

# Is One Epoch All You Need For Multi-Fidelity Hyperparameter Optimization?

Romain Egelé<sup>1,3</sup>, Isabelle Guyon<sup>1,2</sup>, Yixuan Sun<sup>3</sup>, Prasanna Balaprakash<sup>4</sup>

1- LISN, U. Paris-Saclay, France; 2- Google, USA & ChaLearn, USA;  
3- MCS, Argonne Ntl Lab, USA; 4- CCSD, Oak Ridge Ntl Lab, USA

## Abstract.

Hyperparameter optimization (HPO) is crucial for fine-tuning machine learning models, but it can be computationally expensive. To reduce costs, Multi-fidelity HPO (MF-HPO) leverages intermediate accuracy levels in the learning process and discards low-performing models early on. We conducted a comparison of various representative MF-HPO methods against a simple baseline on classical benchmark data. The baseline involved discarding all models except the Top- $K$  after training for only one epoch, followed by further training to select the best model. Surprisingly, this baseline achieved similar results to its counterparts, while requiring an order of magnitude less computation. Upon analyzing the learning curves of the benchmark data, we observed a few dominant learning curves, which explained the success of our baseline. This suggests that researchers should (1) always use the suggested baseline in benchmarks and (2) broaden the diversity of MF-HPO benchmarks to include more complex cases.

## 1 Introduction

Hyperparameter optimization (HPO) is a key component of AutoML systems. It aims to find the best configuration of a machine learning (ML) pipeline, which consists of data and model components. Hyperparameters (HP) are parameters that cannot be learned directly with the primary learning algorithm (e.g., gradient descent), but affect the learning process and the performance of the pipeline. Those may include number of layers and units per layer in deep networks, learning rates, etc. HPO is usually formulated as a black-box optimization problem[1], where a function maps a HP configuration to a performance score. However, evaluating this function can be costly and time-consuming. Therefore, Multi-fidelity HPO (MF-HPO) methods have been proposed [2, 3], which use intermediate learning machine performance estimates to reduce the computational cost and speed up the optimization process. In multi-fidelity algorithms, fidelity refers to the level of accuracy or resolution at which a machine learning model is trained and evaluated. Different levels of fidelity can be used to balance the trade-off between computational cost and accuracy. For example, a low fidelity level could involve training and evaluating a model on a small subset of the data, while a high fidelity level would involve using the entire dataset. Multi-fidelity algorithms use this concept to optimize hyperparameters and other model parameters efficiently by exploring the model space using a combination of different fidelity levels. The accuracy obtained at each fidelity level is used to guide

the optimization process towards better performing models while minimizing the computational cost.

Fig. 1 presents the generic components of a multi-fidelity HPO. The input to HPO is the HP search space. The output is a trained model with the corresponding hyperparameters. HPO is a bi-level optimization problem: the upper level (**outer loop**) optimizes the HP of the ML pipeline while the lower level (**inner loop**) optimizes the parameters of the learning machine, given the HP configuration of the upper level. After completing the search process, a final step of **model selection** is performed before returning the trained model. In regular HPO, the inner loop is halted by a given stopping criterion. A natural improvement is to consider the multi-fidelity setting,

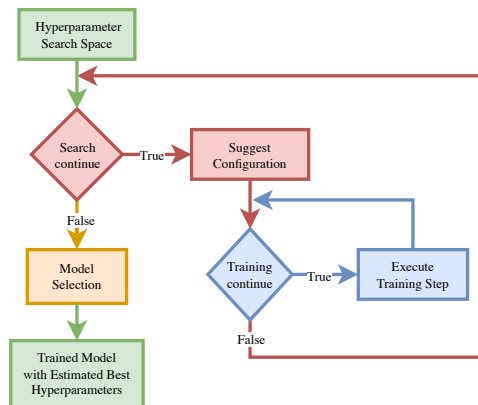


Fig. 1: Hyperparameter optimization and its components including **input/output**, **outer optimization loop exploring new HP configurations**, **inner optimization loop incrementally allocating new training “steps”** and **model selection**.

where intermediate performance estimates are obtained at different stages of the inner optimization process (as a function of the number of training examples or training epochs). This involves monitoring a “learning curve.”

