72
cost_count (mloop.learners.GaussianProcessLearner attribute), 50
cost_count (mloop.learners.NeuralNetLearner attribute), 62
cost_range (mloop.learners.GaussianProcessLearner attribute), 49
cost_range (mloop.learners.MachineLearner attribute), 58
cost_range (mloop.learners.NeuralNetLearner attribute), 62
cost_scaler (mloop.learners.GaussianProcessLearner attribute), 50
cost_scaler (mloop.learners.NeuralNetLearner attribute), 62
cost_scaler_init_index
(mloop.learners.NeuralNetLearner attribute), 62
costs_generations
(mloop.learners.DifferentialEvolutionLearner attribute), 46
costs_in_queue (mloop.controllers.Controller attribute), 37
costs_in_queue (mloop.learners.Learner attribute), 54
create_controller() (in module *mloop.controllers*), 42

`create_controller_visualizations()` (in module `mloop.visualizations`), 78
`create_differential_evolution_learner_visualizations()` (in module `mloop.visualizations`), 78
`create_gaussian_process()` (`mloop.learners.GaussianProcessLearner` method), 52
`create_gaussian_process_learner_visualizations()` (in module `mloop.visualizations`), 79
`create_interface()` (in module `mloop.interfaces`), 44
`create_learner_visualizations()` (in module `mloop.visualizations`), 79
`create_learner_visualizer_from_archive()` (in module `mloop.visualizations`), 79
`create_neural_net()` (`mloop.learners.NeuralNetLearner` method), 62
`create_neural_net_learner_visualizations()` (in module `mloop.visualizations`), 80
`create_visualizations()` (`mloop.visualizations.ControllerVisualizer` method), 72
`create_visualizations()` (`mloop.visualizations.DifferentialEvolutionVisualizer` method), 73
`create_visualizations()` (`mloop.visualizations.GaussianProcessVisualizer` method), 74
`create_visualizations()` (`mloop.visualizations.NelderMeadVisualizer` method), 76
`create_visualizations()` (`mloop.visualizations.NeuralNetVisualizer` method), 76
`create_visualizations()` (`mloop.visualizations.RandomVisualizer` method), 78
`curr_std` (`mloop.learners.DifferentialEvolutionLearner` attribute), 46

D

`datetime_to_string()` (in module `mloop.utilities`), 70
`dict_to_txt_file()` (in module `mloop.utilities`), 70
`DifferentialEvolutionController` (class in `mloop.controllers`), 40
`DifferentialEvolutionLearner` (class in `mloop.learners`), 45
`DifferentialEvolutionVisualizer` (class in `mloop.visualizations`), 73
`do_cross_sections()` (`mloop.visualizations.NeuralNetVisualizer` method), 76

E

`end_event` (`mloop.learners.Learner` attribute), 54
`end_interface()` (`mloop.controllers.Controller` attribute), 37
`end_learner` (`mloop.controllers.Controller` attribute), 37

F

`FakeExperiment` (class in `mloop.testing`), 65
`FileInterface` (class in `mloop.interfaces`), 42
`find_global_minima()` (`mloop.learners.GaussianProcessLearner` method), 52
`find_global_minima()` (`mloop.learners.NeuralNetLearner` method), 62
`find_next_parameters()` (`mloop.learners.GaussianProcessLearner` method), 52
`find_next_parameters()` (`mloop.learners.NeuralNetLearner` method), 63
`fit_count` (`mloop.learners.GaussianProcessLearner` attribute), 50
`fit_gaussian_process()` (`mloop.learners.GaussianProcessLearner` method), 52

G

`gaussian_process` (`mloop.learners.GaussianProcessLearner` attribute), 50
`GaussianProcessController` (class in `mloop.controllers`), 40
`GaussianProcessLearner` (class in `mloop.learners`), 47
`GaussianProcessVisualizer` (class in `mloop.visualizations`), 73
`generate_filename_suffix()` (in module `mloop.utilities`), 70
`generate_population()` (`mloop.learners.DifferentialEvolutionLearner` method), 47
`generation_num` (`mloop.learners.GaussianProcessLearner` attribute), 49
`generation_num` (`mloop.learners.NeuralNetLearner` attribute), 62
`get_controller_type_from_learner_archive()` (in module `mloop.utilities`), 70
`get_cost_dict()` (`mloop.testing.TestLandscape` method), 66
`get_dict_from_file()` (in module `mloop.utilities`), 70
`get_file_type()` (in module `mloop.utilities`), 71

