

---

# Reshuffling Resampling Splits Can Improve Generalization of Hyperparameter Optimization

---

Thomas Nagler\*  
t.nagler@lmu.de

Lennart Schneider\*

Bernd Bischl

Matthias Feurer

Department of Statistics, LMU Munich  
Munich Center for Machine Learning (MCML)

## Abstract

Hyperparameter optimization is crucial for obtaining peak performance of machine learning models. The standard protocol evaluates various hyperparameter configurations using a resampling estimate of the generalization error to guide optimization and select a final hyperparameter configuration. Without much evidence, paired resampling splits, i.e., either a fixed train-validation split or a fixed cross-validation scheme, are often recommended. We show that, surprisingly, reshuffling the splits for every configuration often improves the final model’s generalization performance on unseen data. Our theoretical analysis explains how reshuffling affects the asymptotic behavior of the validation loss surface and provides a bound on the expected regret in the limiting regime. This bound connects the potential benefits of reshuffling to the signal and noise characteristics of the underlying optimization problem. We confirm our theoretical results in a controlled simulation study and demonstrate the practical usefulness of reshuffling in a large-scale, realistic hyperparameter optimization experiment. While reshuffling leads to test performances that are competitive with using fixed splits, it drastically improves results for a single train-validation holdout protocol and can often make holdout become competitive with standard CV while being computationally cheaper.

## 1 Introduction

Hyperparameters have been shown to strongly influence the performance of machine learning models (van Rijn & Hutter, 2018; Probst et al., 2019). The primary goal of hyperparameter optimization (HPO; also called tuning) is the identification and selection of a hyperparameter configuration (HPC) that minimizes the estimated generalization error (Feurer & Hutter, 2019; Bischl et al., 2023). Typically, this task is challenged by the absence of a closed-form mathematical description of the objective function, the unavailability of an analytic gradient, and the large cost to evaluate HPCs, categorizing HPO as a noisy, black-box optimization problem. An HPC is evaluated via resampling, such as a holdout split or  $M$ -fold cross-validation (CV), during tuning.

These resampling splits are usually constructed in a fixed and instantiated manner, i.e., the same training and validation splits are used for the internal evaluation of all configurations. On the one hand, this is an intuitive approach, as it should facilitate a fair comparison between HPCs and reduce the variance in the comparison.<sup>1</sup> On the other hand, such a fixing of train and validation splits might steer the optimization, especially after a substantial budget of evaluations, towards favoring HPCs

---

\*Equal contribution.

<sup>1</sup>This approach likely originates from the concept of paired statistical tests and the resulting variance reduction, but in our literature search we did not find any references discussing this in the context of HPO. For example, when comparing the performance of two classifiers on one dataset, paired tests are commonly

which are specifically tailored to the chosen splits. Such and related effects, where we "overoptimize" the validation performance without effective reward in improved generalization performance have been sometimes dubbed "overtuning" or "oversearching". For a more detailed discussion of this topic, including related work, see Section 5 and Appendix B. The practice of reshuffling resampling splits during HPO is generally neither discussed in the scientific literature nor HPO software tools.<sup>2</sup> To the best of our knowledge, only Lévesque (2018) investigated reshuffling train-validation splits for every new HPC. For both holdout and  $M$ -fold CV using reshuffled resampling splits resulted in, on average, slightly lower generalization error when used in combination with Bayesian optimization (BO, Garnett, 2023) or CMA-ES (Hansen & Ostermeier, 2001) as HPO algorithms. Additionally, reshuffling was used by a solution to the NeurIPS 2006 performance prediction challenge to estimate the final generalization performance (Guyon et al., 2006). Recently, in the context of evolutionary optimization, reshuffling was applied after every generation (Larcher & Barbosa, 2022).

In this paper, we systematically examine the effect of reshuffling on HPO performance. Our contributions can be summarized as follows:

1. We show theoretically that reshuffling resampling splits during HPO can result in finding a configuration with better overall generalization performance, especially when the loss surface is rather flat and its estimate is noisy (Section 2).
2. We confirm these theoretical insights through controlled simulation studies (Section 3).
3. We demonstrate in realistic HPO benchmark experiments that reshuffling splits can lead to a real-world improvement of HPO (Section 4). Especially in the case of reshuffled holdout, we find that the final generalization performance is often on par with 5-fold CV under a wide range of settings.

We discuss results, limitations, and avenues for future research in Section 5.

## 2 Theoretical Analysis

### 2.1 Problem Statement and Setup

Machine learning (ML) aims to fit a model to data, so that it generalizes well to new observations of the same distribution. Let  $\mathcal{D} = \{\mathbf{Z}_i\}_{i=1}^n$  be the observed dataset consisting of *i.i.d.* random variables from a distribution  $P$ , i.e., in the supervised setting  $\mathbf{Z}_i = (\mathbf{X}_i, Y_i)$ .<sup>3,4</sup> Formally, an inducer  $g$  configured by an HPC  $\lambda \in \Lambda$  maps a dataset  $\mathcal{D}$  to a model from our hypothesis space  $h = g_\lambda(\mathcal{D}) \in \mathcal{H}$ . During HPO, we want to find a HPC that minimizes the expected generalization error, i.e., find

$$\lambda^* = \arg \min_{\lambda \in \Lambda} \mu(\lambda), \quad \text{where} \quad \mu(\lambda) = \mathbb{E}[\ell(\mathbf{Z}, g_\lambda(\mathcal{D}))],$$

where  $\ell(\mathbf{Z}, h)$  is the loss of model  $h$  on a fresh observation  $\mathbf{Z}$ . In practice, there is usually a limited computational budget for each HPO run, so we assume that there is only a finite number of distinct HPCs  $\Lambda = \{\lambda_1, \dots, \lambda_J\}$  to be evaluated, which also simplifies the subsequent analysis. Naturally, we cannot optimize the generalization error directly, but only an estimate of it. To do so, a resampling is constructed. For every HPC  $\lambda_j$ , draw  $M$  random sets  $\mathcal{I}_{1,j}, \dots, \mathcal{I}_{M,j} \subset \{1, \dots, n\}$  of validation indices with  $n_{\text{valid}} = \lceil \alpha n \rceil$  instances each. The random index draws are assumed to be independent of the observed data. The data is then split accordingly into pairs  $\mathcal{V}_{m,j} = \{\mathbf{Z}_i\}_{i \in \mathcal{I}_{m,j}}, \mathcal{T}_{m,j} = \{\mathbf{Z}_i\}_{i \notin \mathcal{I}_{m,j}}$  of disjoint validation and training sets. Define the validation loss on the  $m$ -th fold

$$L(\mathcal{V}_{m,j}, g_{\lambda_j}(\mathcal{T}_{m,j})) = \frac{1}{n_{\text{valid}}} \sum_{i \in \mathcal{I}_{m,j}} \ell(\mathbf{Z}_i, g_{\lambda_j}(\mathcal{T}_{m,j})),$$

employed that implicitly assume that differences between the performance of classifiers on a given CV fold are comparable (Dietterich, 1998; Nadeau & Bengio, 1999, 2003; Demšar, 2006).

<sup>2</sup>In Appendix B, we present an overview of how resampling is addressed in tutorials and examples of standard HPO libraries and software. We conclude that usually fixed splits are used or recommended.

<sup>3</sup>Throughout, we use bold letters to indicate (fixed and random) vectors.

<sup>4</sup>We provide a notation table for symbols used in the main paper in Table 2 in the appendix.

and the  $M$ -fold validation loss as

$$\hat{\mu}(\boldsymbol{\lambda}_j) = \frac{1}{M} \sum_{m=1}^M L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})).$$

Since  $\mu$  is unknown, we minimize  $\hat{\boldsymbol{\lambda}} = \arg \min_{\boldsymbol{\lambda} \in \Lambda} \hat{\mu}(\boldsymbol{\lambda})$ , hoping that  $\mu(\hat{\boldsymbol{\lambda}})$  will also be small. Typically, the same splits are used for every HPC, so  $\mathcal{I}_{m,j} = \mathcal{I}_m$  for all  $j = 1, \dots, J$  and  $m = 1, \dots, M$ . In the following, we investigate how reshuffling train-validation splits (i.e.,  $\mathcal{I}_{m,j} \neq \mathcal{I}_{m,j'}$  for  $j \neq j'$ ) affects the HPO problem.

## 2.2 How Reshuffling Affects the Loss Surface

We first investigate how different validation and reshuffling strategies affect the empirical loss surface  $\hat{\mu}$ . In particular, we derive the limiting distribution of the sequence  $\sqrt{n}(\hat{\mu}(\boldsymbol{\lambda}_j) - \mu(\boldsymbol{\lambda}_j))_{j=1}^J$ . This limiting regime will not only reveal the effect of reshuffling on the loss surface, but also give us a tractable setting to study HPO performance.

**Theorem 2.1.** *Under regularity conditions stated in Appendix C.1, it holds*

$$\sqrt{n}(\hat{\mu}(\boldsymbol{\lambda}_j) - \mu(\boldsymbol{\lambda}_j))_{j=1}^J \rightarrow \mathcal{N}(0, \Sigma) \quad \text{in distribution,}$$

where

$$\Sigma_{i,j} = \tau_{i,j,M} K(\boldsymbol{\lambda}_i, \boldsymbol{\lambda}_j), \quad \tau_{i,j,M} = \lim_{n \rightarrow \infty} \frac{1}{nM^2\alpha^2} \sum_{s=1}^n \sum_{m=1}^M \sum_{m'=1}^M \Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}),$$

and

$$K(\boldsymbol{\lambda}_i, \boldsymbol{\lambda}_j) = \lim_{n \rightarrow \infty} \text{Cov}[\bar{\ell}_n(\mathbf{Z}', \boldsymbol{\lambda}_i), \bar{\ell}_n(\mathbf{Z}', \boldsymbol{\lambda}_j)], \quad \bar{\ell}_n(\mathbf{z}, \boldsymbol{\lambda}) = \mathbb{E}[\ell(\mathbf{z}, g_{\boldsymbol{\lambda}}(\mathcal{T}))] - \mathbb{E}[\ell(\mathbf{Z}, g_{\boldsymbol{\lambda}}(\mathcal{T}))],$$

where the expectation is taken over a training set  $\mathcal{T}$  of size  $n$  and two fresh samples  $\mathbf{Z}, \mathbf{Z}'$  from the same distribution.

The regularity conditions are rather mild and discussed further in Appendix C.1. The kernel  $K$  reflects the (co-)variability of the losses caused by validation samples. The contribution of training samples only has a higher-order effect. The validation scheme enters the distribution through the quantities  $\tau_{i,j,M}$ . In what follows, we compute explicit expressions for some popular examples. The following list provides formal definitions for the index sets  $\mathcal{I}_{m,j}$ .

- (i) (holdout) Let  $M = 1$  and  $\mathcal{I}_{1,j} = \mathcal{I}_1$  for all  $j = 1, \dots, J$ , and some size- $\lceil \alpha n \rceil$  index set  $\mathcal{I}_1$ .
- (ii) (reshuffled holdout) Let  $M = 1$  and  $\mathcal{I}_{1,1}, \dots, \mathcal{I}_{1,J}$  be independently drawn from the uniform distribution over all size- $\lceil \alpha n \rceil$  subsets from  $\{1, \dots, n\}$ .
- (iii) ( $M$ -fold CV) Let  $\alpha = 1/M$  and  $\mathcal{I}_1, \dots, \mathcal{I}_M$  be a disjoint partition of  $\{1, \dots, n\}$ , and  $\mathcal{I}_{m,j} = \mathcal{I}_m$  for all  $j = 1, \dots, J$ .
- (iv) (reshuffled  $M$ -fold CV) Let  $\alpha = 1/M$  and  $(\mathcal{I}_{1,j}, \dots, \mathcal{I}_{M,j}), j = 1, \dots, J$ , be independently drawn from the uniform distribution over disjoint partitions of  $\{1, \dots, n\}$ .
- (v) ( $M$ -fold holdout) Let  $\mathcal{I}_m, m = 1, \dots, M$ , be independently drawn from the uniform distribution over size- $\lceil \alpha n \rceil$  subsets of  $\{1, \dots, n\}$  and set  $\mathcal{I}_{m,j} = \mathcal{I}_m$  for all  $m = 1, \dots, M, j = 1, \dots, J$ .
- (vi) (reshuffled  $M$ -fold holdout) Let  $\mathcal{I}_{m,j}, m = 1, \dots, M, j = 1, \dots, J$ , be independently drawn from the uniform distribution over size- $\lceil \alpha n \rceil$  subsets of  $\{1, \dots, n\}$ .

The value of  $\tau_{i,j,M}$  for each example is computed explicitly in Appendix E. In all these examples, we in fact have

$$\tau_{i,j,M} = \begin{cases} \sigma^2, & i = j \\ \tau^2 \sigma^2, & i \neq j. \end{cases} \quad (1)$$

for some method-dependent parameters  $\sigma, \tau$  shown in Table 1. The parameter  $\sigma^2$  captures any increase in variance caused by omitting an observation from the validation sets. The parameter  $\tau$  quantifies a potential decrease in correlation in the loss surface due to reshuffling. More precisely,

Table 1: Exemplary parametrizations in Equation (1) for resamplings; see Appendix E for details.

Method	$\sigma^2$	$\tau^2$
holdout (HO)	$1/\alpha$	1
reshuffled HO	$1/\alpha$	$\alpha$
$M$ -fold CV	1	1
reshuffled $M$ -fold CV	1	1
$M$ -fold HO (subsampling / Monte Carlo CV)	$1 + (1 - \alpha)/M\alpha$	1
reshuffled $M$ -fold HO	$1 + (1 - \alpha)/M\alpha$	$1/(1 + (1 - \alpha)/M\alpha)$

the observed losses  $\widehat{\mu}(\lambda_i), \widehat{\mu}(\lambda_j)$  at distinct HPCs  $\lambda_i \neq \lambda_j$  become less correlated when  $\tau$  is small. Generally, an increase in variance leads to worse generalization performance. The effect of a correlation decrease is less obvious and is studied in detail in the following section.

We make the following observations about the differences between methods in Table 1:

- $M$ -fold CV incurs no increase in variance ( $\sigma^2 = 1$ ) and — because every HPC uses the same folds — no decrease in correlation. Interestingly, the correlation does not even decrease when reshuffling the folds. In any case, all samples are used exactly once as validation and training instance. At least asymptotically, this leads to the same behavior, and reshuffling should have almost no effect on  $M$ -fold CV.
- The two (1-fold) holdout methods bear the same  $1/\alpha$  increase in variance. This is caused by only using a fraction  $\alpha$  of the data as validation samples. Reshuffled holdout also decreases the correlation parameter  $\tau^2$ . In fact, if HPCs  $\lambda_i \neq \lambda_j$  are evaluated on largely distinct samples, the validation losses  $\widehat{\mu}(\lambda_i)$  and  $\widehat{\mu}(\lambda_j)$  become almost independent.
- $M$ -fold holdout also increases the variance, because some samples may still be omitted from validation sets. This increase is much smaller for large  $M$ . Accordingly, the correlation is also decreased by less in the reshuffled variant.

### 2.3 How Reshuffling Affects HPO Performance

In practice, we are mainly interested in the performance of a model trained with the optimal HPC  $\widehat{\lambda}$ . To simplify the analysis, we explore this in the large-sample regime derived in the previous section. Assume

$$\widehat{\mu}(\lambda_j) = \mu(\lambda_j) + \epsilon(\lambda_j) \tag{2}$$

where  $\epsilon(\lambda)$  is a zero-mean Gaussian process with covariance kernel

$$\text{Cov}(\epsilon(\lambda), \epsilon(\lambda')) = \begin{cases} K(\lambda, \lambda) & \text{if } \lambda = \lambda', \\ \tau^2 K(\lambda, \lambda') & \text{else.} \end{cases} \tag{3}$$

Let  $\Lambda \subseteq \{\lambda \in \mathbb{R}^d: \|\lambda\| \leq 1\}$  with  $|\Lambda| = J < \infty$  be the set of hyperparameters. Theorem 2.2 ahead gives a bound on the expected regret  $\mathbb{E}[\mu(\widehat{\lambda}) - \mu(\lambda^*)]$ . It depends on several quantities characterizing the difficulty of the HPO problem. The constant

$$\kappa = \sup_{\|\lambda\|, \|\lambda'\| \leq 1} \frac{|K(\lambda, \lambda) - K(\lambda, \lambda')|}{K(\lambda, \lambda) \|\lambda - \lambda'\|^2}.$$

can be interpreted as a measure of correlation of the process  $\epsilon$ . In particular,  $\text{Corr}(\epsilon(\lambda), \epsilon(\lambda')) \geq 1 - \kappa \|\lambda - \lambda'\|^2$ . The constant is small when  $\epsilon$  is strongly correlated, and large otherwise. Further, define  $\eta$  as the minimal number such that any  $\eta$ -ball contained in  $\{\|\lambda\| \leq 1\}$  contains at least one element of  $\Lambda$ . It measures how densely the set of candidate HPCs  $\Lambda$  covers set of all possible HPCs. If  $\Lambda$  is a deterministic uniform grid, we have about  $\eta \approx J^{-1/d}$ . Similarly, Lemma D.1 in the Appendix shows that  $\eta \lesssim J^{-1/2d}$  when randomly sampling HPCs. Finally, the constant

$$m = \sup_{\lambda \in \Lambda} \frac{|\mu(\lambda) - \mu(\lambda^*)|}{\|\lambda - \lambda^*\|^2},$$

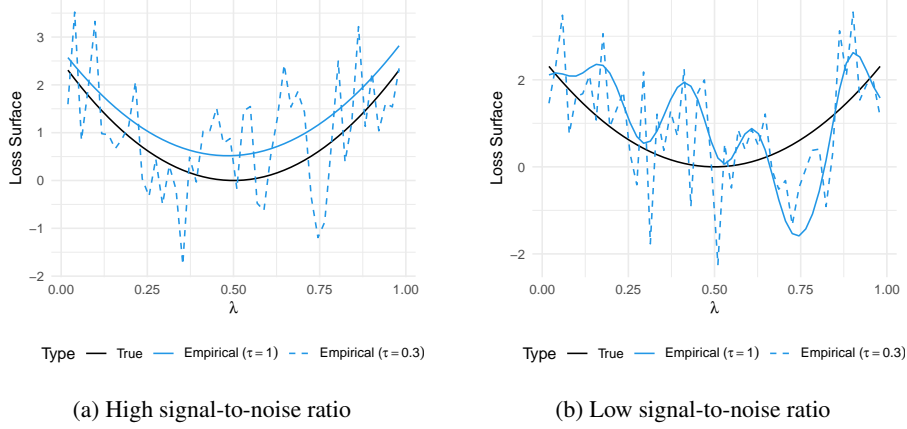


Figure 1: Example of reshuffled empirical loss yielding a worse (left) and better (right) minimizer.

measures the local curvature at the minimum of the loss surface  $\mu$ . Finding an HPC  $\lambda$  close to the theoretical optimum  $\lambda^*$  is easier when the minimum is more pronounced (large  $m$ ). On the other hand, the regret  $\mu(\lambda) - \mu(\lambda^*)$  is also punishing mistakes more quickly. Defining  $\log(x)_+ = \max\{0, \log(x)\}$ , we can now state our main result.