This study evaluates the effectiveness of early discarding policies in the **inner loop**, while random search is used for the **outer loop**. We use widely-used approaches for early discarding, performing denovo search without prior knowledge of learning curve typology or meta-learning involvement.

We consider two baselines at each extreme. The first evaluates each HP configuration candidate at maximum fidelity, i.e., after being fully trained within the overall time budget (100-Epochs). The second evaluates them at minimum fidelity, after only one epoch of training. We demonstrate the effectiveness of the 1-Epoch approach and explain why it is successful. Our findings provide a new perspective on previous studies that omitted this baseline. Only one study [4] included a similar baseline and reported results similar to ours. Therefore, we advocate for including the 1-Epoch baseline in further MF-HPO benchmarks.

## 2 Related Work

Our study focuses on methods, which train only a single model at a time, but keep all check-points for further reference. Early discarding means switching to training a model with another HP configuration, before attaining the maximum number of epochs allowed. A typical example of such strategy, some-

times referred to as “vertical” model selection [5] is Asynchronous Successive Halving [6] (SHA). Several methods can be adapted to this setting, including Hyperband [3], which can explore different levels of fidelity to differentiate noisy learning curves; Learning Curve Extrapolation [7] (LCE), which can observe early performances and extrapolate future performances to decide whether training should continue; FABOLAS [8], which learns correlations in the candidates’ ranking between different levels of fidelity; Bayesian Optimization Hyperband [9], which embeds Bayesian optimization in Hyperband to sample candidates more efficiently; Learning Curve with Support Vector Regression [10], which predicts final performance based on the configuration and early observations; Learning Curve with Bayesian neural network [11], which is similar to the previous method but switches the model with a Bayesian neural network; and Trace Aware Knowledge-Gradient [12], which leverages an observed curve to update the posterior distribution of a Gaussian process more efficiently.

The previous works had some limitations, which were primarily due to the assumptions made about the learning curve. For example, methods based on SHA assume that the discarded learning curves will not cross in the future, since only the Top- $K$  models are allowed to continue at any given step. In this context, models that start slowly are often discarded. This phenomenon is known as “short-horizon bias” [13]. For methods based on LCE, they either assume a “biased” parametric model, or they are more generic but require a more expensive initialization [14]. The second limitation is that the configuration returned is often the lowest validation error observed during the search. However, this approach is quite limiting because it does not consider more sophisticated model selection schemes such as cross-validation or Hyperband. Therefore, we added the **model selection** block in Fig. 1.

Benchmarks play a critical role in the design and development of HPO methods. We have surveyed several recent benchmarks for learning curves that are continuously evolving, such as HPOBench [15, 16], LCBench [17], JAHS-Bench-201 [18], and YAHPO-Gym [19].

### 3 Method

Our experiments include representative state-of-the-art methods in multi-fidelity HPO research. The first method considered is **Successive Halving** (SHA) [2] which we consider in its asynchronous variant [6] to be able to handle a stream of learners. SHA decides to continue the training loop at exponentially increasing steps called rungs (based on learning curves [20] theory). At each rung, it continues training only if the current observed score is among the Top- $K$  scores and discards the model otherwise. This method is error-prone when learning curves are noisy. To make SHA more robust **Hyperband** (HB) [3] enforces the exploration of more levels of fidelity on top of SHA to reduce the impact of noise in decisions. However this is at the expense of consuming more training steps than SHA. Finally, the last approach we consider consists in using a surrogate model to predict the future of the learning curve, called **Learning Curve Extrapolation**

(LCE) [7]. In theory, this approach is less likely to suffer from “short-horizon” bias unless the learning curve surrogate is biased. The method also suffers from instabilities in its original implementation [7]. To address this shortcoming, we are using our own implementation of this method (See the RoBER method described in supplemental material<sup>1</sup>).