`get_losses()` (*mloop.learners.NeuralNetLearner* method), 63
`get_next_cost_dict()` (*mloop.interfaces.FileInterface* method), 42
`get_next_cost_dict()` (*mloop.interfaces.Interface* method), 43
`get_next_cost_dict()` (*mloop.interfaces.ShellInterface* method), 44
`get_next_cost_dict()` (*mloop.interfaces.TestInterface* method), 44
`get_params_and_costs()` (*mloop.learners.MachineLearner* method), 59
`get_regularization_histories()` (*mloop.learners.NeuralNetLearner* method), 63

H

`has_trust_region()` (*mloop.learners.DifferentialEvolutionLearner* attribute), 46
`has_trust_region()` (*mloop.learners.GaussianProcessLearner* attribute), 50
`has_trust_region()` (*mloop.learners.MachineLearner* attribute), 58
`has_trust_region()` (*mloop.learners.NeuralNetLearner* attribute), 62

I

`import_neural_net()` (*mloop.learners.NeuralNetLearner* method), 63
`in_costs()` (*mloop.controllers.Controller* attribute), 37
`in_uncers()` (*mloop.controllers.Controller* attribute), 38
`init_simplex_corner()` (*mloop.learners.NelderMeadLearner* attribute), 60
`init_simplex_disp()` (*mloop.learners.NelderMeadLearner* attribute), 60
`init_std()` (*mloop.learners.DifferentialEvolutionLearner* attribute), 46
`Interface` (class in *mloop.interfaces*), 42
`interface_error_queue()` (*mloop.controllers.Controller* attribute), 37
`InterfaceInterrupt`, 43

L

`launch_extras()` (in module *mloop.launchers*), 45
`launch_from_file()` (in module *mloop.launchers*), 45

`Learner` (class in *mloop.learners*), 54
`learner` (*mloop.controllers.Controller* attribute), 37
`learner_costs_queue` (*mloop.controllers.Controller* attribute), 37
`learner_params_queue` (*mloop.controllers.Controller* attribute), 37
`LearnerInterrupt`, 56
`length_scale_history` (*mloop.learners.GaussianProcessLearner* attribute), 49

M

`MachineLearner` (class in *mloop.learners*), 56
`MachineLearnerController` (class in *mloop.controllers*), 40
`mloop` (module), 35
`mloop.controllers` (module), 36
`mloop.interfaces` (module), 42
`mloop.launchers` (module), 45
`mloop.learners` (module), 45
`mloop.testing` (module), 65
`mloop.utilities` (module), 67
`mloop.visualizations` (module), 72
`mutate()` (*mloop.learners.DifferentialEvolutionLearner* method), 47

N

`NelderMeadController` (class in *mloop.controllers*), 41
`NelderMeadLearner` (class in *mloop.learners*), 59
`NelderMeadVisualizer` (class in *mloop.visualizations*), 76
`neural_net` (*mloop.learners.NeuralNetLearner* attribute), 62
`NeuralNetController` (class in *mloop.controllers*), 41
`NeuralNetLearner` (class in *mloop.learners*), 60
`NeuralNetVisualizer` (class in *mloop.visualizations*), 76
`next_generation()` (*mloop.learners.DifferentialEvolutionLearner* method), 47
`noise_level_history` (*mloop.learners.GaussianProcessLearner* attribute), 50
`NullQueueListener` (class in *mloop.utilities*), 67
`num_in_costs` (*mloop.controllers.Controller* attribute), 37
`num_out_params` (*mloop.controllers.Controller* attribute), 37
`num_population_members` (*mloop.learners.DifferentialEvolutionLearner* attribute), 49

