

HPOBench: A Collection of Reproducible Multi-Fidelity Benchmark Problems for HPO



**Katharina
Eggensperger**



**Philipp
Müller**



**Neeratyoy
Mallik**



**Matthias
Feurer**



**René
Sass**



**Aaron
Klein**




**Noor
Awad**



**Marius
Lindauer**



**Frank
Hutter**



[/automl/HPOBench](https://github.com/automl/HPOBench)

Why should you care?

Available data is growing.

Applications of ML are growing.

Model size and complexity is growing.

We need efficient hyperparameter optimization methods!

→ **Multi-fidelity optimization**

BUT


To **develop, improve, understand** and **compare** methods we need benchmark problems that are **realistic, efficient** and **available for a long time**.

→ **HPOBench**




Photo by [frank.mckenna](#) on [Unsplash](#)

- The first collection of **containerized multi-fidelity** HPO benchmarks with **100+ benchmark problems**



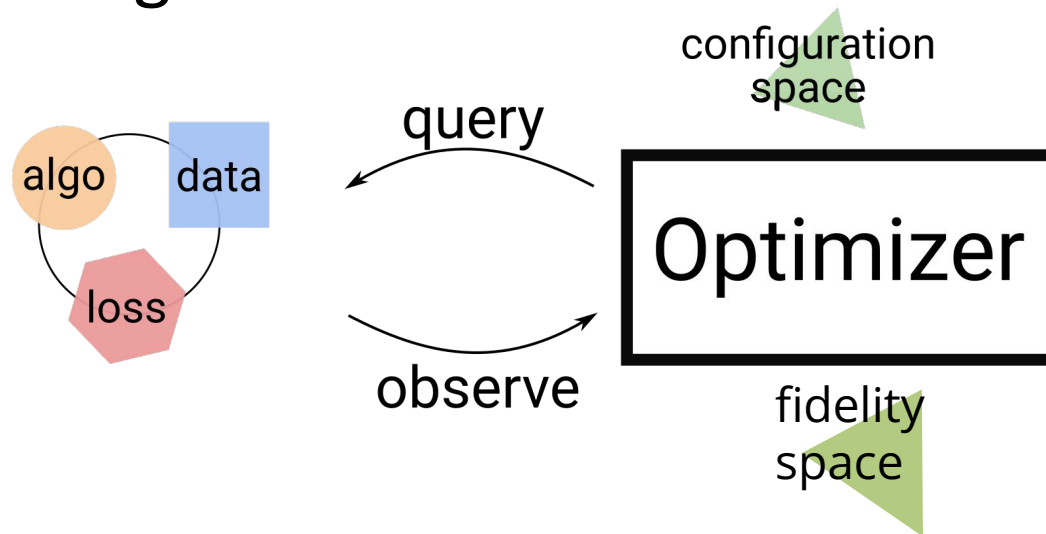
- The first set of HPO benchmarks that
 - are available as **raw and tabular** versions
 - which also support **multi-objective optimization** and **transfer-HPO across datasets**

- An exemplary **large-scale study** evaluating **>10 optimization methods** on all benchmarks



HPO Benchmarks

Benchmark ingredients:



Ideal features:

Efficiency



Photo by [Federico Bottos](#) on [Unsplash](#)

Reproducibility



Photo by [frank mckenna](#) on [Unsplash](#)


Flexibility




Photo by [Edz Norton](#) on [Unsplash](#)

