

# Towards an Empirical Foundation for Assessing Bayesian Optimization of Hyperparameters

Katharina Eggenberger, Matthias Feurer, Frank Hutter  
Freiburg University  
{eggenpk, feurerm, fh}@informatik.uni-freiburg.de

James Bergstra  
University of Waterloo  
james.bergstra@uwaterloo.ca


Jasper Snoek  
Harvard University  
jsnoek@seas.harvard.edu

Holger H. Hoos and Kevin Leyton-Brown  
University of British Columbia  
{hoos, kevinlb}@cs.ubc.ca



UNI  
FREIBURG

UNIVERSITY OF  
WATERLOO



... in 15 sec

**Bayesian Optimization** techniques are becoming readily applicable for hyperparameter tuning.

Many approaches exist, but it remains unclear which one performs best in which case.

For further progress we need an easy-to-use **benchmark library** and a **baseline comparison**. We provide both.

Contribute

Our **publicly available** software wraps various benchmarks and the Bayesian optimizers SMAC, Spearmint and TPE. This is just a first step towards a canonical benchmark library. You can help us:

- **Contribute your algorithm** to add another benchmark for future research
- **Include your optimizer** to find out its strengths and weaknesses

Ready To Use

**SMAC** (Sequential Model-based algorithm configuration) is based on random forests and can handle continuous, discrete and conditional hyperparameters.

[Hutter, Hoos, and Leyton-Brown, 2011]

**Spearmint** uses Gaussian Process (GP) models and performs slice sampling over the GP's hyperparameters. Can handle continuous and discrete hyperparameters.

[Snoek, Larochelle, and Adams, 2012]

**TPE** (Tree Parzen Estimator) is based on Gaussian Mixture Models. Supports conditional, continuous and discrete parameters and also priors over them.

[Bergstra, Bardenet, Bengio, and Kégl, 2011]

Add your optimizer, executable from command line interface

```
>python wrapping.py <optimizer> <benchmark> [-s <seed>]
```