**Theorem 2.2.** *Let  $\hat{\mu}$  follow the Gaussian process model (2). Suppose  $\kappa < \infty$ ,  $0 < \underline{\sigma}^2 \leq \text{Var}[\epsilon(\lambda)] \leq \sigma^2 < \infty$  for all  $\lambda \in \Lambda$ , and  $m > 0$ . Then*

$$\mathbb{E}[\mu(\hat{\lambda}) - \mu(\lambda^*)] \leq \sigma\sqrt{d}[8 + B(\tau) - A(\tau)],$$

where

$$B(\tau) = 48 \left[ \sqrt{1 - \tau^2} \sqrt{\log J} + \tau \sqrt{1 + \log(3\kappa)_+} \right], \quad A(\tau) = \sqrt{1 - \tau^2} (\underline{\sigma}/\sigma) \sqrt{\log \left( \frac{\sigma}{2m\eta^2} \right)_+}.$$

The numeric constants result from several simplifications in a worst-case analysis, which lowers their practical relevance. A qualitative analysis of the bound is still insightful. The bound is increasing in  $\sigma$  and  $d$ , indicating that the HPO problem is harder when there is a lot of noise or there are many parameters to tune. The terms  $B(\tau)$  and  $A(\tau)$  have conceptual interpretations:

- The term  $B(\tau)$  quantifies how likely it is to pick a bad  $\hat{\lambda}$  because of bad luck: a  $\lambda$  far away from  $\lambda^*$  had such a small  $\epsilon(\lambda)$  that it outweighs the increase in  $\mu$ . Such events are more likely when the process  $\epsilon$  is weakly correlated. Accordingly,  $B(\tau)$  is decreasing in  $\tau$  and increasing in  $\kappa$ .
- The term  $A(\tau)$  quantifies how likely it is to pick a good  $\hat{\lambda}$  by luck: a  $\lambda$  close to  $\lambda^*$  had such a small  $\epsilon(\lambda)$  that it overshoots all the other fluctuations. Also such events are more likely when the process  $\epsilon$  is weakly correlated. Accordingly, the term  $A(\tau)$  is decreasing in  $\tau$ .

The  $B$ , as stated, is unbounded, but a closer inspection of the proof shows that it is upper bounded by  $\sqrt{\log J}$ . This bound is attained only in the unrealistic scenario when the validation losses are essentially uncorrelated across all HPCs. The term  $A$  is bounded from below by zero, which is also the worst case because the term enters our regret bound with a negative sign.

Both  $A$  and  $B$  are decreasing in the reshuffling parameter  $\tau$ . There are two regimes. If  $\sigma/2m\eta^2 \leq e$ , then  $A(\tau) = 0$  and reshuffling cannot lead to an improvement of the bound. The term  $\sigma/m\eta^2$  can be interpreted as noise-to-signal ratio (relative to the grid density). If the signal is much stronger than the noise, the HPO problem is so easy that reshuffling will not help. This situation is illustrated in Figure 1a.

If on the other hand  $\sigma/m\eta^2 > e$ , the terms  $A(\tau)$  and  $B(\tau)$  enter the bound with opposing signs. This creates tension: reshuffling between HPCs increases  $B(\tau)$ , which is countered by a decrease in  $A(\tau)$ . So which scenarios favor reshuffling? When the process  $\epsilon$  is strongly correlated,  $\kappa$  is small and reshuffling (decreasing  $\tau$ ) incurs a high cost in  $B(\tau)$ . This is intuitive: When there is strong

correlation, the validation loss surface  $\hat{\mu}$  is essentially just a vertical shift of  $\mu$ . Finding the optimal  $\lambda$  is then almost as easy as if we would know  $\mu$ , and decorrelating the surface through reshuffling would make it unnecessarily hard. When  $\epsilon$  is less correlated ( $\kappa$  large) however, reshuffling does not hurt the term  $B(\tau)$  as much, but we can reap all the benefits of increasing  $A(\tau)$ . Here, the effect of reshuffling can be interpreted as hedging against the catastrophic case where all  $\hat{\mu}(\lambda)$  close to the optimal  $\lambda^*$  are simultaneously dominated by a region of bad hyperparameters. This is illustrated in Figure 1b.

### 3 Simulation Study

To test our theoretical understanding of the potential benefits of reshuffling resampling splits during HPO, we conduct a simulation study. This study helps us explore the effects of reshuffling in a controlled setting.

#### 3.1 Design

We construct a univariate quadratic loss surface function  $\mu : \Lambda \subset \mathbb{R} \mapsto \mathbb{R}, \lambda \rightarrow m(\lambda - 0.5)^2/2$  which we want to minimize. The global minimum is given at  $\mu(0.5) = 0$ . Combined with a kernel for the noise process  $\epsilon$  as in Equation (3), this allows us to simulate an objective as observed during HPO by sampling  $\hat{\mu}(\lambda) = \mu(\lambda) + \epsilon(\lambda)$ . We use a squared exponential kernel  $K(\lambda, \lambda') = \sigma_K^2 \exp(-\kappa(\lambda - \lambda')^2/2)$  that is plugged into the covariance kernel of the noise process  $\epsilon$  in Equation (3). The parameters  $m$  and  $\kappa$  in our simulation setup correspond exactly to the curvature and correlation constants from the previous sections. Recall that Theorem 2.2 states that the effect of reshuffling strongly depends on the curvature  $m$  of the loss surface  $\mu$  (a larger  $m$  implies a stronger curvature) and the constant  $\kappa$  as a measure of correlation of the noise  $\epsilon$  (a larger  $\kappa$  implies weaker correlation). Combined with the possibility to vary  $\tau$  in the covariance kernel of  $\epsilon$ , we can systematically investigate how curvature of the loss surface, correlation of the noise and the extent of reshuffling affect optimization performance. In each simulation run, we simulate the observed objective  $\hat{\mu}(\lambda)$ , identify the minimizer  $\hat{\lambda} = \arg \min_{\lambda \in \Lambda} \hat{\mu}(\lambda)$ , and calculate its true risk,  $\mu(\hat{\lambda})$ . We repeat this process 10000 times for various combinations of  $\tau$ ,  $m$ , and  $\kappa$ .

#### 3.2 Results

Figure 2 visualizes the true risk of the configuration  $\hat{\lambda}$  that minimizes the observed objective. We observe that for a loss surface with low curvature (i.e.,  $m \leq 2$ ), reshuffling is beneficial (lower values of  $\tau$  resulting in a better true risk of the configuration that optimizes the observed objective) as long as the noise process is not too correlated (i.e.,  $\kappa \geq 1$ ). As soon as the noise process is more strongly correlated, even flat valleys of the true risk  $\mu$  remain clearly visible in the observed risk  $\hat{\mu}$ , and reshuffling starts to hurt the optimization performance. Moving to scenarios of high curvature, the general relationship of  $m$  and  $\kappa$  remains the same, but reshuffling starts to hurt optimization performance already with weaker correlation in the noise. In summary, the simulations show that in cases of low curvature of the loss surface, reshuffling (reducing  $\tau$ ) tends to improve the true risk of the optimized configuration, especially when the loss surface is flat (small  $m$ ) and the noise is not strongly correlated (i.e.,  $\kappa$  is large). This exactly confirms our theoretical predictions from the previous section.

## 4 Benchmark Experiments

In this section, we present benchmark experiments of real-world HPO problems where we investigate the effect of reshuffling resampling splits during HPO. First, we discuss the experimental setup. Second, we present results for HPO using random search (Bergstra & Bengio, 2012). Third, we also show the effect of reshuffling when applied in BO using HEBO (Cowen-Rivers et al., 2022) and SMAC3 (Lindauer et al., 2022). Recall that our theoretical insight suggests that 1) reshuffling might be beneficial during HPO and 2) holdout should be affected the most by reshuffling and other resamplings should only be affected to a lesser extent.

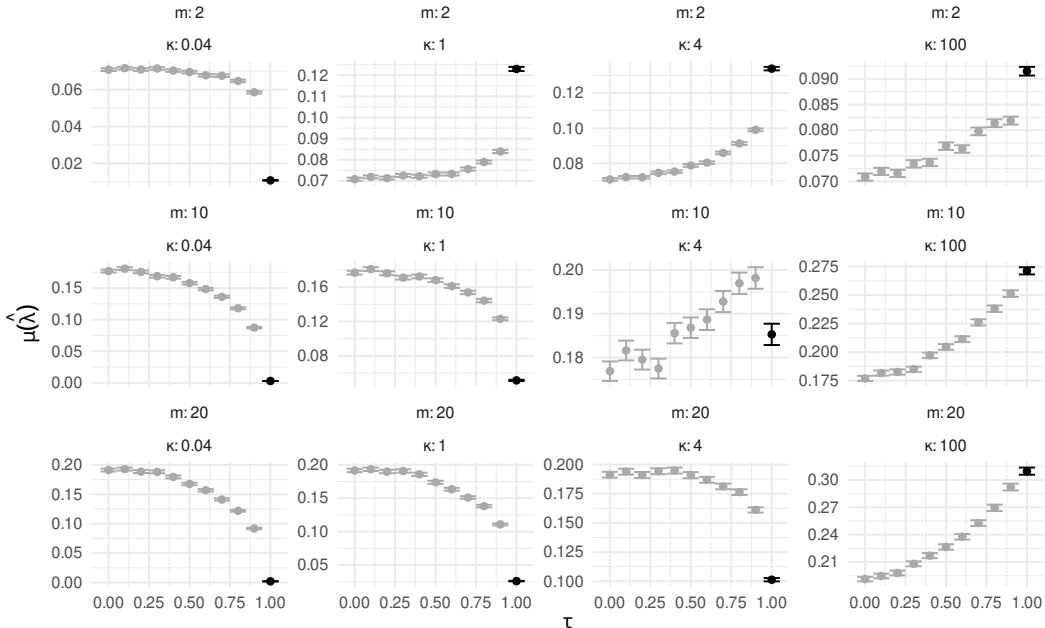


Figure 2: Mean true risk (lower is better) of the configuration minimizing the observed objective systematically varied with respect to curvature  $m$ , correlation strength  $\kappa$  of the noise (a larger  $\kappa$  implying weaker correlation), and extent of reshuffling  $\tau$  (lower  $\tau$  increasing reshuffling). A  $\tau$  of 1 indicates no reshuffling. Error bars represent standard errors.

#### 4.1 Experimental Setup

As benchmark tasks, we use a set of standard HPO problems defined on small- to medium-sized tabular datasets for binary classification. We suspect the effect of the resampling variant used and whether the resampling is reshuffled to be larger for smaller datasets, where the variance of the validation loss estimator is naturally higher. Furthermore, from a practical perspective, this also ensures computational feasibility given the large number of HPO runs in our experiments. We systematically vary the learning algorithm, optimized performance metric, resampling method, whether the resampling is reshuffled, and the size of the dataset used for training and validation during HPO. Below, we outline the general experimental design and refer to Appendix F for details.

We used a subset of the datasets defined by the AutoML benchmark (Gijbbers et al., 2024), treating these as data generating processes (DGPs; Hothorn et al., 2005). We only considered datasets with less than 100 features to reduce the required computation time and required the number of observations to be between 10000 and 1000000; for further details see Appendix F.1. Our aim was to robustly measure the generalization performance when varying the size  $n$ , which, as defined in Section 2 denotes the size of the combined data for model selection, so one training and validation set combined. First, we sampled 5000 data points per dataset for robust assessment of the generalization error; these points are not used during HPO in any way. Then, from the remaining points we sampled tasks with  $n \in \{500, 1000, 5000\}$ .

We selected CatBoost (Prokhorenkova et al., 2018) and XGBoost (Chen & Guestrin, 2016) for their state-of-the-art performance on tabular data (Grinsztajn et al., 2022; Borisov et al., 2022; McElfresh et al., 2023; Kohli et al., 2024). Additionally, we included an Elastic Net (Zou & Hastie, 2005) to represent a linear baseline with a smaller search space and a funnel-shaped MLP (Zimmer et al., 2021) as a cost-effective neural network baseline. We provide details regarding training pipelines and search spaces in Appendix F.2.

We conduct a random search with 500 HPC evaluations for every resampling strategy we described in Table 1, for both fixed and reshuffled splits. We always use 80/20 train-validation splits for holdout




Figure 3: Average test performance (negative ROC AUC) of the incumbent for XGBoost on dataset albert for increasing  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

and 5-fold CVs, so that training set size (and negative estimation bias) are the same. Anytime test performance of an HPO run is assessed by re-training the current incumbent (i.e. the best HPC until the current HPO iteration based on validation performance) on all available train and validation data and evaluating its performance on the outer test set. Note we do this for scientific evaluation in this experiment; obviously, this is not possible in practice. Using random search allows us to record various metrics and afterwards simulate optimizing for different ones, specifically, we recorded accuracy, area under the ROC curve (ROC AUC) and logloss.

We also investigated the effect of reshuffling on two state-of-the-art BO variants (Eggensperger et al., 2021; Turner et al., 2021), namely HEBO (Cowen-Rivers et al., 2022) and SMAC3 (Lindauer et al., 2022). The experimental design was the same as for random search, except for the budget, which we reduced from 500 HPCs to 250 HPCs, and only optimized ROC AUC.

## 4.2 Experimental Results

In the following, we focus on the results obtained using ROC AUC. We present aggregated results over different tasks, learning algorithms and replications to get a general understanding of the effects. Unaggregated results and results involving accuracy and logloss can be found in Appendix G.

**Results of Reshuffling Different Resamplings** For each resampling (holdout, 5-fold holdout, 5-fold CV, and 5x 5-fold CV), we empirically analyze the effect of reshuffling train and validation splits during HPO.

In Figure 3 we exemplarily show how test performance develops over the course of an HPO run on a single task for different resamplings (with and without reshuffling). Naturally, test performance does not necessarily increase in a monotonic fashion, and especially holdout without reshuffling tends to be unstable. Its reshuffled version results in substantially better test performance.

Next, we look at the relative *improvement* (compared to standard 5-fold CV, which we consider our baseline) with respect to *test* ROC AUC performance of the incumbent over time in Figure 4, i.e., the difference in test performance of the incumbent between standard 5-fold CV and a different resampling protocol; hence a positive difference tells us how much better in test error we are, if we would have chosen the other protocol instead 5-fold CV. We observe that reshuffling generally results in equal or better performance compared to the same resampling protocol without reshuffling. For 5-fold holdout and especially 5-fold CV and 5x 5-fold CV, reshuffling has a smaller effect on relative test performance improvement, as expected. Holdout is affected the most by reshuffling and results in substantially better relative test performance compared to standard holdout. We also observe that an HPO protocol based on reshuffled holdout results in similar final test performance as standard 5-fold CV while overall being substantially cheaper due to requiring less model fits per HPC evaluation. In Appendix G.2, we further provide an ablation study on the number of folds when using  $M$ -fold holdout, where we observed that – in line with our theory – the more folds are used, the less reshuffling affects  $M$ -fold holdout.






Figure 4: Average improvement (compared to standard 5-fold CV) with respect to test performance (ROC AUC) of the incumbent over different tasks, learning algorithms and replications separately for increasing  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

However, this general trend can vary for certain combinations of classifier and performance metric, see Appendix G. Especially for logloss, we observed that reshuffling rarely is beneficial; see the discussion in Section 5. Finally, the different resamplings generally behave as expected. The more we are willing to invest compute resources into a more intensive resampling like 5-fold CV or 5x 5-fold CV, the better the generalization performance of the final incumbent.