Efficiency

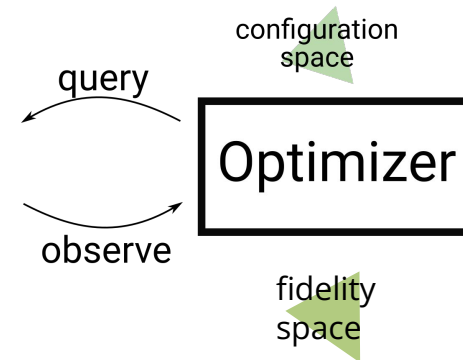
raw
benchmark




tabular
benchmark



surrogate
benchmark





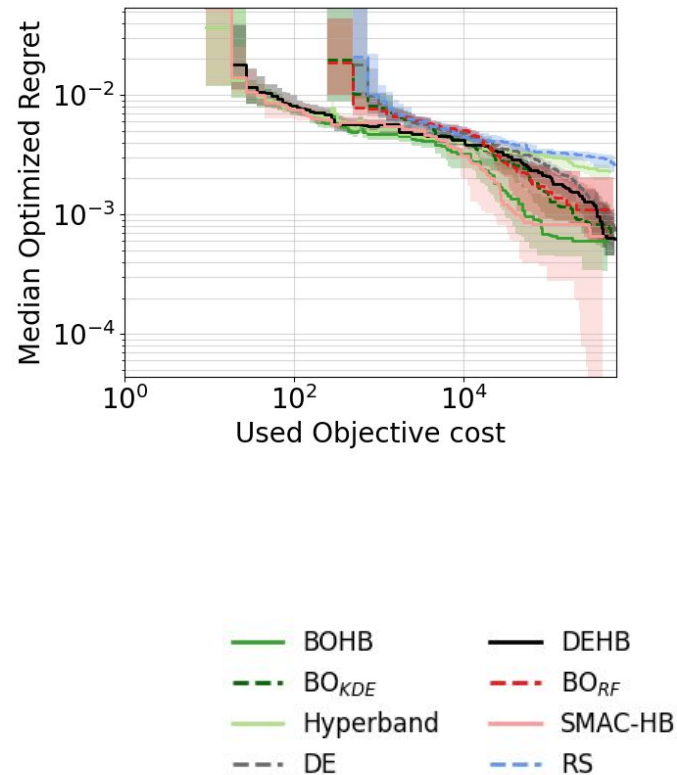
Family	#benchs	#cont(log)	#int(log)	#cat	#ord	fidelity	type
<i>Cartpole</i>	1	4(1)	3(3)	-	-	repetitions	✓
<i>BNN</i>	2	3(1)	2(2)	-	-	samples	✓
<i>Net</i>	6	5	1	-	-	time	✗
<i>NBHPO</i>	4	-	-	3	6	epochs	(✓)
<i>NB101</i>	3	-	-	26	-	epochs	(✓)
		21	1	14	-		
<i>NB201</i>	3	-	-	6	-	epochs	(✓)
<i>NB1Shot1</i>	3	-	-	9	-	epochs	(✓)
		-	-	9	-		
		-	-	11	-		
<i>LogReg</i>	20	2(2)	-	-	-	iter	✓, (✓)
<i>SVM</i>	20	2(2)	-	-	-	data	✓, (✓)
<i>RandomForest</i>	20	1	3(2)	-	-	#trees	✓, (✓)
<i>XGBoost</i>	20	3(2)	1(1)	-	-	#trees	✓, (✓)
<i>MLP</i>	8	2(2)	3(2)	-	-	epochs	✓, (✓)

✓	raw
(✓)	tabular
✗	surrogate

We ran >10 optimization methods on all benchmarks and studied the following:

1. Do advanced methods improve over random baselines?
2. Do multi-fidelity methods improve over single-fidelity methods

→ Short answer: **Yes**




→ **HPOBench** provides **>100 containerized benchmarks** for **multi-fidelity HPO**

What else you can do with HPOBench:

- multi-objective optimization and transfer-HPO across datasets
- compare raw, tabular and surrogate benchmarks
- ...

Thank you!

 /AutoML_org

 /automl/HPOBench



**Katharina
Eggensperger¹**




**Philipp
Müller¹**



**Neeratyoy
Mallik¹**



**Matthias
Feurer¹**



You



**René
Sass²**



**Aaron
Klein³**



**Noor
Awad¹**



**Marius
Lindauer²**



**Frank
Hutter^{1,4}**

¹ Albert-Ludwigs-Universität Freiburg

² Leibniz Universität Hannover

³ Amazon (work done prior to joining Amazon)

⁴ Bosch Center for Artificial
Intelligence