Bayesian optimization algorithm

Searchspace

hyperparameter setting, CV fold

Wrapper: limits memory & time of this evaluation

Loss achieved, time spent

Benchmark

Machine learning algorithm

Data

Empirical Results

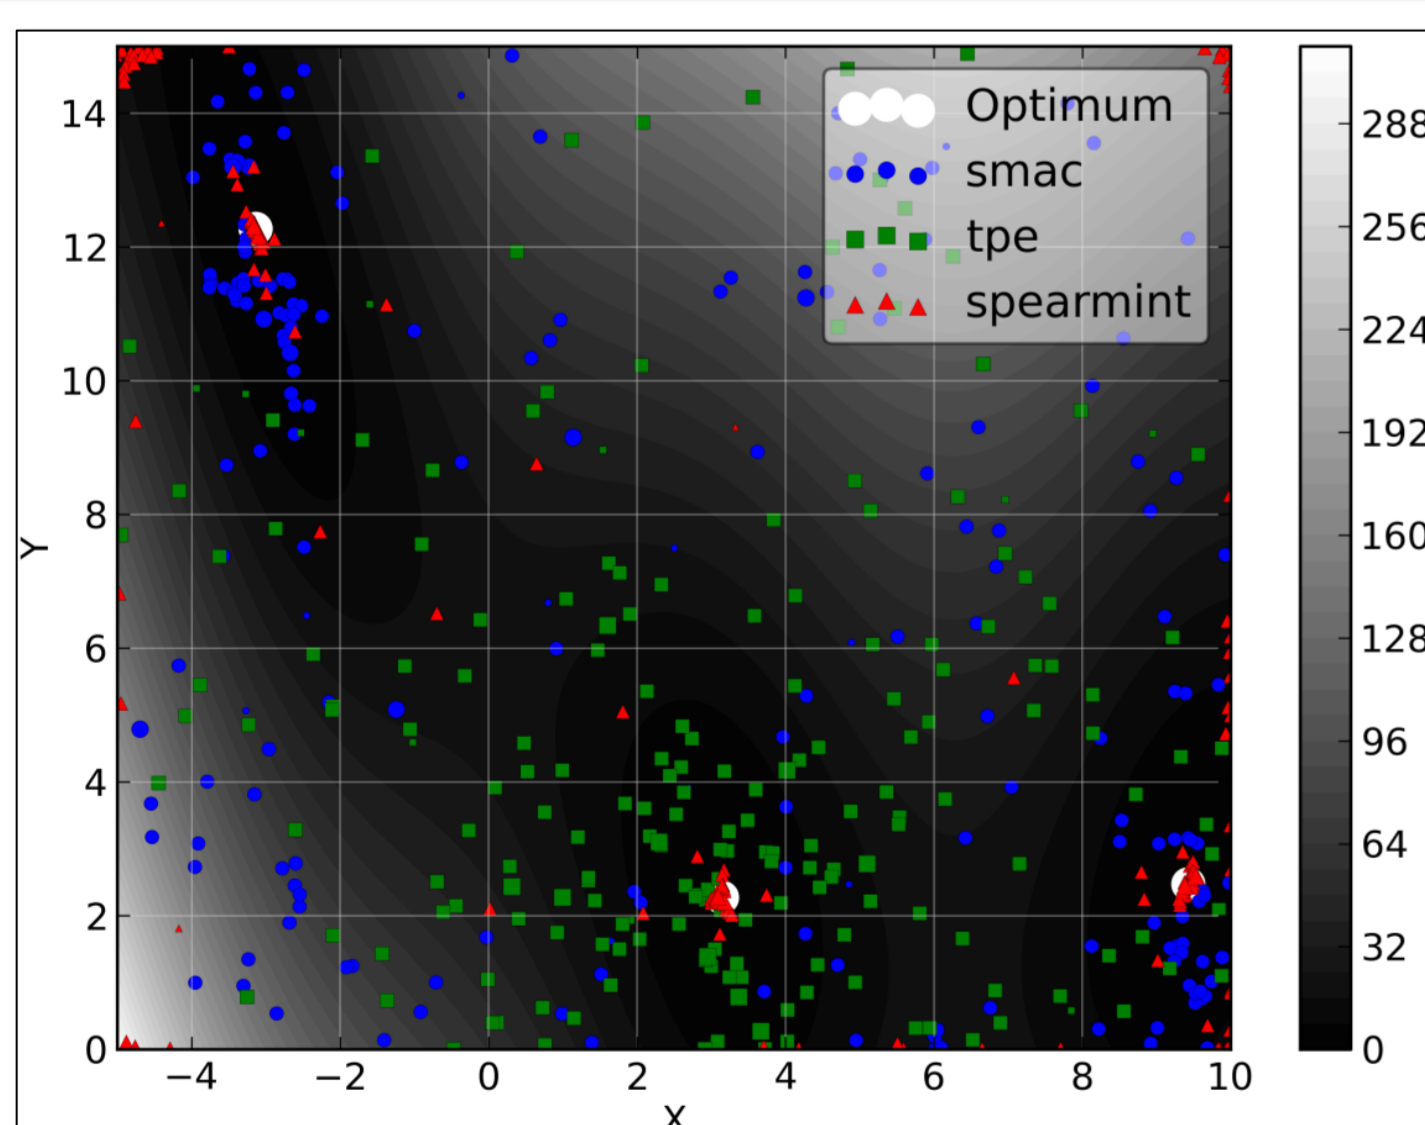



Fig. 1: Evaluated points for the 2-dimensional branin test function. The function's three global optima are visualized as white circles.

Fig. 2: Representative traces of loss function values over time when optimizing the 38 hyperparameters of the HP-NNET on the convex dataset.




