

Bayesian optimization with Scikit-Optimize

Gilles Louppe

@glouppe



PyData Amsterdam 2017

Pause café?

You have an espresso machine with many buttons and knobs to tweak.

Your task is to **brew the best cup of espresso** before dying of caffeine overdose.



Credits: [Steam punk coffee machine](#)

$$x^* = \arg \max_x f(x)$$


- f is a black box function, with no closed form nor gradients.
 - f is expensive to evaluate.
 - You may only have noisy observations of f .
-

If you do not have
these constraints,
do not use Bayesian
optimization.

Bayesian optimisation, step-by-step


Start with observations of the objective f :

$$\{(x_i, f(x_i)) \mid i = 1, \dots, t\}$$



Bayesian optimisation, step-by-step


Build a probabilistic model for f
(typically, a Gaussian process).



Bayesian optimisation, step-by-step


Optimize a cheap utility function u based on the posterior distribution for sampling the next point.

$$x_{t+1} = \arg \max_x u(x)$$



Bayesian optimisation, step-by-step

Evaluate f and repeat.



What is Bayesian about Bayesian optimization?

- The unknown objective is considered as a *random function* (a stochastic process) on which we place a **prior** (here defined by a Gaussian process capturing our beliefs about the function behaviour).
- Function evaluations are treated as data and used to update the prior to form the **posterior** distribution over the objective function.



$$P(A|B) = \frac{P(B|A)P(A)}{P(B)}$$

Scikit-Optimize

A simple library for black-box optimization with a `scipy.optimize` interface.

Sprouted from a (never-ending) Scikit-Learn [pull request](#).

<https://scikit-optimize.github.io>

```
pip install scikit-optimize
```


Minimize

- Takes the objective function and the bounds of the parameter space.
- That's it!
- The `r` tuple contains all the results (final and intermediate).

```
from skopt import gp_minimize

space = [(-5.0, 0.0),      # learning_rate
         (1, 5),          # max_depth
         (2, 100),        # min_samples_split
         (1, X.shape[1])] # max_features

r = gp_minimize(objective, space, n_calls=50, random_state=1)
```

```
r.x
```

```
[-0.96542858133704623, 5, 9, 7]
```

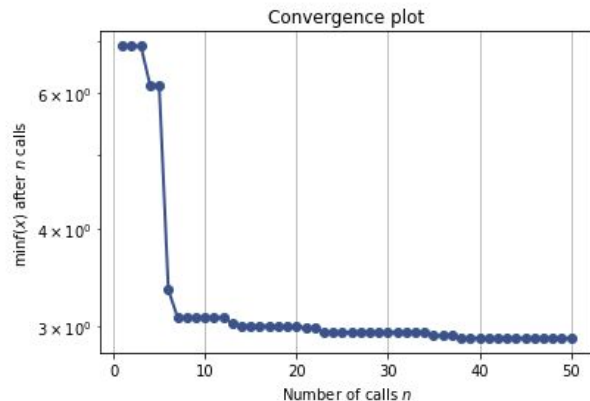
```
r.fun
```

```
2.8944939733595256
```


Visualization

- Convergence plot
- Pairwise partial dependence plot of the surrogate objective.

```
from skopt.plots import plot_convergence
plot_convergence(r, yscale="log")
```



```
from skopt.plots import plot_objective
plot_objective(r)
```



Acquisition function

Specifies the next sample x .