We compare these methods with two baselines: the maximum fidelity baseline (**100-Epochs**), which trains all methods for 100 Epochs and selects the best one and a minimal fidelity baseline (**1-Epoch**), which comprises in fact two phases: (i) exploration search phase pre-selecting the Top- $K$  after only 1 Epoch (minimum fidelity blue loop in Fig. 1); (ii) model selection phase evaluating the Top- $K$  at maximum fidelity and returning the best model. 1-Epoch is expected to perform poorly if there is a strong “short-horizon” bias (see previous section for a definition) and noisy learning curves.

## 4 Experimental Results

In our analysis we looked at four benchmarks: HPOBench, LCBench, YAHPO-Gym and JAHS-Bench. But, since they all yielded similar results, for brevity we only expose the Naval Propulsion problem from HPOBench. We provide complementary materials <sup>1</sup> with our full set of results.

To visualize the algorithm progress, we display the test error as a function of training epochs [19] (Fig.1). This type of learning curve weighs training iterations (epochs) equally for all HP configurations, which may be deceiving since they can vary in computational cost. Still, this is a convenient simple method abstracting from details of implementation. We configured the number of iterations for the outer-loop optimization to 200 (red loop). As a result, the maximum fidelity policy (100-Epochs), consumes the maximum number of 20,000 training epochs. In contrast, its counterpart at minimum fidelity (1-Epoch, performed with a Top-3 model selection ( $K = 3$ ) at maximum fidelity after the search), consumes a fixed number of  $200 \times 1 + 3 \times 100 = 500$  training epochs.

Figure 2 shows the evolution of the test RMSE as a function of the number of training epochs used. The key observation is that the test RMSE for the 1-Epoch (blue curve) and 100-Epochs (orange curve) strategies is similar at the final point. However, the 1-Epoch policy uses 40 times fewer training epochs than its counterpart, the 100-Epoch policy ( $20,000/500=40$ ), as per our experimental setup, which is evident from the last point of the blue and orange





Fig. 2: Comparing the performance (mean and standard error) of various early discarding strategies (all combined with with random search) on the Naval Propulsion problem from the HPOBench benchmark.

<sup>1</sup>[https://github.com/deephyper/scalable-bo/blob/main/esann-23/One\\_Epoch\\_Is\\_Often\\_All\\_You\\_Need\\_Extended.pdf](https://github.com/deephyper/scalable-bo/blob/main/esann-23/One_Epoch_Is_Often_All_You_Need_Extended.pdf)



(a) Learning curves colored by final rank. (b) Evolution of the ranking with epochs.

Fig. 3: Visualizing the ranking for good (blue) and bad (red) models for 1,000 randomly sampled learning curves from the Naval Propulsion benchmark. (a) Models can be selected from the first epoch. (b) The ranking stabilizes faster among the best models.

curves. Additionally, we note that the more complex agents (SHA, HB, and LCE) do not differ significantly from each other and all result in similar final test RMSE, but they consume significantly more epochs. Hyperband consumes slightly more resources than SHA, which is consistent with its design. SHA consumes the least training epochs, about 10 times fewer epochs than 100-Epochs, but still, four times more than 1-Epoch.

After observing the performance of training for only one epoch during the search, we will explain why. In Fig. 3a, we display 1,000 randomly sampled learning curves from the same benchmark, colored by their ranking at maximum fidelity. Low ranks, colored in blue, correspond to good models, while high ranks, colored in red, correspond to bad models. This visualization reveals that the groups of bad and good models can be identified in the first epoch of training, making selection with one epoch effective. Some noise exists between models, as seen in the blue curves, necessitating the top- $K$  tournament selection to differentiate among different models. The noise in the ranking for the same benchmark can be visualized in Fig.3b, where the ranking for the same models at each training epoch is displayed. The first half (before 50 epochs) has more noise than the second half, and the ranking of good models stabilizes faster, already after 5 epochs, compared to the ranking of bad models, which remains relatively noisy until the end.