Mean loss (typically: validation error rate) and standard deviation across 10 runs of each optimizer. Bold face indicates best value, underlined values are not statistically significant worse than the best value.

Experiment	#evals	SMAC	Spearmint	TPE
<b>Branin (0.398)</b>	200	0.655 ± 0.27	<b>0.398</b> ± 0.0	0.526 ± 0.12
<b>Har6 (-3.322)</b>	200	-2.98 ± 0.11	<b>-3.13</b> ± 0.41	-2.82 ± 0.18
Log. Regr.	100	8.6 ± 0.9	<u>7.3</u> ± 0.2	8.2 ± 0.6
LDA ongrid	50	<u>1270</u> ± 3	<u>1273</u> ± 10	<u>1272</u> ± 4
SVM ongrid	100	<b>24.2</b> ± 0.1	24.6 ± 0.9	<b>24.2</b> ± 0.0
HP-NNET	100	<u>19.5</u> ± 1.5	20.6 ± 0.3	<b>19.5</b> ± 1.6
HP-NNET	200	<u>18.3</u> ± 1.9	20.0 ± 0.9	<u>18.5</u> ± 1.4
HP-NNET	100	<u>51.5</u> ± 2.8	<u>52.2</u> ± 3.3	<b>50.0</b> ± 1.7
HP-NNET	200	<b>48.3</b> ± 1.8	51.4 ± 3.2	<b>48.9</b> ± 1.4
HP-DBNET	100	<u>16.4</u> ± 1.2	<u>20.74</u> ± 6.9	<u>17.29</u> ± 1.7
HP-DBNET	200	<u>15.4</u> ± 0.8	<u>17.45</u> ± 5.6	16.1 ± 0.5
Auto-WEKA	30h	<u>27.49</u> ± 4.9	40.64 ± 7.2	35.33 ± 2.93
Log. Regr. 5CV	500 folds	<u>8.1</u> ± 0.2	<u>8.2</u> ± 0.1	8.9 ± 0.5
HP-NNET 5CV	500 folds	<b>18.2</b> ± 1.5	23.0 ± 5.0	20.9 ± 1.3
HP-NNET 5CV	500 folds	<b>47.9</b> ± 0.7	52.8 ± 5.1	50.8 ± 1.4

Ready To Use

# hyp.params (conditional)	continuous /discrete	Dataset	Citation
2 (-)	2 / -	-	[Hedar]
6 (-)	6 / -	-	
4 (-)	4 / -	MNIST	[Snoek, Larochelle, and Adams, 2012]
3 (-)	- / 3	Wikipedia Articles	[LeCun, Bottou, Bengio, and Haffner, 1998]
3 (-)	- / 3	UniProbe	[Hoffman, Blei, and Bach, 2010]
14 (4)	7 / 7	convex	[Miller, Kumer, Packer, Goodman, and Koller, 2012]
14 (4)	7 / 7	MRBI	[Bergstra, Bardenet, Bengio, and Kégl, 2011]
36 (27)	19 / 17	convex	[Larochelle, Erhan, Courville, Bergstra, and Bengio, 2007]
786 (784)	296 / 490	convex	[Thornton, Hutter, Hoos, and Leyton-Brown, 2012]
4 (-)	4 / -	MNIST	[Larochelle, Erhan, Courville, Bergstra, and Bengio, 2007]
14 (4)	7 / 7	convex	[Hall, Frank, Holmes, Pfahringer, Reutemann, and Witten, 2009]
14 (4)	7 / 7	MRBI	[Snoek, Larochelle, and Adams, 2012]
14 (4)	7 / 7	convex	[Bergstra, Bardenet, Bengio, and Kégl, 2011]
14 (4)	7 / 7	MRBI	[Larochelle, Erhan, Courville, Bergstra, and Bengio, 2007]

Add your algorithm, executable from command line interface



Implementation Publicly Available:

[www.automl.org/hpolib](http://www.automl.org/hpolib)