- Upper confidence bound

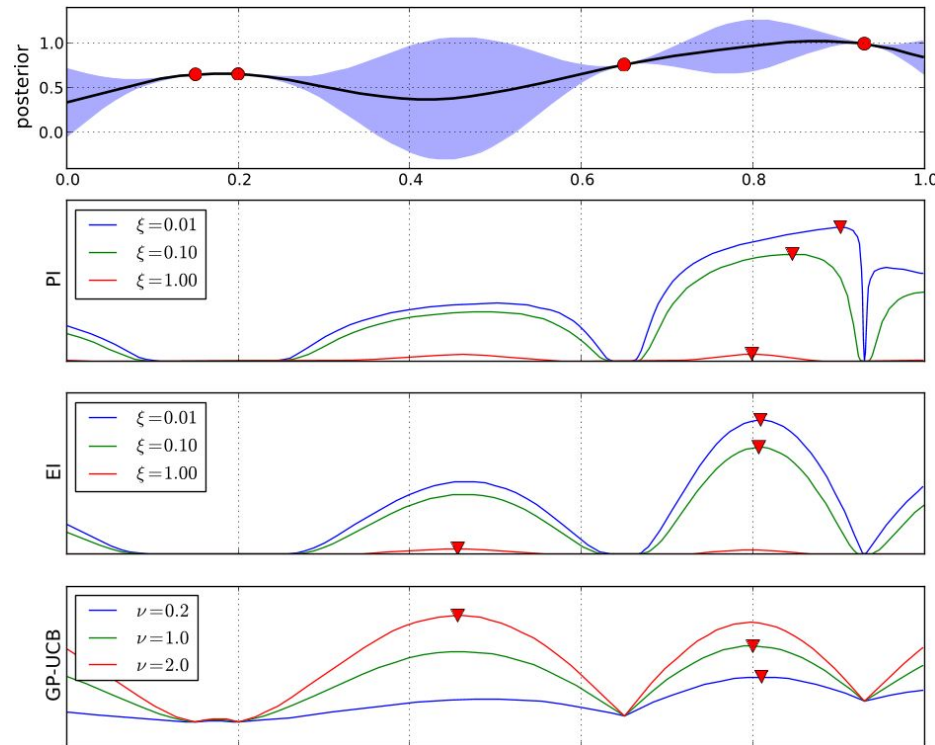
$$UCB(x) = \mu_{GP}(x) + \kappa\sigma_{GP}(x)$$

- Probability of improvement

$$PI(x) = P(f(x) \geq f(x_t^+) + \kappa)$$

- Expected improvement

$$EI(x) = \mathbb{E}[f(x) - f(x_t^+)]$$



κ provides a knob for controlling the exploration-exploitation trade-off.


Surrogate model

The probabilistic model for f is usually built as a Gaussian Process.

Scikit-Optimize supports any Scikit-Learn regressor that can also return the variance of the predictions (`return_std=True`).

- Random forests / Extra-trees
- Gradient boosting

Tree-based optimization is fast and usually better on discontinuous high-dimensional spaces.



*Random forests as
a probabilistic model.*

```
from skopt import forest_minimize
r = forest_minimize(objective, space)
```

```
r.x
```

```
[-0.89303546067450945, 5, 33, 9]
```

```
r.fun
```

```
2.8175239612455112
```

Is this better than
random?

Was this actually worth it?

Random search

- If nothing works, try random search.
- Works surprisingly well, often only slightly worse than “smart” algorithms.

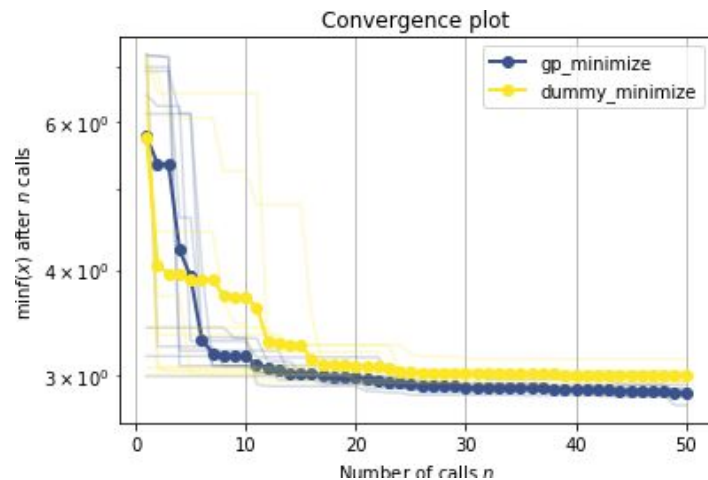
```
from skopt import dummy_minimize
r_dummy = dummy_minimize(objective, space, n_calls=50)
```

```
r_dummy.x
```

```
[-0.46028641219053501, 4, 80, 11]
```

```
r_dummy.fun
```


```
3.1180178726663654
```



Out!

Wait.

That's nice, but how do I put
my coffee maker into a
Python function?