**Results for BO and Reshuffling** Figure 5 shows that, generally HEBO and SMAC3 outperform random search with respect to generalization performance (i.e., comparing HEBO and SMAC3 to random search under standard holdout, or comparing under reshuffled holdout). More interestingly, HEBO, SMAC3 and random search all strongly benefit from reshuffling. Moreover, the performance gap between HEBO and random search but also SMAC3 and random search narrows when the resampling is reshuffled, which is an interesting finding of its own: As soon as we are concerned with generalization performance of HPO and not only investigate validation performance during optimization, the choice of optimizer might have less impact on final generalization performance compared to other choices such as whether the resampling is reshuffled during HPO or not. We present results for BO and reshuffling for different resamplings in Appendix G.




Figure 5: Average improvement (compared to random search on standard holdout) with respect to test performance (ROC AUC) of the incumbent over tasks, learning algorithms and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

## 5 Discussion

In the previous sections, we have shown theoretically and empirically that reshuffling can enhance generalization performance of HPO. The main purpose of this article is to draw attention to this

surprising fact about a technique that is simple but rarely discussed. Our work goes beyond a preliminary experimental study on reshuffling (Lévesque, 2018), in that we also study the effect of reshuffling on random search, multiple metrics and learning algorithms, and most importantly, for the first time, we provide a theoretical analysis that explains why reshuffling can be beneficial.

**Limitations** To unveil the mechanisms underlying the reshuffling procedures, our theoretical analysis relies on an asymptotic approximation of the empirical loss surface. This allows us to operate on Gaussian loss surfaces, which exhibit convenient concentration and anti-concentration properties required in our proof. The latter are lacking for general distributions, which explains our asymptotic approach. The analysis was further facilitated by a loss stability assumption regarding the learning algorithms that is generally rather mild; see the discussion in Bayle et al. (2020). However, it typically fails for highly sensitive losses, which has practical consequences. In fact, Figure 9 in Appendix G shows that reshuffling usually hurts generalization for the logloss and small sample sizes. It is still an open question whether this problem can be fixed by less naive implementations of the technique. Another limitation is our focus on generalization after search through a fixed, finite set of candidates. This largely ignores the dynamic nature of many HPO algorithms, which would greatly complicate our analysis. Finally, our experiments are limited in that we restricted ourselves to tabular data and binary classification and we avoided extremely small or large datasets.

**Relation to Overfitting** The fact that generalization performance can decrease during HPO (or computational model selection in general) is sometimes known as oversearching, overtuning, or overfitting to the validation set (Quinlan & Cameron-Jones, 1995; Escalante et al., 2009; Koch et al., 2010; Igel, 2012; Bischl et al., 2023), but has arguably not been studied very thoroughly. Given recent theoretical (Feldman et al., 2019) and empirical (Purucker & Beel, 2023) findings, we expect less overtuning on multi-class datasets, making it interesting to see how reshuffling would affect the generalization performance.

Several works suggest strategies to counteract this effect. First, LOOCVCV proposes a conservative choice of incumbents (Ng, 1997) at the cost of leave-one-out analysis or an additional hyperparameter. Second, it is possible to use an extra *selection set* (Igel, 2012; Lévesque, 2018; Mohr et al., 2018) at the cost of reduced training data, which was found to lead to reduced overall performance (Lévesque, 2018). Third, by using early stopping one can stop hyperparameter optimization before the generalization performance degrades again. This was so far demonstrated to be able to save compute budget at only marginally reduced performance, but also requires either a sensitivity hyperparameter or correct estimation of the variance of the generalization estimate and was only developed for cross-validation so far (Makarova et al., 2022). Reshuffling itself is orthogonal to these proposals and a combination with the above-mentioned methods might result in further improvements.

**Outlook** Generally, the related literature detects overfitting to the validation set either visually (Ng, 1997) or by measuring it (Koch et al., 2010; Igel, 2012; Fabris & Freitas, 2019). Developing a unified formal definition of the above-mentioned terms and thoroughly analyzing the effect of decreased generalization performance after many HPO iterations and how it relates to our measurements of the validation performance is an important direction for future work.

We further found, both theoretically and experimentally, that investing more resources when evaluating each HPC can result in better final HPO performance. To reduce the computational burden on HPO again, we suggest further investigating the use of adaptive CV techniques, as proposed by AutoWEKA (Thornton et al., 2013) or under the name Lazy Paired Hyperparameter Tuning (Zheng & Bilenko, 2013). Designing more advanced HPO algorithms exploiting the reshuffling effect should be a promising avenue for further research.

## Acknowledgments and Disclosure of Funding

We thank Martin Binder and Florian Karl for helpful discussions. Lennart Schneider is supported by the Bavarian Ministry of Economic Affairs, Regional Development and Energy through the Center for Analytics - Data - Applications (ADACenter) within the framework of BAYERN DIGITAL II (20-3410-2-9-8). Lennart Schneider acknowledges funding from the LMU Mentoring Program of the Faculty of Mathematics, Informatics and Statistics.

## References

- Arlot, S. and Celisse, A. A survey of cross-validation procedures for model selection. *Statistics Surveys*, 4:40–79, 2010. B
- Austern, M. and Zhou, W. Asymptotics of cross-validation. *arXiv:2001.11111 [math.ST]*, 2020. C.1
- Awad, N., Mallik, N., and Hutter, F. DEHB: Evolutionary hyperband for scalable, robust and efficient Hyperparameter Optimization. In Zhou, Z. (ed.), *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI'21)*, pp. 2147–2153, 2021. B
- Bayle, P., Bayle, A., Janson, L., and Mackey, L. Cross-validation confidence intervals for test error. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.-F., and Lin, H. (eds.), *Proceedings of the 33rd International Conference on Advances in Neural Information Processing Systems (NeurIPS'20)*, pp. 16339–16350. Curran Associates, 2020. 5, C.1, C.1, C.1
- Bergman, E., Purucker, L., and Hutter, F. Don't waste your time: Early stopping cross-validation. In Eggenberger, K., Garnett, R., Vanschoren, J., Lindauer, M., and Gardner, J. (eds.), *Proceedings of the Third International Conference on Automated Machine Learning*, volume 256 of *Proceedings of Machine Learning Research*, pp. 9/1–31. PMLR, 2024. B
- Bergstra, J. and Bengio, Y. Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, 13:281–305, 2012. 4, B
- Bischi, B., Binder, M., Lang, M., Pielok, T., Richter, J., Coors, S., Thomas, J., Ullmann, T., Becker, M., Boulesteix, A., Deng, D., and Lindauer, M. Hyperparameter optimization: Foundations, algorithms, best practices, and open challenges. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, pp. e1484, 2023. 1, 5, B
- Blum, A., Kalai, A., and Langford, J. Beating the hold-out: Bounds for k-fold and progressive cross-validation. In *Proceedings of the Twelfth Annual Conference on Computational Learning Theory, COLT '99*, pp. 203–208, 1999. B
- Borisov, V., Leemann, T., Seßler, K., Haug, J., Pawelczyk, M., and Kasneci, G. Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–21, 2022. 4.1
- Bouckaert, Remco and Frank, E. Evaluating the Replicability of Significance Tests for Comparing Learning Algorithms. In Dai, H., Srikant, R., and Zhang, C. (eds.), *Advances in Knowledge Discovery and Data Mining*, pp. 3–12. Springer, 2004. B
- Bousquet, O. and Zhivotovskiy, N. Fast classification rates without standard margin assumptions. *Information and Inference: A Journal of the IMA*, 10(4):1389–1421, 2021. C.1
- Bouthillier, X., Delaunay, P., Bronzi, M., Trofimov, A., Nichyporuk, B., Szeto, J., Sepahvand, N. M., Raff, E., Madan, K., Voleti, V., Kahou, S. E., Michalski, V., Arbel, T., Pal, C., Varoquaux, G., and Vincent, P. Accounting for variance in machine learning benchmarks. In Smola, A., Dimakis, A., and Stoica, I. (eds.), *Proceedings of Machine Learning and Systems 3*, volume 3, pp. 747–769, 2021. B
- Buczak, P., Groll, A., Pauly, M., Rehof, J., and Horn, D. Using sequential statistical tests for efficient hyperparameter tuning. *AStA Advances in Statistical Analysis*, 108(2):441–460, 2024. B
- Cawley, G. and Talbot, N. On Overfitting in Model Selection and Subsequent Selection Bias in Performance Evaluation. *Journal of Machine Learning Research*, 11:2079–2107, 2010. B
- Chen, T. and Guestrin, C. XGBoost: A scalable tree boosting system. In Krishnapuram, B., Shah, M., Smola, A., Aggarwal, C., Shen, D., and Rastogi, R. (eds.), *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'16)*, pp. 785–794. ACM Press, 2016. 4.1
- Cowen-Rivers, A., Lyu, W., Tutunov, R., Wang, Z., Grosnit, A., Griffiths, R., Maraval, A., Jianye, H., Wang, J., Peters, J., and Ammar, H. HEBO: Pushing the limits of sample-efficient hyper-parameter optimisation. *Journal of Artificial Intelligence Research*, 74:1269–1349, 2022. 4, 4.1, B

- Demšar, J. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research*, 7:1–30, 2006. 1
- Dietterich, T. G. Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7):1895–1923, 1998. 1
- Dunias, Z., Van Calster, B., Timmerman, D., Boulesteix, A.-L., and van Smeden, M. A comparison of hyperparameter tuning procedures for clinical prediction models: A simulation study. *Statistics in Medicine*, 43(6):1119–1134, 2024. B
- Eggenesperger, K., Lindauer, M., Hoos, H., Hutter, F., and Leyton-Brown, K. Efficient benchmarking of algorithm configurators via model-based surrogates. *Machine Learning*, 107(1):15–41, 2018. 5
- Eggenesperger, K., Lindauer, M., and Hutter, F. Pitfalls and best practices in algorithm configuration. *Journal of Artificial Intelligence Research*, pp. 861–893, 2019. B
- Eggenesperger, K., Müller, P., Mallik, N., Feurer, M., Sass, R., Klein, A., Awad, N., Lindauer, M., and Hutter, F. HPOBench: A collection of reproducible multi-fidelity benchmark problems for HPO. In Vanschoren, J. and Yeung, S. (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran Associates, 2021. 4.1, B
- Escalante, H., Montes, M., and Sucar, E. Particle Swarm Model Selection. *Journal of Machine Learning Research*, 10:405–440, 2009. 5
- Fabris, F. and Freitas, A. Analysing the overfit of the auto-sklearn automated machine learning tool. In Nicosia, G., Pardalos, P., Umeton, R., Giuffrida, G., and Sciacca, V. (eds.), *Machine Learning, Optimization, and Data Science*, volume 11943 of *Lecture Notes in Computer Science*, pp. 508–520, 2019. 5
- Falkner, S., Klein, A., and Hutter, F. BOHB: Robust and efficient Hyperparameter Optimization at scale. In Dy, J. and Krause, A. (eds.), *Proceedings of the 35th International Conference on Machine Learning (ICML’18)*, volume 80, pp. 1437–1446. Proceedings of Machine Learning Research, 2018. B
- Feldman, V., Frostig, R., and Hardt, M. The advantages of multiple classes for reducing overfitting from test set reuse. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning (ICML’19)*, volume 97, pp. 1892–1900. Proceedings of Machine Learning Research, 2019. 5
- Feurer, M. and Hutter, F. Hyperparameter Optimization. In Hutter, F., Kotthoff, L., and Vanschoren, J. (eds.), *Automated Machine Learning: Methods, Systems, Challenges*, chapter 1, pp. 3 – 38. Springer, 2019. Available for free at <http://automl.org/book>. 1, B
- Feurer, M., Eggenesperger, K., Falkner, S., Lindauer, M., and Hutter, F. Auto-Sklearn 2.0: Hands-free automl via meta-learning. *Journal of Machine Learning Research*, 23(261):1–61, 2022. B
- Garnett, R. *Bayesian Optimization*. Cambridge University Press, 2023. 1, B
- Gijssbers, P., Bueno, M., Coors, S., LeDell, E., Poirier, S., Thomas, J., Bischl, B., and Vanschoren, J. AMLB: an automl benchmark. *Journal of Machine Learning Research*, 25(101):1–65, 2024. 4.1
- Giné, E. and Nickl, R. *Mathematical Foundations of Infinite-Dimensional Statistical Models*, volume 40. Cambridge University Press, 2016. C.2
- Grinsztajn, L., Oyallon, E., and Varoquaux, G. Why do tree-based models still outperform deep learning on typical tabular data? In *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*, pp. 507–520, 2022. 4.1
- Guyon, I., Alamdari, A., Dror, G., and Buhmann, J. Performance prediction challenge. In *The 2006 IEEE International Joint Conference on Neural Network Proceedings*, 2006. 1
- Guyon, I., Saffari, A., Dror, G., and Cawley, G. Model selection: Beyond the Bayesian/Frequentist divide. *Journal of Machine Learning Research*, 11:61–87, 2010. B

- Guyon, I., Bennett, K., Cawley, G., Escalante, H. J., Escalera, S., Ho, T. K., Macià, N., Ray, B., Saeed, M., Statnikov, A., and Viegas, E. Design of the 2015 ChaLearn AutoML challenge. In *2015 International Joint Conference on Neural Networks (IJCNN'15)*, pp. 1–8. International Neural Network Society and IEEE Computational Intelligence Society, IEEE, 2015. B
- Guyon, I., Sun-Hosoya, L., Boullé, M., Escalante, H., Escalera, S., Liu, Z., Jajetic, D., Ray, B., Saeed, M., Sebag, M., Statnikov, A., Tu, W., and Viegas, E. Analysis of the AutoML Challenge Series 2015-2018. In Hutter, F., Kotthoff, L., and Vanschoren, J. (eds.), *Automated Machine Learning: Methods, Systems, Challenges*, chapter 10, pp. 177–219. Springer, 2019. Available for free at <http://automl.org/book>. B
- Hansen, N. and Ostermeier, A. Completely derandomized self-adaptation in evolution strategies. *Evolutionary C.*, 9(2):159–195, 2001. 1
- Hothorn, T., Leisch, F., Zeileis, A., and Hornik, K. The design and analysis of benchmark experiments. *Journal of Computational and Graphical Statistics*, 14(3):675–699, 2005. 4.1, F.1
- Igel, C. A note on generalization loss when evolving adaptive pattern recognition systems. *IEEE Transactions on Evolutionary Computation*, 17(3):345–352, 2012. 5, 5
- Jamieson, K. and Talwalkar, A. Non-stochastic best arm identification and Hyperparameter Optimization. In Gretton, A. and Robert, C. (eds.), *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics (AISTATS'16)*, volume 51. Proceedings of Machine Learning Research, 2016. B
- Kadra, A., Janowski, M., Wistuba, M., and Grabocka, J. Scaling laws for hyperparameter optimization. In Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Advances in Neural Information Processing Systems*, volume 36, pp. 47527–47553, 2023. B
- Kallenberg, O. *Foundations of modern probability*, volume 2. Springer, 1997. D
- Klein, A., Falkner, S., Bartels, S., Hennig, P., and Hutter, F. Fast Bayesian optimization of machine learning hyperparameters on large datasets. In Singh, A. and Zhu, J. (eds.), *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics (AISTATS'17)*, volume 54. Proceedings of Machine Learning Research, 2017. B
- Koch, P., Konen, W., Flasch, O., and Bartz-Beielstein, T. Optimizing support vector machines for stormwater prediction. Technical Report TR10-2-007, Technische Universität Dortmund, 2010. Proceedings of Workshop on Experimental Methods for the Assessment of Computational Systems joint to PPSN2010. 5, 5
- Kohli, R., Feurer, M., Bischl, B., Eggenberger, K., and Hutter, F. Towards quantifying the effect of datasets for benchmarking: A look at tabular machine learning. In *Data-centric Machine Learning (DMLR) workshop at the International Conference on Learning Representations (ICLR)*, 2024. 4.1
- Lang, M., Kotthaus, H., Marwedel, P., Weihs, C., Rahnenführer, J., and Bischl, B. Automatic model selection for high-dimensional survival analysis. *Journal of Statistical Computation and Simulation*, 85:62–76, 2015. B
- Larcher, C. and Barbosa, H. Evaluating models with dynamic sampling holdout in auto-ml. *SN Computer Science*, 3(506), 2022. 1
- Li, L., Jamieson, K., DeSalvo, G., Rostamizadeh, A., and Talwalkar, A. Hyperband: A novel bandit-based approach to Hyperparameter Optimization. *Journal of Machine Learning Research*, 18(185):1–52, 2018. B
- Lindauer, M., Eggenberger, K., Feurer, M., Biedenkapp, A., Deng, D., Benjamins, C., Ruhkopf, T., Sass, R., and Hutter, F. SMAC3: A versatile bayesian optimization package for Hyperparameter Optimization. *Journal of Machine Learning Research*, 23(54):1–9, 2022. 4, 4.1, B
- Loshchilov, I. and Hutter, F. CMA-ES for Hyperparameter Optimization of deep neural networks. In *International Conference on Learning Representations Workshop track*, 2016. Published online: [iclr.cc](https://arxiv.org/abs/1603.04467). B