attribute), 46

O

`optimize()` (*mloop.controllers.Controller* method), 39

`out_extras` (*mloop.controllers.Controller* attribute), 37

`out_params` (*mloop.controllers.Controller* attribute), 37

`OUT_TYPE` (*mloop.learners.DifferentialEvolutionLearner* attribute), 46

`OUT_TYPE` (*mloop.learners.GaussianProcessLearner* attribute), 50

`OUT_TYPE` (*mloop.learners.Learner* attribute), 55

`OUT_TYPE` (*mloop.learners.NelderMeadLearner* attribute), 60

`OUT_TYPE` (*mloop.learners.NeuralNetLearner* attribute), 62

`OUT_TYPE` (*mloop.learners.RandomLearner* attribute), 65

P

`ParameterScaler` (class in *mloop.utilities*), 67

`params_count` (*mloop.learners.GaussianProcessLearner* attribute), 50

`params_count` (*mloop.learners.MachineLearner* attribute), 58

`params_count` (*mloop.learners.NeuralNetLearner* attribute), 62

`params_generations` (*mloop.learners.DifferentialEvolutionLearner* attribute), 46

`params_out_queue` (*mloop.controllers.Controller* attribute), 37

`params_out_queue` (*mloop.learners.Learner* attribute), 54

`params_scaler` (*mloop.learners.GaussianProcessLearner* attribute), 50

`partial_fit()` (*mloop.utilities.ParameterScaler* method), 68

`plot_cost_vs_run()` (*mloop.visualizations.ControllerVisualizer* method), 72

`plot_costs_vs_generations()` (*mloop.visualizations.DifferentialEvolutionVisualizer* method), 73

`plot_cross_sections()` (*mloop.visualizations.GaussianProcessVisualizer* method), 74

`plot_density_surface()` (*mloop.visualizations.NeuralNetVisualizer* method), 77

`plot_hyperparameters_vs_fit()` (*mloop.visualizations.GaussianProcessVisualizer* method), 74

`plot_hyperparameters_vs_run()` (*mloop.visualizations.GaussianProcessVisualizer* method), 75

`plot_losses()` (*mloop.visualizations.NeuralNetVisualizer* method), 77

`plot_noise_level_vs_fit()` (*mloop.visualizations.GaussianProcessVisualizer* method), 75

`plot_noise_level_vs_run()` (*mloop.visualizations.GaussianProcessVisualizer* method), 75

`plot_parameters_vs_cost()` (*mloop.visualizations.ControllerVisualizer* method), 72

`plot_parameters_vs_run()` (*mloop.visualizations.ControllerVisualizer* method), 73

`plot_params_vs_generations()` (*mloop.visualizations.DifferentialEvolutionVisualizer* method), 73

`plot_regularization_history()` (*mloop.visualizations.NeuralNetVisualizer* method), 77

`plot_surface()` (*mloop.visualizations.NeuralNetVisualizer* method), 77

`predict_biased_cost()` (*mloop.learners.GaussianProcessLearner* method), 52

`predict_cost()` (*mloop.learners.GaussianProcessLearner* method), 53

`predict_cost()` (*mloop.learners.NeuralNetLearner* method), 63

`predict_cost_gradient()` (*mloop.learners.NeuralNetLearner* method), 64

`predict_costs_from_param_array()` (*mloop.learners.NeuralNetLearner* method), 64

`predicted_best_cost` (*mloop.learners.GaussianProcessLearner* attribute), 52

`predicted_best_cost` (*mloop.learners.NeuralNetLearner* attribute), 63

`predicted_best_parameters` (*mloop.learners.GaussianProcessLearner* attribute), 52

`predicted_best_parameters` (*mloop.learners.NeuralNetLearner* attribute), 63

`predicted_best_uncertainty` (*mloop.learners.GaussianProcessLearner* attribute), 52

`print_results()` (*mloop.controllers.Controller*

method), 40
`print_results()` (*mloop.controllers.MachineLearnerController* *tribute*), 61
method), 41
`put_params_and_get_cost()`
(*mloop.learners.Learner* *method*), 56