Ask and tell API

- Factor out the evaluation of the objective function.
- `ask`: query next point to evaluate
`tell`: provide function value
- Resume long optimization by pickling/unpickling the `Optimizer` object.

```
from skopt import Optimizer
opt = Optimizer(space, gp)
```

```
# Next query point
next_x = opt.ask()
next_x
```

```
[-0.69339327310821286, 3, 74, 3]
```

```
# Tell value at x
r = opt.tell(next_x, objective(next_x))
```

```
# Shortcut for ask() + tell()
r = opt.run(objective, n_iter=50)
```

```
r.x
```

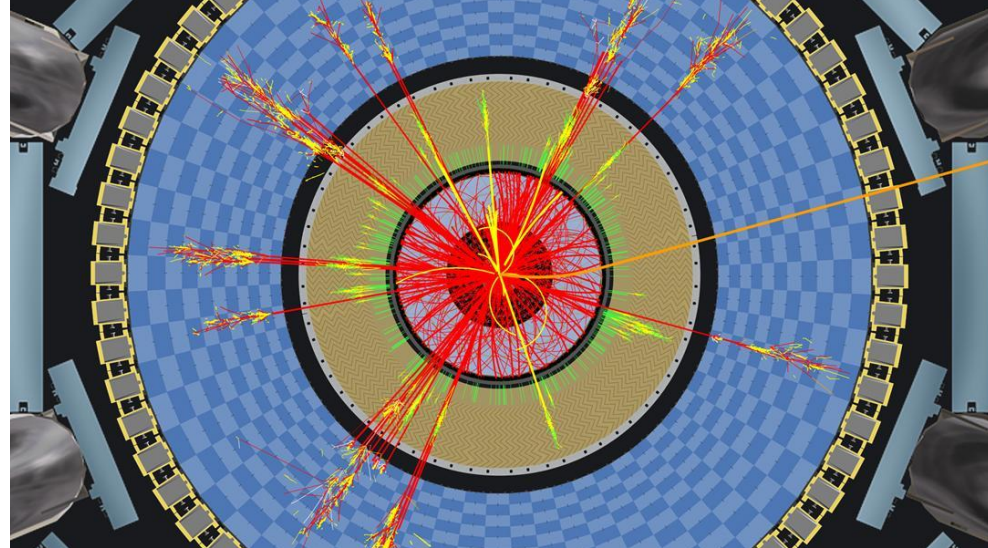


```
[-0.867784969421451, 5, 37, 9]
```

```
r.fun
```

```
2.736335329599823
```

Beyond coffee (or ML)

Tuning high energy physics
simulators to match
experimental data
(e.g., [1610.08328](#))



Good coffee is expensive
and does not come with
gradients.