- Lévesque, J. *Bayesian Hyperparameter Optimization: Overfitting, Ensembles and Conditional Spaces*. PhD thesis, Université Laval, 2018. 1, 5, 5
- Makarova, A., Shen, H., Perrone, V., Klein, A., Faddoul, J., Krause, A., Seeger, M., and Archambeau, C. Automatic termination for hyperparameter optimization. In Guyon, I., Lindauer, M., van der Schaar, M., Hutter, F., and Garnett, R. (eds.), *Proceedings of the First International Conference on Automated Machine Learning*. Proceedings of Machine Learning Research, 2022. 5
- Mallik, N., Bergman, E., Hvarfner, C., Stoll, D., Janowski, M., Lindauer, M., Nardi, L., and Hutter, F. PriorBand: Practical hyperparameter optimization in the age of deep learning. In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Proceedings of the 36th International Conference on Advances in Neural Information Processing Systems (NeurIPS'23)*. Curran Associates, 2023. B
- McElfresh, D., Khandagale, S., Valverde, J., Prasad C., V., Ramakrishnan, G., Goldblum, M., and White, C. When do neural nets outperform boosted trees on tabular data? In Oh, A., Neumann, T., Globerson, A., Saenko, K., Hardt, M., and Levine, S. (eds.), *Proceedings of the 36th International Conference on Advances in Neural Information Processing Systems (NeurIPS'23)*, pp. 76336–76369. Curran Associates, 2023. 4.1, F.2
- Mohr, F., Wever, M., and Hüllermeier, E. ML-Plan: Automated machine learning via hierarchical planning. *Machine Learning*, 107(8-10):1495–1515, 2018. 5, B
- Molinaro, A., Simon, R., and Pfeiffer, R. Prediction error estimation: A comparison of resampling methods. *Bioinformatics*, 21(15):3301–3307, 2005. B
- Nadeau, C. and Bengio, Y. Inference for the generalization error. In Solla, S., Leen, T., and Müller, K. (eds.), *Proceedings of the 13th International Conference on Advances in Neural Information Processing Systems (NeurIPS'99)*. The MIT Press, 1999. 1
- Nadeau, C. and Bengio, Y. Inference for the generalization error. *Machine Learning*, 52:239–281, 2003. 1
- Ng, A. Preventing “overfitting” of cross-validation data. In Fisher, D. H. (ed.), *Proceedings of the Fourteenth International Conference on Machine Learning (ICML'97)*, pp. 245–253. Morgan Kaufmann Publishers, 1997. 5, 5
- Pfisterer, F., Schneider, L., Moosbauer, J., Binder, M., and Bischl, B. YAHPO Gym – an efficient multi-objective multi-fidelity benchmark for hyperparameter optimization. In Guyon, I., Lindauer, M., van der Schaar, M., Hutter, F., and Garnett, R. (eds.), *Proceedings of the First International Conference on Automated Machine Learning*. Proceedings of Machine Learning Research, 2022. B, 5
- Pineda Arango, S., Jomaa, H., Wistuba, M., and Grabocka, J. HPO-B: A large-scale reproducible benchmark for black-box HPO based on OpenML. In Vanschoren, J. and Yeung, S. (eds.), *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks*. Curran Associates, 2021. B, 5
- Probst, P., Boulesteix, A., and Bischl, B. Tunability: Importance of hyperparameters of machine learning algorithms. *Journal of Machine Learning Research*, 20(53):1–32, 2019. 1
- Prokhorenkova, L., Gusev, G., Vorobev, A., Dorogush, A., and Gulin, A. Catboost: Unbiased boosting with categorical features. In Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Proceedings of the 31st International Conference on Advances in Neural Information Processing Systems (NeurIPS'18)*, pp. 6639–6649. Curran Associates, 2018. 4.1
- Purucker, L. and Beel, J. CMA-ES for post hoc ensembling in automl: A great success and salvageable failure. In Faust, A., Garnett, R., White, C., Hutter, F., and Gardner, J. R. (eds.), *Proceedings of the Second International Conference on Automated Machine Learning*, volume 224 of *Proceedings of Machine Learning Research*, pp. 1/1–23. PMLR, 2023. 5
- Quinlan, J. and Cameron-Jones, R. Oversearching and layered search in empirical learning. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, volume 2 of *IJCAI'95*, pp. 1019–1024, 1995. 5

- Rao, R., Fung, G., and Rosales, R. On the dangers of cross-validation. an experimental evaluation. In *Proceedings of the 2008 SIAM International Conference on Data Mining (SDM)*, pp. 588–596, 2008. B
- Salinas, D., Seeger, M., Klein, A., Perrone, V., Wistuba, M., and Archambeau, C. Syne Tune: A library for large scale hyperparameter tuning and reproducible research. In Guyon, I., Lindauer, M., van der Schaar, M., Hutter, F., and Garnett, R. (eds.), *Proceedings of the First International Conference on Automated Machine Learning*, pp. 16–1. Proceedings of Machine Learning Research, 2022. B
- Schaffer, C. Selecting a classification method by cross-validation. *Machine Learning Journal*, 13: 135–143, 1993. B
- Swersky, K., Snoek, J., and Adams, R. Freeze-thaw Bayesian optimization. *arXiv:1406.3896 [stats.ML]*, 2014. B
- Talagrand, M. *The generic chaining: upper and lower bounds of stochastic processes*. Springer Science & Business Media, 2005. C.2
- Thornton, C., Hutter, F., Hoos, H., and Leyton-Brown, K. Auto-WEKA: Combined selection and Hyperparameter Optimization of classification algorithms. In Dhillon, I., Koren, Y., Ghani, R., Senator, T., Bradley, P., Parekh, R., He, J., Grossman, R., and Uthurusamy, R. (eds.), *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD’13)*, pp. 847–855. ACM Press, 2013. 5, B
- Turner, R., Eriksson, D., McCourt, M., Kiili, J., Laaksonen, E., Xu, Z., and Guyon, I. Bayesian optimization is superior to random search for machine learning hyperparameter tuning: Analysis of the Black-Box Optimization Challenge 2020. In Escalante, H. and Hofmann, K. (eds.), *Proceedings of the Neural Information Processing Systems Track Competition and Demonstration*, pp. 3–26. Curran Associates, 2021. 4.1
- van der Vaart, A. *Asymptotic statistics*, volume 3. Cambridge university press, 2000. C.1
- van Erven, T., Grünwald, P., Mehta, N., Reid, M., and Williamson, R. Fast rates in statistical and online learning. *Journal of Machine Learning Research*, 16(54):1793–1861, 2015. C.1
- van Rijn, J. and Hutter, F. Hyperparameter importance across datasets. In Guo, Y. and Farooq, F. (eds.), *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD’18)*, pp. 2367–2376. ACM Press, 2018. 1
- Vanschoren, J., van Rijn, J., Bischl, B., and Torgo, L. OpenML: Networked science in machine learning. *SIGKDD Explorations*, 15(2):49–60, 2014. 4
- Wainer, J. and Cawley, G. Empirical Evaluation of Resampling Procedures for Optimising SVM Hyperparameters. *Journal of Machine Learning Research*, 18:1–35, 2017. B
- Wainwright, M. *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge university press, 2019. C.2
- Wistuba, M., Schilling, N., and Schmidt-Thieme, L. Scalable Gaussian process-based transfer surrogates for Hyperparameter Optimization. *Machine Learning*, 107(1):43–78, 2018. G.1
- Wu, J., Toscano-Palmerin, S., Frazier, P., and Wilson, A. Practical multi-fidelity Bayesian optimization for hyperparameter tuning. In Peters, J. and Sontag, D. (eds.), *Proceedings of The 36th Uncertainty in Artificial Intelligence Conference (UAI’20)*, pp. 788–798. PMLR, 2020. B
- Zheng, A. and Bilenko, M. Lazy paired hyper-parameter tuning. In Rossi, F. (ed.), *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI’13)*, pp. 1924–1931, 2013. 5, B
- Zimmer, L., Lindauer, M., and Hutter, F. Auto-Pytorch: Multi-fidelity metalearning for efficient and robust AutoDL. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43:3079–3090, 2021. 4.1, F.2
- Zou, H. and Hastie, T. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 67(2):301–320, 2005. 4.1

## A Notation

Table 2: Notation table. We discuss all symbols used in the main paper.

$\mathbf{X}_i$	Random vector, describing the features
$Y_i$	Random variable, describing the target
$\mathbf{Z}_i = (\mathbf{X}_i, Y_i)$	Data point
$\mathcal{D} = \{\mathbf{Z}_i\}_{i=1}^n$	Dataset consisting of <i>iid</i> random variables
$n$	Number of observations
$g$	Inducer/ML algorithm
$h$	Model, created by the inducer via $h = g_\lambda(\mathcal{D})$
$\lambda$	Hyperparameter configuration
$\Lambda$	Finite set of all hyperparameter configurations
$J$	$ \Lambda $ , i.e., the number of hyperparameter configurations
$g_\lambda$	Hyperparameterized inducer
$\mu(\lambda)$	Expected loss of a hyperparameterized inducer on the distribution of a dataset
$\ell(\mathbf{Z}, h)$	Loss of a model $h$ on a fresh observation $\mathbf{Z}$
$M$	Number of folds in M-fold cross-validation
$\alpha$	Percentage of samples to be used for validation
$\mathcal{I}_{1,j}, \dots, \mathcal{I}_{M,j} \subset \{1, \dots, n\}$	$M$ sets of validation indices, to be used for evaluating $\lambda_j$
$\mathcal{V}_{m,j}$	Validation data for fold $m$ and configuration $\lambda_j$
$\mathcal{T}_{m,j}$	Training data for fold $m$ and configuration $\lambda_j$
$L(\mathcal{V}_{m,j}, g_\lambda(\mathcal{T}_{m,j}))$	Validation loss for fold $m$ and configuration $\lambda_j$
$\hat{\mu}(\lambda_j)$	M-fold validation loss
$\sigma^2$	Increase in variance of validation loss caused by resampling
$\tau^2$	Decrease in correlation among validation losses caused by reshuffling
$\tau_{i,j,M}$	Resampling-related component of validation loss covariance
$K(\cdot, \cdot)$	Kernel capturing the covariance of the pointwise losses between two HPCs
$\epsilon(\lambda_j)$	Zero-mean Gaussian process, see Equation (2)
$d$	Number of hyperparameters
$\kappa$	Curvature constant of covariance kernel
$\eta$	Density of hyperparameter set $\Lambda$
$m$	Local curvature at the minimum of the loss surface $\mu$
$\sigma$	Lower bound on the noise level
$B(\tau)$	Part of the regret bound penalizing reshuffling
$A(\tau)$	Part of the regret bound rewarding reshuffling

## B Extended Related Work

Due to the black box nature of the HPO problem (Feurer & Hutter, 2019; Bischl et al., 2023), gradient free, zeroth-order optimization algorithms such as BO (Garnett, 2023), Evolutionary Strategies (Loshchilov & Hutter, 2016) or a simple random search (Bergstra & Bengio, 2012) have become standard optimization algorithms to tackle vanilla HPO problems.

In the last decade, most research on HPO has been concerned with constructing new algorithms that excel at finding configurations with a low estimated generalization error. Examples include BO variants such as HEBO (Cowen-Rivers et al., 2022) or SMAC3 (Lindauer et al., 2022). Another direction of HPO research has been concerned with speeding up the HPO process to allow more efficient spending of compute resources. Multifidelity HPO, for example, turns the black box optimization problem into a gray box one by making use of lower fidelity approximations to the target function, i.e., using fewer numbers of epochs or subsets of the data for cheap low-fidelity evaluations that approximate the costly high-fidelity evaluation. Examples include bandit-based budget allocation algorithms such as Successive Halving (Jamieson & Talwalkar, 2016), Hyperband (Li et al., 2018) and their extensions that use non-random search mechanisms (Falkner et al., 2018; Awad et al., 2021; Mallik et al., 2023) or algorithms making use of multi-fidelity information in the context of BO (Swersky et al., 2014; Klein et al., 2017; Wu et al., 2020; Kadra et al., 2023). Several works address the problem of speeding up cross-validation techniques and use techniques that could be described as grey box optimization techniques. Besides the ones mentioned in the main paper (Thornton et al., 2013; Zheng & Bilenko, 2013), it is possible to employ racing techniques for model selection in machine learning as demonstrated by Lang et al. (2015), and there has been a recent interest in methods that adapt the cost of running full cross-validation procedures (Bergman et al., 2024; Buczak et al., 2024).

When addressing the problem of HPO, we must acknowledge an inherent mismatch between the explicit objective we optimize – namely, the estimated generalization performance of a model – and the actual implicit optimization goal, which is to identify a configuration that yields the best



generalization performance on new, unseen data. Typically, evaluations and comparisons of different HPO algorithms focus exclusively on the final best validation performance (i.e., the objective that is directly optimized), even though an unbiased estimate of performance on an external unseen test set might be available. While this approach is logical for assessing the efficacy of an optimization algorithm based on the metric it seeks to improve, relying solely on finding an optimal validation configuration is beneficial only if there is reason to assume a strong correlation between the optimized validation performance and true generalization ability on new, unseen test data. This discrepancy can be found deeply within the HPO community, where the evaluation of HPO algorithms on standard benchmark libraries is usually done solely with respect to the validation performance (Eggensperger et al., 2021; Pineda Arango et al., 2021; Salinas et al., 2022; Pfisterer et al., 2022).<sup>5</sup> This relationship between validation performance (i.e., the estimated generalization error derived from resampling) and true generalization performance (e.g., assessed through an outer holdout test set or additional resampling) of an optimal validation configuration found during HPO remains a largely unexplored area of research.

In general, little research has focused on the selection of resampling types, let alone the automated selection of resampling types (Guyon et al., 2010; Feurer et al., 2022). While we usually expect that a more intensive resampling will reduce the variance of the estimated generalization error and thereby improve the (rank) correlation between optimized validation and unbiased outer test performance within HPCs, this benefit is naturally offset by a higher computational expense. Overall, there is little research on which resampling method to use in practice for model selection, and we only know of a study for support vector machines (Wainer & Cawley, 2017), a simulation study for clinical prediction models (Dunias et al., 2024), a study on feature selection (Molinario et al., 2005) and a study on fast CV (Bergman et al., 2024). In addition, ML-Plan (Mohr et al., 2018) proposed a two-stage procedure. In a first stage (search), the tool uses planning on hierarchical task networks to find promising machine learning pipelines on 70% of the training data. In a second step (selection), it uses 100% of the training data and retrains the most promising candidates from the search step. Finally, it uses a combination of the internal generalization error estimation that was used during search and the 0.75 percentile of the generalization error estimation from the selection step to make a more unbiased selection of the final model. The paper found that this improves performance over using only regular cross-validation for search and selection. The general consensus, that is in agreement with our findings, is that CV or repeated CV generally leads to better generalization performance. In addition, while there are theoretical works that compare the accuracy of estimating the generalization error of holdout and CV (Blum et al., 1999), our goal is to correctly identify a single solution, which generalizes well, see the excellent survey by Arlot & Celisse (2010) for a discussion on this topic.

Bouthillier et al. (2021) studied the sources of variance in machine learning experiments, and find that the split into training and test data has the largest impact. Consequently, they suggest to reshuffle the data prior to splitting it into the training, which is then used for HPO, and the test set. We followed their suggestion when designing our experiments and draw a new test sample for every replication, see Section 4.1 and Appendix F. This dependence on the exact split was further already discussed in the context of how much the outcome of a statistical test on results of machine learning experiments depended on the exact train-test split (Bouckaert, 2004).

Finally, the first warning against comparing too many hypothesis using cross-validation was raised by Schaffer (1993), and in addition to the works discussed in Section 5 in the main paper, also picked up by Rao et al. (2008); Cawley & Talbot (2010). Moreover, the problem of finding a correct "upper objective" in a bilevel optimization problem has been noted (Guyon et al., 2010, 2015, 2019). Also, in the related field of algorithm configuration the problem has been identified (Eggensperger et al., 2019).

## B.1 Current Treatment of Resamplings in HPO Libraries and Software

In Table 3, we provide a brief summary of how resampling is handled in popular HPO libraries and software.<sup>6</sup> For each library, we checked whether the core functionality, examples, or tutorials mention

<sup>5</sup>We admit that these benchmark libraries implement efficient benchmarking methods such as surrogate (Eggensperger et al., 2018; Pfisterer et al., 2022) or tabular benchmarks (Pineda Arango et al., 2021). It would be possible to adapt them to return the test performance, however, changes in the HPO evaluation protocol, such as the one we propose, would not be feasible.

<sup>6</sup>This summary is not exhaustive but reflects the general consensus observed in widely-used software.

the possibility of reshuffling the resampling during HPO or if the resampling is considered fixed. If reshuffling is used in an example, mentioned, or if core functionality uses it, we mark it with a ✓. If it is unclear or inconsistent across examples and core functionality, we mark it with a ?. Otherwise, we use a ✗. Our conclusion is that the concept of reshuffling resampling generally receives little attention.