R

`random_index_sample()`
(*mloop.learners.DifferentialEvolutionLearner* *method*), 47
`RandomController` (*class in mloop.controllers*), 41
`RandomLearner` (*class in mloop.learners*), 65
`RandomVisualizer` (*class in mloop.visualizations*), 78
`return_cross_sections()`
(*mloop.visualizations.GaussianProcessVisualizer* *method*), 75
`return_cross_sections()`
(*mloop.visualizations.NeuralNetVisualizer* *method*), 77
`run()` (*mloop.interfaces.Interface* *method*), 43
`run()` (*mloop.learners.DifferentialEvolutionLearner* *method*), 47
`run()` (*mloop.learners.GaussianProcessLearner* *method*), 53
`run()` (*mloop.learners.NelderMeadLearner* *method*), 60
`run()` (*mloop.learners.NeuralNetLearner* *method*), 65
`run()` (*mloop.learners.RandomLearner* *method*), 65
`run()` (*mloop.testing.FakeExperiment* *method*), 66
`run()` (*mloop.visualizations.GaussianProcessVisualizer* *method*), 76
`run()` (*mloop.visualizations.NeuralNetVisualizer* *method*), 78

S

`safe_cast_to_array()` (*in module mloop.utilities*), 71
`safe_cast_to_list()` (*in module mloop.utilities*), 71
`save_archive()` (*mloop.controllers.Controller* *method*), 40
`save_archive()` (*mloop.learners.Learner* *method*), 56
`save_dict_to_file()` (*in module mloop.utilities*), 71
`save_generation()`
(*mloop.learners.DifferentialEvolutionLearner* *method*), 47
`scaled_costs` (*mloop.learners.GaussianProcessLearner* *attribute*), 49
`scaled_costs` (*mloop.learners.MachineLearner* *attribute*), 57

`scaled_costs` (*mloop.learners.NeuralNetLearner* *attribute*), 61
`set_bad_region()` (*mloop.testing.TestLandscape* *method*), 66
`set_default_landscape()`
(*mloop.testing.TestLandscape* *method*), 66
`set_landscape()` (*mloop.testing.FakeExperiment* *method*), 66
`set_legend_location()` (*in module mloop.visualizations*), 80
`set_noise_function()`
(*mloop.testing.TestLandscape* *method*), 66
`set_quadratic_landscape()`
(*mloop.testing.TestLandscape* *method*), 67
`set_random_quadratic_landscape()`
(*mloop.testing.TestLandscape* *method*), 67
`ShellInterface` (*class in mloop.interfaces*), 43
`show_all_default_visualizations()` (*in module mloop.visualizations*), 80
`show_all_default_visualizations_from_archive()`
(*in module mloop.visualizations*), 80
`simplex_costs` (*mloop.learners.NelderMeadLearner* *attribute*), 60
`simplex_params` (*mloop.learners.NelderMeadLearner* *attribute*), 60
`start()` (*mloop.utilities.NullQueueListener* *method*), 67
`stop()` (*mloop.utilities.NullQueueListener* *method*), 67

T

`TestInterface` (*class in mloop.interfaces*), 44
`TestLandscape` (*class in mloop.testing*), 66
`txt_file_to_dict()` (*in module mloop.utilities*), 71

U

`update_archive()` (*mloop.learners.DifferentialEvolutionLearner* *method*), 47
`update_archive()` (*mloop.learners.GaussianProcessLearner* *method*), 54
`update_archive()` (*mloop.learners.Learner* *method*), 56
`update_archive()` (*mloop.learners.MachineLearner* *method*), 59
`update_archive()` (*mloop.learners.NelderMeadLearner* *method*), 60
`update_archive()` (*mloop.learners.NeuralNetLearner* *method*), 65
`update_bads()` (*mloop.learners.MachineLearner* *method*), 59
`update_bias_function()`
(*mloop.learners.GaussianProcessLearner* *method*), 54
`update_search_params()`
(*mloop.learners.MachineLearner* *method*),

59
`update_search_region()`
 (*mloop.learners.MachineLearner* *method*),
 59

W

`wait_for_new_params_event()`
 (*mloop.learners.MachineLearner* *method*),
 59
`worst_cost (mloop.learners.GaussianProcessLearner`
 attribute), 49
`worst_cost (mloop.learners.MachineLearner` *at-*
 tribute), 57
`worst_cost (mloop.learners.NeuralNetLearner`
 attribute), 62
`worst_index (mloop.learners.GaussianProcessLearner`
 attribute), 49
`worst_index (mloop.learners.MachineLearner`
 attribute), 58
`worst_index (mloop.learners.NeuralNetLearner` *at-*
 tribute), 62