**Bayesian optimisation
can help.**

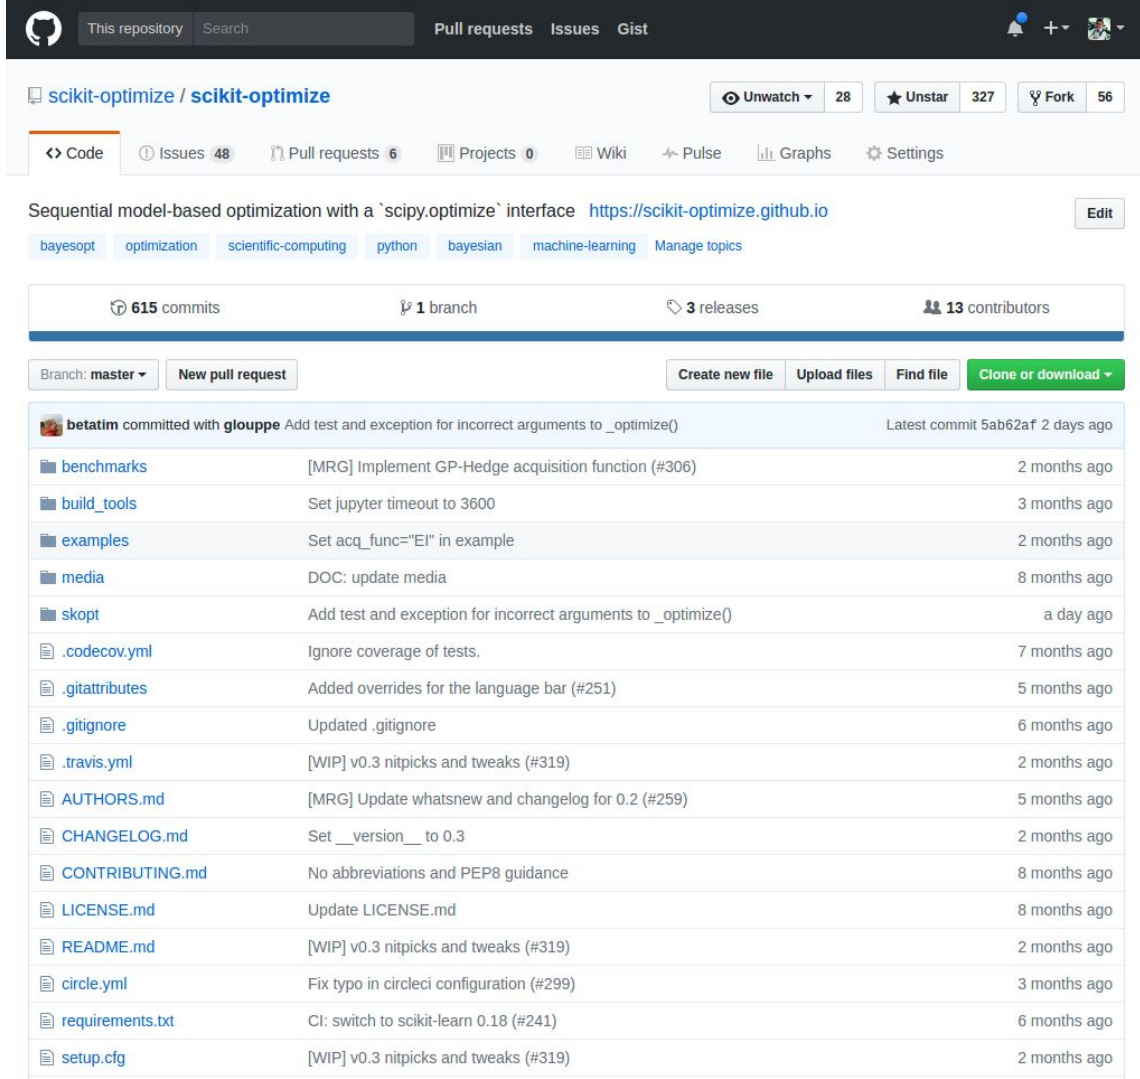
Summary

- Bayesian optimisation is a principled approach for optimising an expensive function f .
- Scikit-Optimize provides an easy-to-use set of tools for this.

We want you!

Help us improve Scikit-Optimize

<https://github.com/scikit-optimize/scikit-optimize>



The screenshot shows the GitHub repository page for `scikit-optimize / scikit-optimize`. At the top, there are navigation links for `Code`, `Issues 48`, `Pull requests 6`, `Projects 0`, `Wiki`, `Pulse`, `Graphs`, and `Settings`. Below this, the repository description reads: "Sequential model-based optimization with a `scipy.optimize` interface" with a link to `https://scikit-optimize.github.io`. There are also tags for `bayesopt`, `optimization`, `scientific-computing`, `python`, `bayesian`, and `machine-learning`.

Statistics for the repository include: `615` commits, `1` branch, `3` releases, and `13` contributors. A navigation bar shows the current branch is `master` and provides options to `Create new file`, `Upload files`, `Find file`, and `Clone or download`.

The commit history table is as follows:

Commit	Description	Time
<code>betatim</code> committed with <code>gloupe</code>	Add test and exception for incorrect arguments to <code>_optimize()</code>	Latest commit 5ab62af 2 days ago
<code>benchmarks</code>	[MRG] Implement GP-Hedge acquisition function (#306)	2 months ago
<code>build_tools</code>	Set jupyter timeout to 3600	3 months ago
<code>examples</code>	Set <code>acq_func="EI"</code> in example	2 months ago
<code>media</code>	DOC: update media	8 months ago
<code>skopt</code>	Add test and exception for incorrect arguments to <code>_optimize()</code>	a day ago
<code>.codecov.yml</code>	Ignore coverage of tests.	7 months ago
<code>.gitattributes</code>	Added overrides for the language bar (#251)	5 months ago
<code>.gitignore</code>	Updated <code>.gitignore</code>	6 months ago
<code>.travis.yml</code>	[WIP] v0.3 nitpicks and tweaks (#319)	2 months ago
<code>AUTHORS.md</code>	[MRG] Update whatsnew and changelog for 0.2 (#259)	5 months ago
<code>CHANGELOG.md</code>	Set <code>__version__</code> to 0.3	2 months ago
<code>CONTRIBUTING.md</code>	No abbreviations and PEP8 guidance	8 months ago
<code>LICENSE.md</code>	Update <code>LICENSE.md</code>	8 months ago
<code>README.md</code>	[WIP] v0.3 nitpicks and tweaks (#319)	2 months ago
<code>circle.yml</code>	Fix typo in circleci configuration (#299)	3 months ago
<code>requirements.txt</code>	CI: switch to scikit-learn 0.18 (#241)	6 months ago
<code>setup.cfg</code>	[WIP] v0.3 nitpicks and tweaks (#319)	2 months ago