Table 3: Exemplary Treatment of Resamplings in HPO Libraries and Software

Software	Reshuffled?	Reference(s)
sklearn	✗	GridSearchCV <sup>1</sup> / RandomizedSearchCV <sup>2</sup>
HEBO	✗	sklearn_tuner <sup>3</sup>
optuna	?	Inconsistency between examples <sup>4,5,6</sup>
bayesian-optimization	✗	sklearn Example <sup>7,8</sup>
ax	✗	CNN Example <sup>9</sup>
spearmint	✗	No official HPO Examples
scikit-optimize	✗	BO for GBT Example <sup>7,10</sup>
SMAC3	✗	SVM Example <sup>7,11</sup>
dragonfly	✗	Tree Based Ensemble Example <sup>12</sup>
aws sagemaker	✗	Blog Post <sup>13</sup>
raytune	?	Inconsistency between examples <sup>14,15</sup>
hyperopt(-sklearn)	?	Cost Function Logic <sup>16</sup>

✗: no reshuffling, ?: both reshuffling and no reshuffling or unclear, ✓: reshuffling

- <sup>1</sup> [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_search.py#L1263](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_search.py#L1263)
- <sup>2</sup> [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_search.py#L1644](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_search.py#L1644)
- <sup>3</sup> [https://github.com/huawei-noah/HEBO/blob/b60f41aa862b4c5148e31ab4981890da6d41f2b1/HEBO/hebo/sklearn\\_tuner.py#L73](https://github.com/huawei-noah/HEBO/blob/b60f41aa862b4c5148e31ab4981890da6d41f2b1/HEBO/hebo/sklearn_tuner.py#L73)
- <sup>4</sup> [https://github.com/optuna/optuna-integration/blob/15e6b0ec6d9a0d7f572ad387be8478c56257bef7/optuna\\_integration/sklearn/sklearn.py#L223](https://github.com/optuna/optuna-integration/blob/15e6b0ec6d9a0d7f572ad387be8478c56257bef7/optuna_integration/sklearn/sklearn.py#L223) here sklearn's `cross_validate` is used which by default does not reshuffle the resampling [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_validation.py#L186](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_validation.py#L186)
- <sup>5</sup> [https://github.com/optuna/optuna-examples/blob/dd56b9692e6d1f4fa839332edbcdd93fd48c16d8/pytorch/pytorch\\_simple.py#L79](https://github.com/optuna/optuna-examples/blob/dd56b9692e6d1f4fa839332edbcdd93fd48c16d8/pytorch/pytorch_simple.py#L79) here, data loaders for train and valid are instantiated within the objective of the trial but the data within the loaders is fixed
- <sup>6</sup> [https://github.com/optuna/optuna-examples/blob/dd56b9692e6d1f4fa839332edbcdd93fd48c16d8/xgboost/xgboost\\_simple.py#L22](https://github.com/optuna/optuna-examples/blob/dd56b9692e6d1f4fa839332edbcdd93fd48c16d8/xgboost/xgboost_simple.py#L22) here, the train validation split is performed within the objective of the trial and no seed is set which results in reshuffling [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_split.py#L2597](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_split.py#L2597)
- <sup>7</sup> functionality relies on sklearn's `cross_val_score` which by default does not reshuffle the resampling [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_validation.py#L631](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_validation.py#L631)
- <sup>8</sup> [https://github.com/bayesian-optimization/BayesianOptimization/blob/c7e5c3926944fc6011ae7ace29f7b5ed0f9c983b/examples/sklearn\\_example.py#L32](https://github.com/bayesian-optimization/BayesianOptimization/blob/c7e5c3926944fc6011ae7ace29f7b5ed0f9c983b/examples/sklearn_example.py#L32)
- <sup>9</sup> [https://github.com/facebook/Ax/blob/ac44a6661f535dd3046954f8fd8701327f4a53e2/tutorials/tune\\_cnn\\_service.ipynb#L39](https://github.com/facebook/Ax/blob/ac44a6661f535dd3046954f8fd8701327f4a53e2/tutorials/tune_cnn_service.ipynb#L39) and [https://github.com/facebook/Ax/blob/ac44a6661f535dd3046954f8fd8701327f4a53e2/ax/utilis/tutorials/cnn\\_utils.py#L154](https://github.com/facebook/Ax/blob/ac44a6661f535dd3046954f8fd8701327f4a53e2/ax/utilis/tutorials/cnn_utils.py#L154)
- <sup>10</sup> <https://github.com/scikit-optimize/scikit-optimize/blob/a2369ddb332d16d8ff173b12404b03fea472492/examples/hyperparameter-optimization.py#L82C21-L82C36>
- <sup>11</sup> [https://github.com/automl/SMAC3/blob/9aaa8e94a5b3a9657737a87b903ee96c683cc42c/examples/1\\_basics/2\\_svm\\_cv.py#L63](https://github.com/automl/SMAC3/blob/9aaa8e94a5b3a9657737a87b903ee96c683cc42c/examples/1_basics/2_svm_cv.py#L63)
- <sup>12</sup> [https://github.com/dragonfly/dragonfly/blob/3eef7d30bcc2e56f2221a624bd8ec7f933f81e40/examples/tree\\_reg/skltree.py#L111](https://github.com/dragonfly/dragonfly/blob/3eef7d30bcc2e56f2221a624bd8ec7f933f81e40/examples/tree_reg/skltree.py#L111)
- <sup>13</sup> <https://aws.amazon.com/blogs/architecture/field-notes-build-a-cross-validation-machine-learning-model-pipeline-at-scale-with-amazon-sagemaker/>
- <sup>14</sup> <https://github.com/ray-project/ray/blob/3f5aa5c4642eeb12447d9de5dce2208512312f3/doc/source/tune/examples/tune-pytorch-cifar.ipynb#L120> here, data loaders for train and valid are instantiated within the objective but the data within the loaders are fixed
- <sup>15</sup> <https://github.com/ray-project/ray/blob/3f5aa5c4642eeb12447d9de5dce2208512312f3/doc/source/tune/examples/tune-xgboost.ipynb#L335> here, the train validation split is performed within the objective and no seed is set which results in reshuffling [https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model\\_selection/\\_split.py#L2597](https://github.com/scikit-learn/scikit-learn/blob/8721245511de2f225ff5f9aa5f5fadce663cd4a3/sklearn/model_selection/_split.py#L2597)
- <sup>16</sup> [https://github.com/hyperopt/hyperopt-sklearn/blob/4bc286479677a0bfd2178dac4546ea268b3f3b77/hpsklearn/estimator/\\_cost\\_fn.py#L144](https://github.com/hyperopt/hyperopt-sklearn/blob/4bc286479677a0bfd2178dac4546ea268b3f3b77/hpsklearn/estimator/_cost_fn.py#L144) dependence on random seed which by default is not set and there is no discussion of reshuffling and behavior is somewhat unclear

## C Proofs of the Main Results

### C.1 Proof of Theorem 2.1

We impose *stability* assumptions on the learning algorithm similar to Bayle et al. (2020); Austern & Zhou (2020). Let  $\mathbf{Z}, \mathbf{Z}_1, \dots, \mathbf{Z}_n, \mathbf{Z}'_1$ , be *iid* random variables. Define  $\mathcal{T} = \{\mathbf{Z}_i\}_{i=1}^n$ , and  $\mathcal{T}'$  as  $\mathcal{T}$  but with  $\mathbf{Z}_n$  replaced by the independent copy  $\mathbf{Z}'_n$ . Define

$$\tilde{\ell}_n(\mathbf{z}, \boldsymbol{\lambda}) = \ell(\mathbf{z}, g_{\boldsymbol{\lambda}}(\mathcal{T})) - \mathbb{E}[\ell(\mathbf{Z}, g_{\boldsymbol{\lambda}}(\mathcal{T})) \mid \mathcal{T}],$$

assume that each  $g_{\boldsymbol{\lambda}}(\mathcal{T})$  is invariant to the ordering in  $\mathcal{T}$ ,  $\ell$  is bounded, and

$$\max_{\boldsymbol{\lambda} \in \Lambda} \mathbb{E}\{[\tilde{\ell}(\mathbf{Z}, g_{\boldsymbol{\lambda}}(\mathcal{T})) - \tilde{\ell}(\mathbf{Z}, g_{\boldsymbol{\lambda}}(\mathcal{T}'))]^2\} = o(1/n). \quad (4)$$

This *loss stability* assumption is rather mild, see Bayle et al. (2020) for an extensive discussion. Further, define the risk  $R(g) = \mathbb{E}[\ell(\mathbf{Z}, g)]$  and assume that for every  $\boldsymbol{\lambda} \in \Lambda$ , there is a prediction rule  $g_{\boldsymbol{\lambda}}^*$  such that

$$\max_{\boldsymbol{\lambda} \in \Lambda} \mathbb{E}[|R(g_{\boldsymbol{\lambda}}(\mathcal{T})) - R(g_{\boldsymbol{\lambda}}^*)|] = o(1/\sqrt{n}). \quad (5)$$

This assumption requires  $g_{\boldsymbol{\lambda}}(\mathcal{T})$  to converge to some fixed prediction rule sufficiently fast and serves as a reasonable working condition for our purposes. It is satisfied, for example, when  $\ell$  is the square loss and  $g_{\boldsymbol{\lambda}}$  is an empirical risk minimizer over a hypothesis class  $\mathcal{G}_{\boldsymbol{\lambda}}$  with finite VC-dimension. For further examples, see, e.g., Bousquet & Zhivotovskiy (2021), van Erven et al. (2015), and references therein. The assumption could be relaxed, but this would lead to a more complicated limiting distribution but with the same essential interpretation.

**Theorem C.1.** *Under assumptions (4) and (5), it holds*

$$\sqrt{n}(\hat{\boldsymbol{\mu}}(\boldsymbol{\lambda}_j) - \boldsymbol{\mu}(\boldsymbol{\lambda}_j))_{j=1}^J \rightarrow_d \mathcal{N}(0, \Sigma),$$

where

$$\begin{aligned} \Sigma_{j,j'} &= \tau_{i,j,M} \lim_{n \rightarrow \infty} \text{Cov}[\bar{\ell}_n(\mathbf{Z}, \boldsymbol{\lambda}_j), \bar{\ell}_n(\mathbf{Z}, \boldsymbol{\lambda}_{j'})], \\ \tau_{j,j',M} &= \lim_{n \rightarrow \infty} \frac{1}{nM^2\alpha^2} \sum_{i=1}^n \sum_{m=1}^M \sum_{m'=1}^M \Pr(i \in \mathcal{I}_{m,j} \cap \mathcal{I}_{m',j'}). \end{aligned}$$

*Proof.* Define

$$\tilde{\boldsymbol{\mu}}(\boldsymbol{\lambda}_j) = \frac{1}{M} \sum_{m=1}^M \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})) \mid \mathcal{T}_{m,j}].$$

By the triangle inequality (first and second step), Jensen's inequality (third step), and (5) (last step),

$$\begin{aligned} & \mathbb{E}[|\tilde{\boldsymbol{\mu}}(\boldsymbol{\lambda}_j) - \boldsymbol{\mu}(\boldsymbol{\lambda}_j)|] \\ & \leq \max_{1 \leq m \leq M} \mathbb{E} \left[ \left| \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})) \mid \mathcal{T}_{m,j}] - \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j}))] \right| \right] \\ & \leq \max_{1 \leq m \leq M} \mathbb{E} \left[ \left| \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})) \mid \mathcal{T}_{m,j}] - \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}^*(\mathcal{T}_{m,j}))] \right| \right] \\ & \quad + \max_{1 \leq m \leq M} \mathbb{E} \left[ \left| \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j}))] - \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}^*(\mathcal{T}_{m,j}))] \right| \right] \\ & \leq 2 \max_{1 \leq m \leq M} \mathbb{E} \left[ \left| \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})) \mid \mathcal{T}_{m,j}] - \mathbb{E}[L(\mathcal{V}_{m,j}, g_{\boldsymbol{\lambda}_j}^*(\mathcal{T}_{m,j}))] \right| \right] \\ & = 2 \max_{1 \leq m \leq M} \mathbb{E} \left[ \left| R(g_{\boldsymbol{\lambda}_j}(\mathcal{T}_{m,j})) - R(g_{\boldsymbol{\lambda}_j}^*) \right| \right] \\ & = o(1/\sqrt{n}). \end{aligned}$$

Next, assumption (4) together with Theorem 2 and Proposition 3 of Bayle et al. (2020) yield

$$\sqrt{n}(\hat{\boldsymbol{\mu}}(\boldsymbol{\lambda}_j) - \tilde{\boldsymbol{\mu}}(\boldsymbol{\lambda}_j)) - \frac{1}{M} \sum_{m=1}^M \frac{1}{\alpha\sqrt{n}} \sum_{i \in \mathcal{I}_{m,j}} \bar{\ell}_n(\mathbf{Z}_i, \boldsymbol{\lambda}_j) \rightarrow_p 0.$$

Now rewrite

$$\frac{1}{M\alpha\sqrt{n}} \sum_{m=1}^M \sum_{i \in \mathcal{I}_{m,j}} \bar{\ell}_n(\mathbf{Z}_i, \boldsymbol{\lambda}_j) = \frac{1}{M\alpha\sqrt{n}} \sum_{i=1}^n \underbrace{\sum_{m=1}^M \mathbb{1}(i \in \mathcal{I}_{m,j})}_{:=\xi_{i,n}^{(j)}} \bar{\ell}_n(\mathbf{Z}_i, \boldsymbol{\lambda}_j).$$

The sequence  $(\xi_{i,n}^{(j)})_{i=1}^n = (\xi_{i,n}^{(j)}, \dots, \xi_{i,n}^{(j)})_{i=1}^n$  is a triangular array of independent, centered, and bounded random vectors. Because  $\mathbb{1}(\mathbf{Z}_i \in \mathcal{V}_{m,j})$  and  $\mathbf{Z}_i$  are independent, it holds

$$\text{Cov}(\xi_{i,n}^{(j)}, \xi_{i,n}^{(j')}) = \sum_{m=1}^M \sum_{m'=1}^M \mathbb{E}[\mathbb{1}(i \in \mathcal{I}_{m,j} \cap \mathcal{I}_{m',j'})] \mathbb{E}[\bar{\ell}_n(\mathbf{Z}_i, \boldsymbol{\lambda}_j) \bar{\ell}_n(\mathbf{Z}_i, \boldsymbol{\lambda}_{j'})],$$

so

$$\lim_{n \rightarrow \infty} \text{Cov} \left[ \frac{1}{M\alpha\sqrt{n}} \sum_{i=1}^n \xi_{i,n}^{(j)}, \frac{1}{M\alpha\sqrt{n}} \sum_{i=1}^n \xi_{i,n}^{(j')} \right] = \lim_{n \rightarrow \infty} \frac{1}{nM^2\alpha^2} \sum_{i=1}^n \text{Cov} \left[ \xi_{i,n}^{(j)}, \xi_{i,n}^{(j')} \right] = \Sigma_{j,j'}.$$

Now the result follows from Lindeberg's central limit theorem for triangular arrays (e.g., van der Vaart, 2000, Proposition 2.27).  $\square$

## C.2 Proof of Theorem 2.2

We want to bound the probability that  $\mu(\hat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*)$  is large. For some  $\delta > 0$ , define the set of 'good' hyperparameters

$$\Lambda_\delta = \{\boldsymbol{\lambda}_j : \mu(\boldsymbol{\lambda}_j) - \mu(\boldsymbol{\lambda}^*) \leq \delta\}.$$

Now

$$\begin{aligned} \Pr \left( \mu(\hat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta \right) &= \Pr \left( \hat{\boldsymbol{\lambda}} \notin \Lambda_\delta \right) \\ &= \Pr \left( \min_{\boldsymbol{\lambda} \notin \Lambda_\delta} \hat{\mu}(\boldsymbol{\lambda}) < \min_{\boldsymbol{\lambda} \in \Lambda_\delta} \hat{\mu}(\boldsymbol{\lambda}) \right) \\ &\leq \Pr \left( \min_{\boldsymbol{\lambda} \notin \Lambda_\delta} \hat{\mu}(\boldsymbol{\lambda}) < \min_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \hat{\mu}(\boldsymbol{\lambda}) \right) \\ &= \Pr \left( \min_{\boldsymbol{\lambda} \notin \Lambda_\delta} \mu(\boldsymbol{\lambda}) + \epsilon(\boldsymbol{\lambda}) < \min_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \mu(\boldsymbol{\lambda}) + \epsilon(\boldsymbol{\lambda}) \right) \\ &\leq \Pr \left( \delta + \min_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) < \delta/2 + \min_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) \right) \\ &= \Pr \left( \min_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) - \min_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) < -\delta/2 \right) \\ &= \Pr \left( \max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) - \max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) > \delta/2 \right). \quad (\epsilon \stackrel{d}{=} -\epsilon) \end{aligned}$$

There is a tension between the two maxima. The more  $\boldsymbol{\lambda}$ 's there are in  $\Lambda_{\delta/2}$  and the less they are correlated, the more likely it is to find one  $\epsilon(\boldsymbol{\lambda})$  that is large. This makes the probability small. However, the less  $\epsilon$  is correlated, the larger is  $\max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda})$ , making the probability large. To formalize this, use the Gaussian concentration inequality (Talagrand, 2005, Lemma 2.1.3):

$$\begin{aligned} &\Pr \left( \max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) - \max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) > \delta/2 \right) \\ &\leq \Pr \left( 2 \left| \max_{\boldsymbol{\lambda} \in \Lambda} \epsilon(\boldsymbol{\lambda}) - \mathbb{E} \left[ \max_{\boldsymbol{\lambda} \in \Lambda} \epsilon(\boldsymbol{\lambda}) \right] \right| > \delta/2 - \mathbb{E} \left[ \max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) \right] + \mathbb{E} \left[ \max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) \right] \right) \\ &\leq 2 \exp \left\{ - \frac{(\delta/2 - \mathbb{E} [\max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda})] + \mathbb{E} [\max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda})])^2}{8\sigma^2} \right\}, \end{aligned}$$

provided  $\delta/2 - \mathbb{E} [\max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda})] + \mathbb{E} [\max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda})] \geq 0$ . We bound the two maxima separately.