## 5 Conclusion

In this study, we investigated the optimization of hyperparameters for machine learning models with minimal training steps using MF-HPO agents. We compared various state-of-the-art multi-fidelity policies and discovered that a new simple baseline method, not considered in prior benchmarks, namely the 1-Epoch strategy assessing candidate configuration at lowest fidelity, performs surprisingly well. This was traced to the fact that good and bad models can be distinguished early in the training process. This phenomenon occurs frequently in benchmarks commonly used in the literature. Therefore, this point to the need to include 1-Epoch in forthcoming benchmarks and eventually designing new harder (possi-

bly more realistic) benchmarks that will defeat it. Our work is limited to using “epoch” as a unit of fidelity. While this is convenient and appealing to conduct studies independent of hardware implementation considerations, practical application settings may require considering wall time or other options as units of fidelity.

## References

- [1] T. Yu et al. Hyper-parameter optimization: A review of algorithms and applications. *arXiv preprint arXiv:2003.05689*, 2020.
- [2] K. Jamieson et al. Non-stochastic best arm identification and hyperparameter optimization. In *Artificial intelligence and statistics*, pages 240–248. PMLR, 2016.
- [3] L. Li et al. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The Journal of Machine Learning Research*, 18(1):6765–6816, 2017.
- [4] O. Bohdal et al. PASHA: Efficient HPO and NAS with progressive resource allocation. In *The Eleventh International Conference on Learning Representations*, 2023.
- [5] F. Mohr et al. Learning curves for decision making in supervised machine learning—a survey. *arXiv preprint arXiv:2201.12150*, 2022.
- [6] L. Li et al. A system for massively parallel hyperparameter tuning. *Proceedings of Machine Learning and Systems*, 2:230–246, 2020.
- [7] T. Domhan et al. Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Twenty-fourth international joint conference on artificial intelligence*, 2015.
- [8] A. Klein et al. Fast bayesian optimization of machine learning hyperparameters on large datasets. In *Artificial intelligence and statistics*, pages 528–536. PMLR, 2017.
- [9] S. Falkner et al. Bohb: Robust and efficient hyperparameter optimization at scale. In *International Conference on Machine Learning*, pages 1437–1446. PMLR, 2018.
- [10] B. Baker et al. Accelerating neural architecture search using performance prediction. *arXiv preprint arXiv:1705.10823*, 2017.
- [11] A. Klein et al. Learning curve prediction with bayesian neural networks. In *International Conference on Learning Representations*, 2017.
- [12] J. Wu et al. Practical multi-fidelity bayesian optimization for hyperparameter tuning. In *Uncertainty in Artificial Intelligence*, pages 788–798. PMLR, 2020.
- [13] Y. Wu et al. Understanding short-horizon bias in stochastic meta-optimization. In *International Conference on Learning Representations*, 2018.
- [14] C. White et al. How powerful are performance predictors in neural architecture search? *Advances in Neural Information Processing Systems*, 34:28454–28469, 2021.
- [15] K. Eggenberger et al. Hpobench: A collection of reproducible multi-fidelity benchmark problems for hpo. *arXiv preprint arXiv:2109.06716*, 2021.
- [16] A. Klein et al. Tabular benchmarks for joint architecture and hyperparameter optimization. *arXiv preprint arXiv:1905.04970*, 2019.
- [17] L. Zimmer et al. Auto-pytorch: Multi-fidelity metalearning for efficient and robust autodl. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3079–3090, 2021.
- [18] A. Bansal et al. Jahs-bench-201: A foundation for research on joint architecture and hyperparameter search. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022.
- [19] F. Pfisterer et al. Yahpo gym—an efficient multi-objective multi-fidelity benchmark for hyperparameter optimization. In *International Conference on Automated Machine Learning*, pages 3–1. PMLR, 2022.
- [20] M. Hutter. Learning curve theory. *arXiv preprint arXiv:2102.04074*, 2021.