### Lower Bound for Maximum over the Good Set

Recall the definition of  $m$  right before Theorem 2.2 and observe

$$\begin{aligned}\Lambda_{\delta/2} &= \{\boldsymbol{\lambda}: \mu(\boldsymbol{\lambda}) - \mu(\boldsymbol{\lambda}^*) \leq \delta/2\} \supset \{\boldsymbol{\lambda}: m\|\boldsymbol{\lambda} - \boldsymbol{\lambda}^*\|^2 \leq \delta/2\} = \{\boldsymbol{\lambda}: \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^*\| \leq (\delta/2m)^{1/2}\} \\ &= B(\boldsymbol{\lambda}^*, (\delta/2m)^{1/2}).\end{aligned}$$

Pack the ball  $B(\boldsymbol{\lambda}^*, (\delta/2m)^{1/2})$  with smaller balls with radius  $\eta$ . We can always construct such a packing with at least  $(\delta/2m\eta^2)^{d/2}$  elements. By assumption, each small ball contains at least one element of  $\Lambda$ . Pick one element from each small ball and collect them into the set  $\Lambda'_{\delta/2}$ . By construction,  $|\Lambda'_{\delta/2}| \geq (\delta/2m\eta^2)^{d/2}$  and

$$\min_{\boldsymbol{\lambda} \neq \boldsymbol{\lambda}' \in \Lambda'_{\delta/2}} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\| \geq \eta.$$

Sudakov's minoration principle (e.g., Wainwright, 2019, Theorem 5.30) gives

$$\begin{aligned}\mathbb{E} \left[ \max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) \right] &\geq \frac{1}{2} \sqrt{\log |\Lambda'_{\delta/2}|} \min_{\{\boldsymbol{\lambda} \neq \boldsymbol{\lambda}'\} \cap \Lambda'_{\delta/2}} \sqrt{\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] } \\ &\geq \frac{1}{2} \sqrt{\log |\Lambda'_{\delta/2}|} \min_{\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\| \geq \eta} \sqrt{\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] }.\end{aligned}$$

In general,

$$\begin{aligned}\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] &= K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) + K(\boldsymbol{\lambda}', \boldsymbol{\lambda}') - 2\tau^2 K(\boldsymbol{\lambda}, \boldsymbol{\lambda}') \\ &= (1 - \tau^2)[K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) + K(\boldsymbol{\lambda}', \boldsymbol{\lambda}')] + \tau^2[K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) - K(\boldsymbol{\lambda}, \boldsymbol{\lambda}')] + \tau^2[K(\boldsymbol{\lambda}', \boldsymbol{\lambda}') - K(\boldsymbol{\lambda}, \boldsymbol{\lambda}')] \\ &\geq 2\underline{\sigma}^2(1 - \tau^2).\end{aligned}$$

Hence, we have

$$\min_{\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\| \geq \eta} \text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] \geq 2\underline{\sigma}^2(1 - \tau^2),$$

which implies

$$\mathbb{E} \left[ \max_{\boldsymbol{\lambda} \in \Lambda_{\delta/2}} \epsilon(\boldsymbol{\lambda}) \right] \geq \frac{1}{2} \underline{\sigma} \sqrt{d} \sqrt{1 - \tau^2} \sqrt{\log(\delta/2m\eta^2)} =: \sigma \sqrt{d} A(\tau, \delta)/2.$$

### Upper Bound for Maximum over the Bad Set

Dudley's entropy bound (e.g., Giné & Nickl, 2016, Theorem 2.3.6) gives

$$\mathbb{E} \left[ \max_{\boldsymbol{\lambda} \notin \Lambda_{\delta}} \epsilon(\boldsymbol{\lambda}) \right] \leq 12 \int_0^{\infty} \sqrt{\log N(s)} ds,$$

where  $N(s)$  is the minimum number of points  $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_{N(s)}$  such that

$$\sup_{\boldsymbol{\lambda} \in \Lambda} \min_{1 \leq k \leq N(s)} \sqrt{\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}_k)]} \leq s.$$

Note that

$$\sup_{\boldsymbol{\lambda}, \boldsymbol{\lambda}' \in \Lambda} \sqrt{\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] } \leq 2\sigma,$$

so  $N(s) = 1$  for all  $s \geq 2\sigma$ . For  $s^2 \leq 4\sigma^2(1 - \tau^2)$ , we can use the trivial bound  $N(s) \leq J$ . For  $s^2 > 4\sigma^2(1 - \tau^2)$ , cover  $\Lambda$  with  $\ell_2$ -balls of size  $(s/2\sigma\tau\kappa)$ . We can do this with less than  $N(s) \leq (6\sigma\kappa/s)^d \vee 1$  such balls. Let  $\boldsymbol{\lambda}_1, \dots, \boldsymbol{\lambda}_N$  be the centers of these balls. In general, it holds

$$\begin{aligned}\text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] &= K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) + K(\boldsymbol{\lambda}', \boldsymbol{\lambda}') - 2\tau^2 K(\boldsymbol{\lambda}, \boldsymbol{\lambda}') \\ &= (1 - \tau^2)[K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) + K(\boldsymbol{\lambda}', \boldsymbol{\lambda}')] + \tau^2[K(\boldsymbol{\lambda}, \boldsymbol{\lambda}) - K(\boldsymbol{\lambda}, \boldsymbol{\lambda}')] + \tau^2[K(\boldsymbol{\lambda}', \boldsymbol{\lambda}') - K(\boldsymbol{\lambda}, \boldsymbol{\lambda}')] \\ &\leq 2(1 - \tau^2)\sigma^2 + 2\tau^2\sigma^2\kappa^2\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|^2.\end{aligned}$$

For  $s^2 > 4\sigma^2(1 - \tau^2)$ , we thus have

$$\begin{aligned} \sup_{\boldsymbol{\lambda} \in \Lambda} \min_{1 \leq k \leq N(s)} \text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}_k)] &\leq \sup_{\|\boldsymbol{\lambda} - \boldsymbol{\lambda}'\|_2 \leq (s/2\tau\sigma\kappa)^2} \text{Var} [\epsilon(\boldsymbol{\lambda}) - \epsilon(\boldsymbol{\lambda}')] \\ &\leq 2(1 - \tau^2)\sigma^2 + 2\tau^2\sigma^2\kappa^2(s/2\tau\sigma\kappa)^2 \\ &\leq s^2, \end{aligned}$$

as desired. Now decompose the integral

$$\begin{aligned} \int_0^\infty \sqrt{\log N(s)} ds &= \int_0^{2\sigma\sqrt{1-\tau^2}} \sqrt{\log N(s)} ds + \int_{2\sigma\sqrt{1-\tau^2}}^{2\sigma} \sqrt{\log N(s)} ds \\ &\leq 2\sigma\sqrt{d}\sqrt{1-\tau^2}\sqrt{\log J} + \int_{2\sigma\sqrt{1-\tau^2}}^{2\sigma} \sqrt{\log N(s)} ds. \end{aligned}$$

For the second term, compute

$$\begin{aligned} \int_{\sigma\sqrt{1-\tau^2}}^{2\sigma} \sqrt{\log N(s)} ds &\leq \sqrt{d} \int_{2\sigma\sqrt{1-\tau^2}}^{2\sigma} \sqrt{\log(6\sigma\kappa/s)_+} ds \\ &= \sigma\sqrt{d} \int_{2\sqrt{1-\tau^2}}^2 \sqrt{\log(6\kappa/s)_+} ds \\ &\leq \sigma\sqrt{d} \left( \int_0^2 \log(6\kappa/s)_+ ds \right)^{1/2} \left( 2(1 - \sqrt{1 - \tau^2}) \right)^{1/2} \\ &= \sigma\sqrt{d} \sqrt{2 + 2\log(3\kappa)_+} \left( 2(1 - \sqrt{1 - \tau^2}) \right)^{1/2} \\ &= 2\sigma\sqrt{d} \sqrt{1 + \log(3\kappa)_+} \frac{\tau}{(1 + \sqrt{1 - \tau^2})^{1/2}} \\ &\leq 2\sigma\sqrt{d}\tau \sqrt{1 + \log(3\kappa)_+}. \end{aligned}$$

We have shown that

$$\mathbb{E} \left[ \max_{\boldsymbol{\lambda} \notin \Lambda_\delta} \epsilon(\boldsymbol{\lambda}) \right] \leq 24\sigma\sqrt{d} \left[ \sqrt{1 - \tau^2} \sqrt{\log J} + \tau \sqrt{1 + \log(3\kappa)_+} \right] =: \sigma\sqrt{d}B(\tau)/4.$$

### Integrating Probabilities

Summarizing the two previous steps, we have

$$\Pr \left( \mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta \right) \leq 2 \exp \left\{ - \frac{\left( \delta - \sigma\sqrt{d}[B(\tau) - A(\tau, \delta)] \right)^2}{36\sigma^2} \right\},$$

provided  $t \geq \sigma\sqrt{d}[B(\tau) - A(\tau, \delta)]$ . Now for any  $s \geq 0$  and  $t \geq 2e^{s^2}m\eta^2$ , it holds

$$A(\tau, s) \geq (\underline{\sigma}/\sigma)\sqrt{1 - \tau^2}s =: A(\tau)s.$$

In particular, if

$$t \geq 2e^{s^2}m\eta^2 + \sigma\sqrt{d}[B(\tau) - A(\tau)s] =: C,$$

we have

$$\Pr \left( \mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta \right) \leq 4 \exp \left\{ - \frac{\left( \delta - \sigma\sqrt{d}[B(\tau) - A(\tau)s] \right)^2}{36\sigma^2} \right\}.$$

Integrating the probability gives

$$\begin{aligned}
\mathbb{E}[\mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*)] &= \int_0^\infty \Pr\left(\mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta\right) d\delta \\
&= \int_0^C \Pr\left(\mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta\right) d\delta + \int_C^\infty \Pr\left(\mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*) > \delta\right) d\delta \\
&\leq C + \int_C^\infty \exp\left\{-\frac{\left(\delta - \sigma\sqrt{d}[B(\tau) - A(\tau)s]\right)^2}{36\sigma^2}\right\} d\delta \\
&\leq C + \sqrt{36}\sigma \\
&= 2e^{s^2} m\eta^2 + \sigma\sqrt{d}[B(\tau) - A(\tau)s] + 6\sigma.
\end{aligned}$$

### Simplifying

The bound can be optimized with respect to  $s$ , but the solution involves the Lambert  $W$ -function, which has no analytical expression. Instead choose  $s$  for simplicity as

$$s = \sqrt{\log\left(\frac{\sigma}{2m\eta^2}\right)_+}$$

which gives

$$\mathbb{E}[\mu(\widehat{\boldsymbol{\lambda}}) - \mu(\boldsymbol{\lambda}^*)] \leq \sigma\sqrt{d} \left[8 + B(\tau) - A(\tau)\sqrt{\log\left(\frac{\sigma}{2m\eta^2}\right)}\right]. \quad \square$$

## D Additional Results on the Density of Random HPC Grids

**Lemma D.1.** *Suppose that the  $J$  elements in  $\Lambda$  are drawn independently from a continuous density  $p$  with  $c := \min_{\|\boldsymbol{\lambda}\| \leq 1} p(\boldsymbol{\lambda}) > 0$ . Then with probability at least  $1 - \delta$ ,*

$$\eta \lesssim \left(\sqrt{\log(1/\delta)/J}\right)^{1/d},$$

and with probability 1,

$$\eta \lesssim \left(\sqrt{\log(J)/J}\right)^{1/d},$$

for all  $J$  sufficiently large.

*Proof.* We want to bound the probability that there is a  $\boldsymbol{\lambda}$  such that  $|B(\boldsymbol{\lambda}, \eta) \cap \Lambda| = 0$ . In what follows  $\boldsymbol{\lambda}$  is silently understood to have norm bounded by 1. Let  $\tilde{\boldsymbol{\lambda}}_1, \dots, \tilde{\boldsymbol{\lambda}}_N$  the centers of  $\eta/2$ -balls covering  $\{\|\boldsymbol{\lambda}\| \leq 1\}$ , for which we may assume  $N \leq (6/\eta)^d$ . For  $\tilde{\boldsymbol{\lambda}}_k$  the closest center to  $\boldsymbol{\lambda}$ , it holds

$$\|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\| \leq \|\boldsymbol{\lambda}' - \tilde{\boldsymbol{\lambda}}_k\| + \|\tilde{\boldsymbol{\lambda}}_k - \boldsymbol{\lambda}\| \leq \|\boldsymbol{\lambda}' - \tilde{\boldsymbol{\lambda}}_k\| + \eta/2,$$

so  $\|\boldsymbol{\lambda}' - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2$  implies  $\|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\| \leq \eta$ . We thus have

$$\begin{aligned}
\Pr(\exists \boldsymbol{\lambda}: |B(\boldsymbol{\lambda}, \eta) \cap \Lambda| = 0) &= \Pr\left(\inf_{\boldsymbol{\lambda}} \sum_{i=1}^J \mathbf{1}\{\|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}\| \leq \eta\} \leq 0\right) \\
&\leq \Pr\left(\min_{1 \leq k \leq N} \sum_{i=1}^J \mathbf{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0\right).
\end{aligned}$$

Further

$$\begin{aligned}
& \Pr \left( \min_{1 \leq k \leq N} \sum_{i=1}^J \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0 \right) \\
&= \Pr \left( \max_{1 \leq k \leq N} \sum_{i=1}^J -\mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \geq 0 \right) \\
&\leq \Pr \left( \max_{1 \leq k \leq N} \sum_{i=1}^J \mathbb{E} \left[ \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \right] - \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \geq J \inf_{\boldsymbol{\lambda}} \mathbb{E} \left[ \mathbb{1}\{\|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}\| \leq \eta/2\} \right] \right).
\end{aligned}$$

It holds

$$\begin{aligned}
\mathbb{E} \left[ \mathbb{1}\{\|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}\| \leq \eta/2\} \right] &= \Pr(\|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}\| \leq \eta/2) = \int_{\|\boldsymbol{\lambda}' - \boldsymbol{\lambda}\| \leq \eta/2} p(\boldsymbol{\lambda}') d\boldsymbol{\lambda}' \geq c \text{vol}(B(0, \eta/2)) \\
&= cv_d(\eta/2)^d,
\end{aligned}$$

where  $v_d = \text{vol}(B(0, 1))$ . Now the union bound and Hoeffding's inequality give

$$\begin{aligned}
\Pr \left( \min_{1 \leq k \leq N} \sum_{i=1}^J \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0 \right) &\leq N \exp \left( -\frac{Jc^2v_d^2(\eta/2)^{2d}}{2} \right) \\
&\leq (6/\eta)^d \exp \left( -\frac{Jc^2v_d^2(\eta/2)^{2d}}{2} \right).
\end{aligned}$$

Choosing

$$\eta = 2 \left( \sqrt{2 \log(3^d \sqrt{J} cv_d / \delta)} / \sqrt{J} cv_d \right)^{1/d}$$

gives

$$\Pr(\exists \boldsymbol{\lambda}: |B(\boldsymbol{\lambda}, \eta) \cap \Lambda| = 0) \leq \delta / \sqrt{2 \log(3^d \sqrt{J} cv_d)},$$

which is bounded by  $\delta$  when  $\sqrt{J} \geq e^{1/2} / 3^d cv_d$ . Further, setting  $\eta = 2(\sqrt{6 \log(J)} / \sqrt{J} cv_d)^{1/d}$  gives

$$\Pr \left( \min_{1 \leq k \leq N} \sum_{i=1}^J \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0 \right) \lesssim J^{-5/2},$$

so that

$$\begin{aligned}
& \sum_{J=1}^{\infty} \Pr \left( \min_{1 \leq j \leq J} \min_{1 \leq k \leq N} \sum_{i=1}^j \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0 \right) \\
&\leq \sum_{J=1}^{\infty} J \Pr \left( \min_{1 \leq k \leq N} \sum_{i=1}^J \mathbb{1}\{\|\boldsymbol{\lambda}_i - \tilde{\boldsymbol{\lambda}}_k\| \leq \eta/2\} \leq 0 \right) \\
&\lesssim \sum_{J=1}^{\infty} \frac{1}{J^{3/2}} < \infty.
\end{aligned}$$

Now the Borel-Cantelli lemma (e.g., Kallenberg, 1997, Theorem 4.18) implies that, with probability 1,

$$|B(\boldsymbol{\lambda}, \eta) \cap \Lambda| \geq 1,$$

for all  $J$  sufficiently large. □



## E Selected Validation Schemes

### E.1 Definition of Index Sets

Recall:

- (i) (holdout) Let  $M = 1$  and  $\mathcal{I}_{1,j} = \mathcal{I}_1$  for all  $j = 1, \dots, J$ , and some size- $\lceil \alpha n \rceil$  index set  $\mathcal{I}_1$ .
- (ii) (reshuffled holdout) Let  $M = 1$  and  $\mathcal{I}_{1,1}, \dots, \mathcal{I}_{1,J}$  be independently drawn from the uniform distribution over all size- $\lceil \alpha n \rceil$  subsets from  $\{1, \dots, n\}$ .
- (iii) ( $M$ -fold CV) Let  $\alpha = 1/M$  and  $\mathcal{I}_1, \dots, \mathcal{I}_M$  be a disjoint partition of  $\{1, \dots, n\}$ , and  $\mathcal{I}_{m,j} = \mathcal{I}_m$  for all  $j = 1, \dots, J$ .
- (iv) (reshuffled  $M$ -fold CV) Let  $\alpha = 1/M$  and  $(\mathcal{I}_{1,j}, \dots, \mathcal{I}_{M,j}), j = 1, \dots, J$ , be independently drawn from the uniform distribution over disjoint partitions of  $\{1, \dots, n\}$ .
- (v) ( $M$ -fold holdout) Let  $\mathcal{I}_m, m = 1, \dots, M$ , be independently drawn from the uniform distribution over size- $\lceil \alpha n \rceil$  subsets of  $\{1, \dots, n\}$  and set  $\mathcal{I}_{m,j} = \mathcal{I}_m$  for all  $m = 1, \dots, M, j = 1, \dots, J$ .
- (vi) (reshuffled  $M$ -fold holdout) Let  $\mathcal{I}_{m,j}, m = 1, \dots, M, j = 1, \dots, J$ , be independently drawn from the uniform distribution over size- $\lceil \alpha n \rceil$  subsets of  $\{1, \dots, n\}$ .

### E.2 Derivation of Reshuffling Parameters in Limiting Distribution

Recall

$$\tau_{i,j,M} = \frac{1}{nM^2\alpha^2} \sum_{s=1}^n \sum_{m=1}^M \sum_{m'=1}^M \Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}).$$

For all schemes in the proposition, the probabilities are independent of the index  $s$ , so the average over  $s = 1, \dots, n$  can be omitted. We now verify the constants  $\sigma, \tau$  from Table 1.

- (i) It holds

$$\Pr(s \in \mathcal{I}_{1,i} \cap \mathcal{I}_{1,j}) = \Pr(s \in \mathcal{I}_1) = \alpha.$$

Hence,

$$\tau_{i,j,1} = 1/\alpha = 1/\alpha \times 1 = \sigma^2 \times \tau^2.$$

- (ii) (reshuffled holdout) This is a special case of part (vi) with  $M = 1$ .
- (iii) ( $M$ -fold CV) It holds

$$\Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}) = \Pr(s \in \mathcal{I}_m \cap \mathcal{I}_{m'}) = \begin{cases} 1/M, & m = m', \\ 0, & m \neq m'. \end{cases}$$

Only  $M$  probabilities in the double sum are non-zero, whence

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times M/M = 1/\alpha^2 M^2 = 1 \times 1 = \sigma^2 \times \tau^2,$$

where we used  $\alpha = 1/M$ .

- (iv) (reshuffled  $M$ -fold CV) It holds

$$\Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}) = \begin{cases} 1/M, & m = m', i = j \\ 0, & m \neq m', i = j \\ 1/M^2, & m = m', i \neq j \\ 1/M^2, & m \neq m', i \neq j. \end{cases}$$

For  $i = j$ , only  $M$  probabilities in the double sum are non-zero. Also using  $\alpha = 1/M$ , we get

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times M \times 1/M = 1 = \sigma^2.$$

For  $i \neq j$ ,

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times M^2 \times 1/M^2 = 1 \times 1 = \sigma^2 \times \tau^2.$$

(v) ( $M$ -fold holdout) It holds

$$\Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}) = \Pr(s \in \mathcal{I}_m \cap \mathcal{I}_{m'}) = \begin{cases} \alpha, & m = m', \\ \alpha^2, & \text{else.} \end{cases}$$

This gives

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times [M \times \alpha + (M-1)M \times \alpha^2] = [1/\alpha M + (M-1)/M] \times 1 = \sigma^2 \times \tau^2.$$

for all  $i, j$ .

(vi) (reshuffled  $M$ -fold holdout) It holds

$$\Pr(s \in \mathcal{I}_{m,i} \cap \mathcal{I}_{m',j}) = \begin{cases} \alpha, & m = m', i = j \\ \alpha^2, & \text{else.} \end{cases}$$

For  $i = j$ , this gives

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times [M \times \alpha + (M-1)M \times \alpha^2] = 1/\alpha M + (M-1)/M.$$

For  $i \neq j$ ,

$$\tau_{i,j,M} = \frac{1}{M^2\alpha^2} \times (M^2 \times \alpha^2) = 1.$$

This implies that (1) holds with  $\sigma^2 = 1/M\alpha + (M-1)/M$ ,  $\tau^2 = 1/(1/M\alpha + (M-1)/M)$ .

**Remark E.1.** Although not technically covered by Theorem 2.1, performing independent bootstraps for each  $\lambda_j$  correspond to reshuffled  $n$ -fold holdout with  $\alpha = 1/n$ . Accordingly,  $\sigma \approx \sqrt{2}$  and  $\tau \approx \sqrt{1/2}$ .

## F Details Regarding Benchmark Experiments

### F.1 Datasets

We list all datasets used in the benchmark experiments in Table 4.

Table 4: List of datasets used in benchmark experiments. All information can be found on [OpenML](#) (Vanschoren et al., 2014).

OpenML Dataset ID	Dataset Name	Size ( $n \times p$ )
23517	numerai28.6	96320 $\times$ 21
1169	airlines	539383 $\times$ 7
41147	albert	425240 $\times$ 78
4135	Amazon_employee_access	32769 $\times$ 9
1461	bank-marketing	45211 $\times$ 16
1590	adult	48842 $\times$ 14
41150	MiniBooNE	130064 $\times$ 50
41162	kick	72983 $\times$ 32
42733	Click_prediction_small	39948 $\times$ 11
42742	porto-seguro	595212 $\times$ 57

Note that datasets serve as data generating processes (DGPs; Hothorn et al., 2005). As we are mostly concerned with the actual generalization performance of the final best HPC found during HPO based on validation performance we rely on a comparably large held out test set that is not used during HPO. We therefore use 5000 data points sampled from a DGP as an outer test set. To further be able to measure the generalization performance robustly for varying data sizes available during HPO, we construct concrete tasks based on the DGPs by sampling subsets of (`train_valid`;  $n$ ) size 500, 1000 and 5000 from the DGPs. This results in 30 tasks in total (10 DGPs  $\times$  3 `train_valid` sizes). For more details and the concrete implementation of this procedure, see Appendix F.3. We also collected another 5000 data points as an external validation set, but did not use it. Therefore, we had to tighten the restriction to 10000 data points mentioned in the main paper to 15000 data points as the lower bound on data points. To allow for stronger variation over different replications, we decided to use 20000 as the final lower bound.

## F.2 Learning Algorithms

Here we briefly present training pipeline details and search spaces of the learning algorithms used in our benchmark experiments.

The funnel-shaped MLP is based on sklearn’s MLP Classifier and is constructed in the following way: The hidden layer size for each layer is determined by `num_layers` and `max_units`. We start with `max_units` and half the number of units for every subsequent layer to create a funnel. `max_batch_size` is the largest power of 2 that is smaller than the number of training samples available. We use ReLU as activation function and train the network optimizing logloss as a loss function via SGD using a constant learning rate and Nesterov momentum for 100 epochs. Table 5 lists the search space (inspired from Zimmer et al. (2021)) used during HPO.

The Elastic Net is based on sklearn’s Logistic Regression Classifier. We train it for a maximum of 1000 iterations using the "saga" solver. Table 6 lists the search space used during HPO.

The XGBoost and CatBoost search spaces are listed in Table 7 and Table 8, both inspired from their search spaces used in McElfresh et al. (2023).

For both the Elastic Net and Funnel MLP, missing values are imputed in the preprocessing pipeline (mean imputation for numerical features and adding a new level for categorical features). Categorical features are target encoded in a cross-validated manner using a 5-fold CV. Features are then scaled to zero mean and unit variance via a standard scaler. For XGBoost, we impute missing values for categorical features (adding a new level) and target encode them in a cross-validated manner using a 5-fold CV. For CatBoost, no preprocessing is performed.

XGBoost and CatBoost models are trained for 2000 iterations and stop early if the validation loss (using the default internal loss function used during training, i.e., logloss) does not improve over a horizon of 20 iterations. For retraining the best configuration on the whole train and validation data, the number of boosting iterations is set to the number of iterations used to find the best validation performance prior to the stopping mechanism taking action.<sup>7</sup>

## F.3 Exact Implementation

In the following, we outline the exact implementation of performing one HPO run for a given learning algorithm on a concrete task (dataset  $\times$  `train_valid` size) and a given resampling. We release all code to replicate benchmark results and reproduce our analyses via [https://github.com/slds-lmu/paper\\_2024\\_resuffling](https://github.com/slds-lmu/paper_2024_resuffling). For a given replication (in total 10):

1. We sample (without replacement) `train_valid` size (500, 1000 or 5000 points) and `test` size (always 5000) points from the DGP (i.e. a concrete dataset in Table 4). These are shared for every learning algorithm (i.e. all learning algorithms are evaluated on the same data).
2. A given HPC is evaluated in the following way:
  - The resampling operates on the train validation<sup>8</sup> set of size `train_valid`.
  - The learning algorithm is configured by the HPC.
  - The learning algorithm is trained on training splits and evaluated on validation splits according to the resampling strategy. In case reshuffling is turned on, the training and validation splits are recreated for every HPO. We compute the Accuracy, ROC AUC and logloss when using a random search and compute ROC AUC when using HEBO or SMAC3 and average performance over all folds for resamplings involving multiple folds.
  - For each HPC we then always re-train the model on all `train_valid` data being available and evaluate the model on the held-out `test` set to compute an outer estimate of generalization performance for each HPC (regardless of whether it is the incumbent for a given iteration or not).

<sup>7</sup>For CV and repeated holdout we take the average number of boosting iterations over the models trained on the different folds.

<sup>8</sup>With train validation we refer to all data being available during HPO which is then further split by a resampling into train and validation sets.

Table 5: Search Space for Funnel-Shaped MLP Classifier.

Parameter	Type	Range	Log
num_layers	Int.	1 to 5	No
max_units	Int.	64, 128, 256, 512	No
learning_rate	Num.	$1 \times 10^{-4}$ to $1 \times 10^{-1}$	Yes
batch_size	Int.	16, 32, ..., max_batch_size	No
momentum	Num.	0.1 to 0.99	No
alpha	Num.	$1 \times 10^{-6}$ to $1 \times 10^{-1}$	Yes

Table 6: Search Space for Elastic Net Classifier.

Parameter	Type	Range	Log
C	Num.	$1 \times 10^{-6}$ to $1 \times 10^4$	Yes
l1_ratio	Num.	0.0 to 1.0	No

Table 7: Search Space for XGBoost Classifier.

Parameter	Type	Range	Log
max_depth	Int.	2 to 12	Yes
alpha	Num.	$1 \times 10^{-8}$ to 1.0	Yes
lambda	Num.	$1 \times 10^{-8}$ to 1.0	Yes
eta	Num.	0.01 to 0.3	Yes

Table 8: Search Space for CatBoost Classifier.

Parameter	Type	Range	Log
learning_rate	Num.	0.01 to 0.3	Yes
depth	Int.	2 to 12	Yes
l2_leaf_reg	Num.	0.5 to 30	Yes

3. We evaluate 500 HPCs when using random search and 250 HPC when using HEBO or SMAC3 (SMAC4HPO facade).

As resamplings, we use holdout with a 80/20 train-validation split and 5 folds for CV, so that the holdout strategy is just one fold of the CV and the fraction of data points being used for training and respectively validation are the same across different resampling strategies. 5-fold holdout simply repeats the holdout procedure five times and 5x 5-fold CV repeats the 5-fold CV five times. Each of the four resamplings can be reshuffled or not (standard).

As mentioned above, the test set is only varied for each of the 10 replica (repetitions with different seeds), but consistent for different tasks (i.e. the different learning algorithms are evaluated on the same test set, similarly, also the different dataset subsets all share the same test set). This allows for fair comparisons of different resamplings on a concrete problem (i.e. a given dataset, train\_valid size and learning algorithm). Additionally, for the random search, the 500 HPCs evaluated for a given learning algorithm are also fixed over different dataset and train\_valid size combinations. This is done to allow for an isolation of the effect, the concrete resampling (and whether it is reshuffled or not) has on generalization performance, reducing noise arising due to different HPCs. Learning algorithms themselves are not explicitly seeded to allow for variation during model training over different replications. Resamplings and partitioning of data are always performed in a stratified manner with respect to the target variable.

For the random search, we only ran (standard and reshuffled) holdout and (standard and reshuffled) 5x 5-fold CV experiments (because we can simulate 5-fold CV and 5-fold holdout experiments based

on the results obtained from the 5x 5-fold CV (by only considering the first repeat or the first fold for each of the five repeats).<sup>9</sup>

For running HEBO or SMAC3, each resampling (standard and reshuffled for holdout, 5-fold holdout, 5-fold CV, 5x 5-fold CV) has to be actually run due to the adaptive nature of BO.

For the random search experiments, this results in  $10$  (DGPs)  $\times 3$  (`train_valid` sizes)  $\times 4$  (learning algorithms)  $\times 2$  (holdout or 5x 5-fold CV)  $\times 2$  (standard or reshuffled)  $\times 10$  (replications) = 4800 HPO runs,<sup>10</sup> each involving the evaluation of 500 HPCs and each evaluation of an HPC involving either 2 (for holdout; due to retraining on train validation data) or 26 (for 5x 5-fold CV; due to retraining on train validation data) model fits. In summary, the random search experiments involve the evaluation of 2.4 Million HPCs with in total 33.6 Million model fits.

Similarly, for the HEBO and SMAC3 experiments, this each results in  $10$  (DGPs)  $\times 3$  (`train_valid` sizes)  $\times 4$  (learning algorithms)  $\times 4$  (holdout, 5-fold CV, 5x 5-fold CV or 5-fold holdout)  $\times 2$  (standard or reshuffled)  $\times 10$  (replications) = 9600 HPO runs<sup>11</sup>, each involving the evaluation of 250 HPCs and each evaluation of an HPC involving either 2 (for holdout; due to retraining on train validation data), 6 (for 5-fold CV or 5-fold holdout; due to retraining on train validation data) or 26 (for 5x 5-fold CV; due to retraining on train validation data) model fits. In summary, the HEBO and SMAC3 experiments *each* involve the evaluation of 2.4 Million HPCs with in total 24 Million model fits.

## F.4 Compute Resources

We estimate our total compute time for the random search, HEBO and SMAC3 experiments to be roughly 11.86 CPU years. Benchmark experiments were run on an internal HPC cluster equipped with a mix of Intel Xeon E5-2670, Intel Xeon E5-2683 and Intel Xeon Gold 6330 instances. Jobs were scheduled to use a single CPU core and were allowed to use up to 16GB RAM. Total emissions are estimated to be an equivalent of roughly 6508.67 kg CO<sub>2</sub>.

# G Additional Benchmark Results Visualizations

## G.1 Main Experiments

In this section, we provide additional visualizations of the results of our benchmark experiments.

Figure 6 illustrates the trade-off between the final number of model fits required by different resamplings and the final average normalized test performance (AUC ROC) after running random search for a budget of 500 hyperparameter configurations. We can see that the reshuffled holdout on average comes close to the final test performance of the overall more expensive 5-fold CV.

Below, we give an overview of the different types of additional analyses and visualizations we provide. Normalized metrics, i.e., normalized validation or test performance refer to the measure being scaled to  $[0, 1]$  based on the empirical observed minimum and maximum values obtained on the raw results level (ADTM; see Wistuba et al., 2018). More concretely, for each scenario consisting of a learning algorithm that is run on a given task (dataset  $\times$  `train_valid` size) given a certain performance metric, the performance values (validation or test) for all resamplings and optimizers are normalized on the replication level to  $[0, 1]$  by subtracting the empirical best value and dividing by the range of performance values. Therefore a normalized performance value of 0 is best and 1 is worst. Note that we additionally provide further aggregated results on the learning algorithm level and raw results of validation and test performance via [https://github.com/slds-lmu/paper\\_2024\\_resuffling](https://github.com/slds-lmu/paper_2024_resuffling).

- Random search
  - Normalized validation performance in Figure 7.

---

<sup>9</sup>We even could have simulated the vanilla holdout from the 5x 5-fold CV experiments by choosing an arbitrary fold and repeat but choose not to do so, to have some sanity checks regarding our implementation by being able to compare the "true" holdout with a the simulated holdout.

<sup>10</sup>Note that we do not have to take the 3 different metrics into account because random search allows us to simulate runs for different metric post hoc.

<sup>11</sup>Note that HEBO and SMAC3 were only run for ROC AUC as the performance metric.



Figure 6: Trade-off between the final number of model fits required by different resamplings and the final average normalized test performance (AUC ROC) after running random search for a budget of 500 hyperparameter configurations. Averaged over different tasks, learning algorithms and replications separately for increasing  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

- Normalized test performance in Figure 8.
- Improvement in test performance over 5-fold CV in Figure 9.
- Rank w.r.t. test performance in Figure 10.
- HEBO and SMAC3 vs. random search holdout
  - Normalized validation performance in Figure 11.
  - Normalized test performance in Figure 12.
  - Improvement in test performance over standard holdout in Figure 13.
  - Rank w.r.t. test performance in Figure 14.
- HEBO and SMAC3 vs. random search 5-fold holdout
  - Normalized validation performance in Figure 15.
  - Normalized test performance in Figure 16.
  - Improvement in test performance over standard 5-fold holdout in Figure 17.
  - Rank w.r.t. test performance in Figure 18.
- HEBO and SMAC3 vs. random search 5-fold CV
  - Normalized validation performance in Figure 19.
  - Normalized test performance in Figure 20.
  - Improvement in test performance over 5-fold CV in Figure 21.
  - Rank w.r.t. test performance in Figure 22.
- HEBO and SMAC3 vs. random search 5x 5-fold CV
  - Normalized validation performance in Figure 23.
  - Normalized test performance in Figure 24.
  - Improvement in test performance over 5x 5-fold CV in Figure 25.
  - Rank w.r.t. test performance in Figure 26.

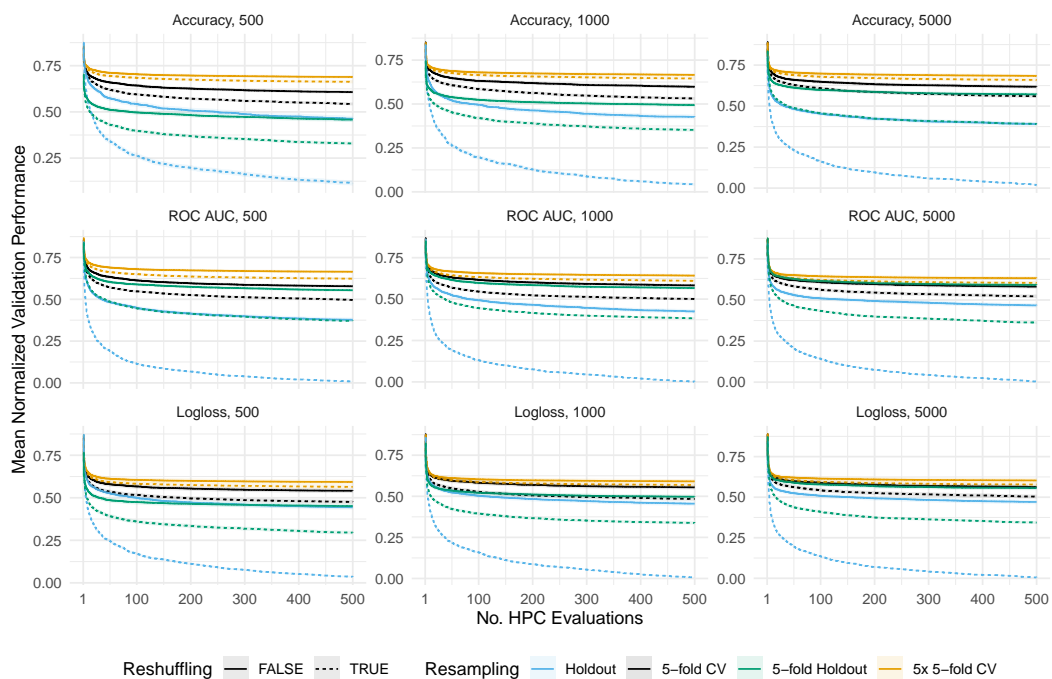


Figure 7: Random search. Average normalized performance over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

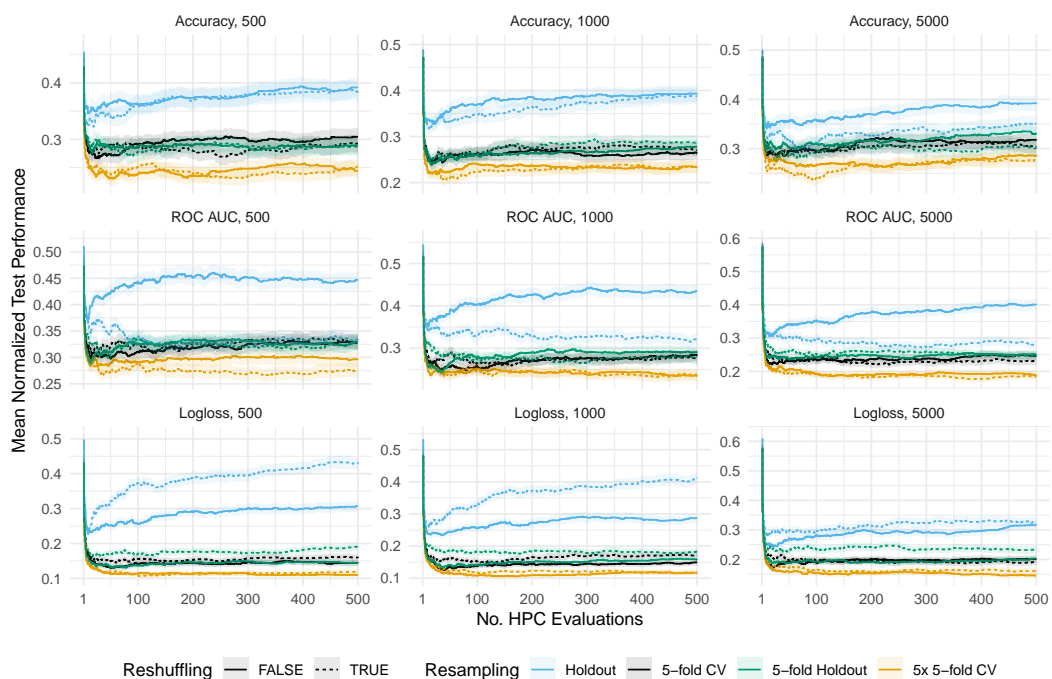


Figure 8: Random search. Average normalized test performance over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

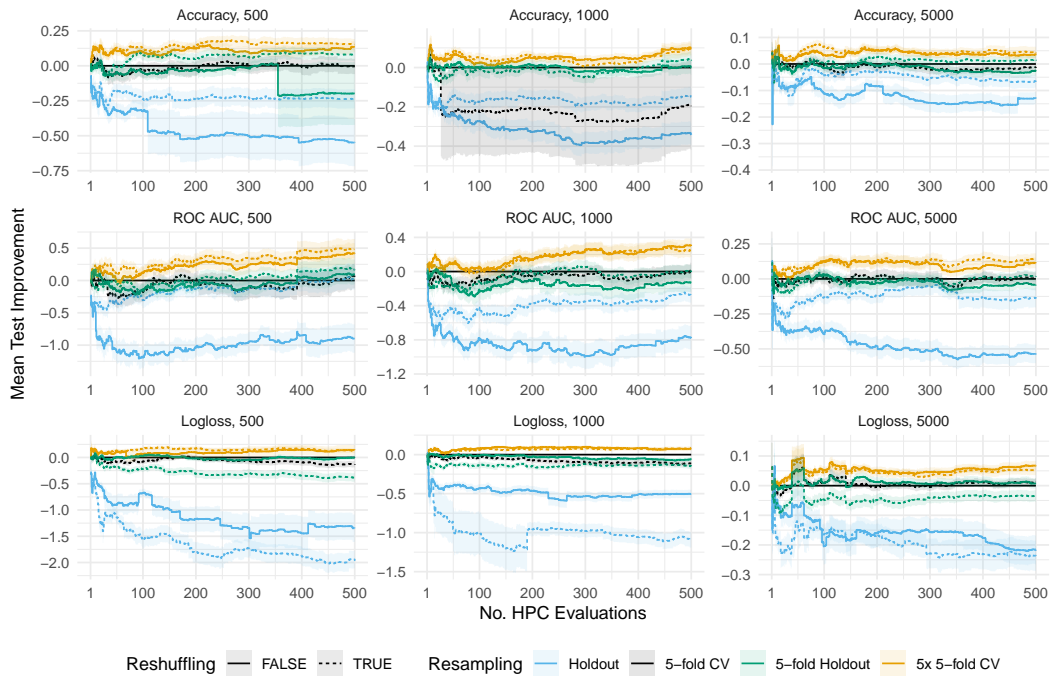


Figure 9: Random search. Average improvement (compared to standard 5-fold CV) with respect to test performance of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

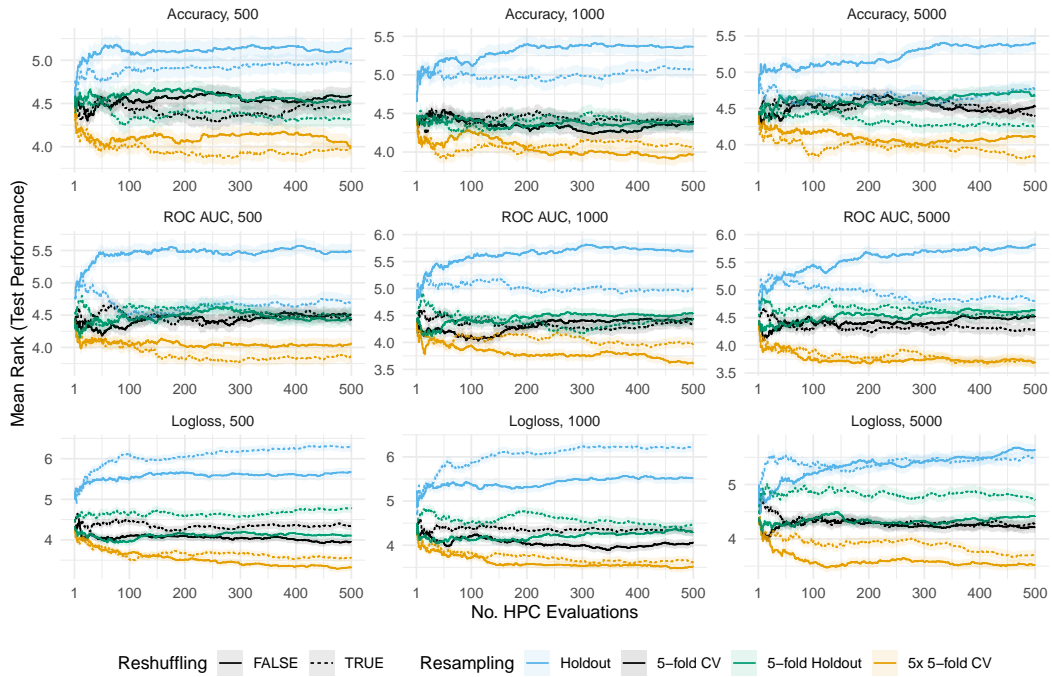


Figure 10: Random search. Average ranks (lower is better) with respect to test performance over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.






Figure 11: HEBO and SMAC3 vs. random search for holdout. Average normalized validation performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 12: HEBO and SMAC3 vs. random search for holdout. Average normalized test performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 13: HEBO and SMAC3 vs. random search for holdout. Average improvement (compared to standard holdout) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 14: HEBO and SMAC3 vs. random search for holdout. Average ranks (lower is better) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 15: HEBO and SMAC3 vs. random search for 5-fold holdout. Average normalized validation performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 16: HEBO and SMAC3 vs. random search for 5-fold holdout. Average normalized test performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 17: HEBO and SMAC3 vs. random search for 5-fold holdout. Average improvement (compared to standard 5-fold holdout) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 18: HEBO and SMAC3 vs. random search for 5-fold holdout. Average ranks (lower is better) with respect to test performance (ROC AUC) of the incumbent tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 19: HEBO and SMAC3 vs. random search for 5-fold CV. Average normalized validation performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 20: HEBO and SMAC3 vs. random search for 5-fold CV. Average normalized test performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 21: HEBO and SMAC3 vs. random search for 5-fold CV. Average improvement (compared to standard 5-fold CV) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 22: HEBO and SMAC3 vs. random search for 5-fold CV. Average ranks (lower is better) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 23: HEBO and SMAC3 vs. random search for 5x 5-fold CV. Average normalized validation performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 24: HEBO and SMAC3 vs. random search for 5x 5-fold CV. Average normalized test performance (ROC AUC) over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 25: HEBO and SMAC3 vs. random search for 5x 5-fold CV. Average improvement (compared to standard 5x 5-fold CV) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.




Figure 26: HEBO and SMAC3 vs. random search for 5x 5-fold CV. Average ranks (lower is better) with respect to test performance (ROC AUC) of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

## G.2 Ablation on M-fold holdout

Based on the 5x 5-fold CV results we further simulated different  $M$ -fold holdout resamplings (standard and reshuffled) by taking  $M$  repeats from the first fold of the 5x 5-fold CV. This allows us to get an understanding of the effect more folds have on  $M$ -fold holdout, especially in the context of reshuffling.

Regarding normalized validation performance we observe that more folds generally result in a less optimistically biased validation performance (see Figure 27). Looking at normalized test performance (Figure 28) we observe the general trend that more folds result in better test performance – which is expected. Reshuffling generally results in better test performance compared to the standard resampling (with the exception of logloss where especially in the case of a single holdout, reshuffling can hurt generalization performance). This effect is smaller, the more folds are used, which is in line with our theoretical results presented in Table 1. Looking at improvement compared to standard 5-fold holdout with respect to test performance and ranks with respect to test performance, we observe that often reshuffled 2-fold holdout results that are highly competitive with standard 3, 4 or 5-fold holdout.

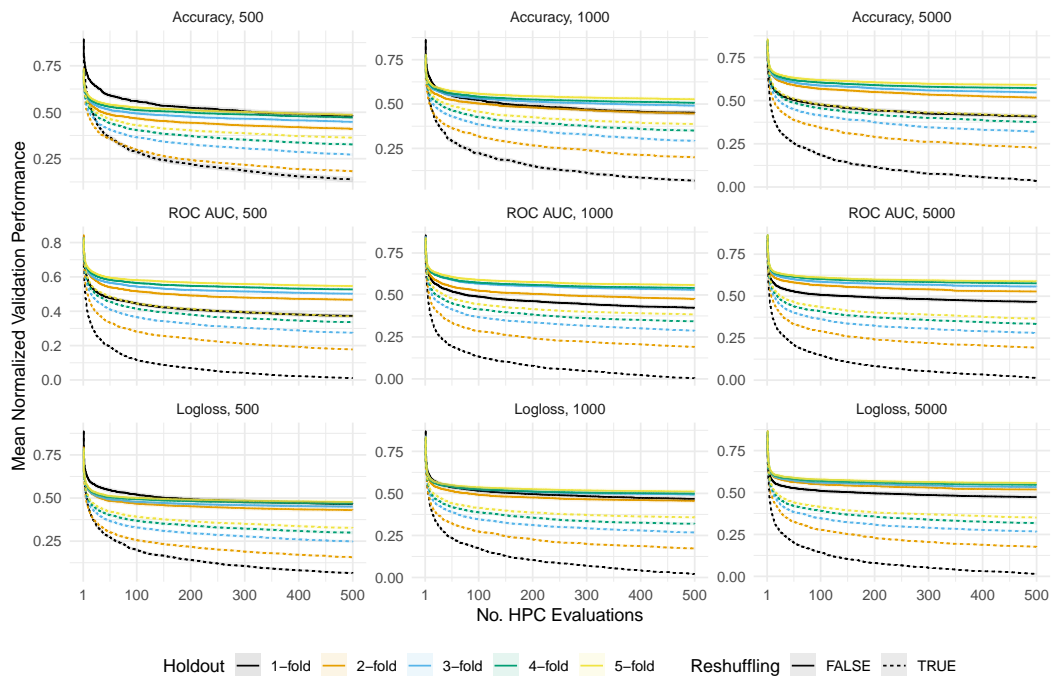


Figure 27: Random search. Average normalized validation performance over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

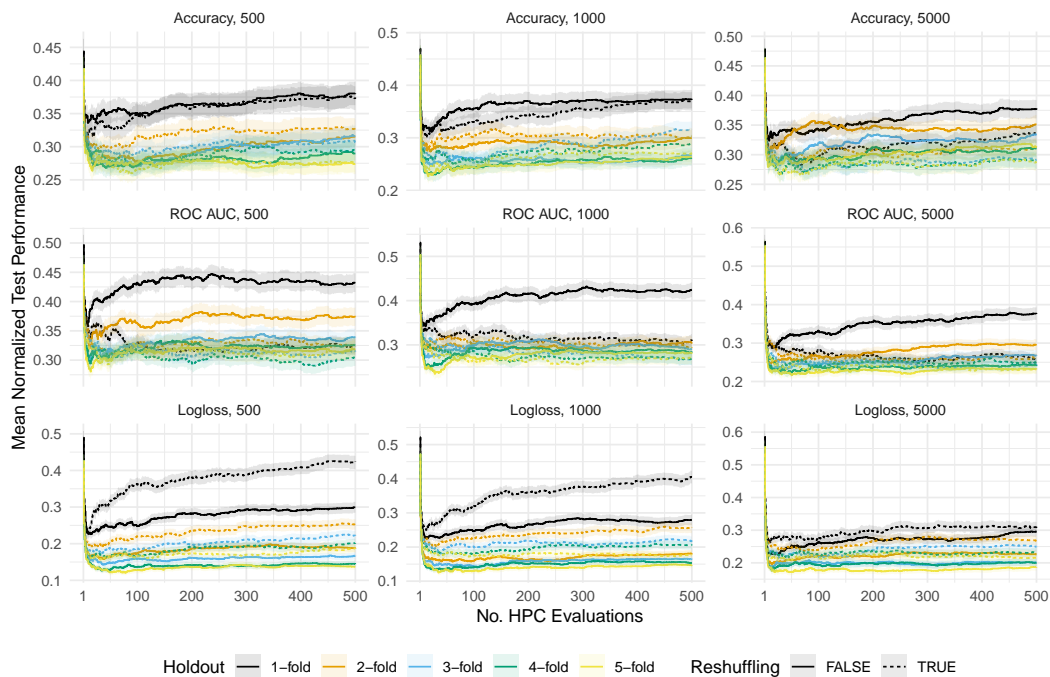


Figure 28: Random search. Average normalized test performance over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

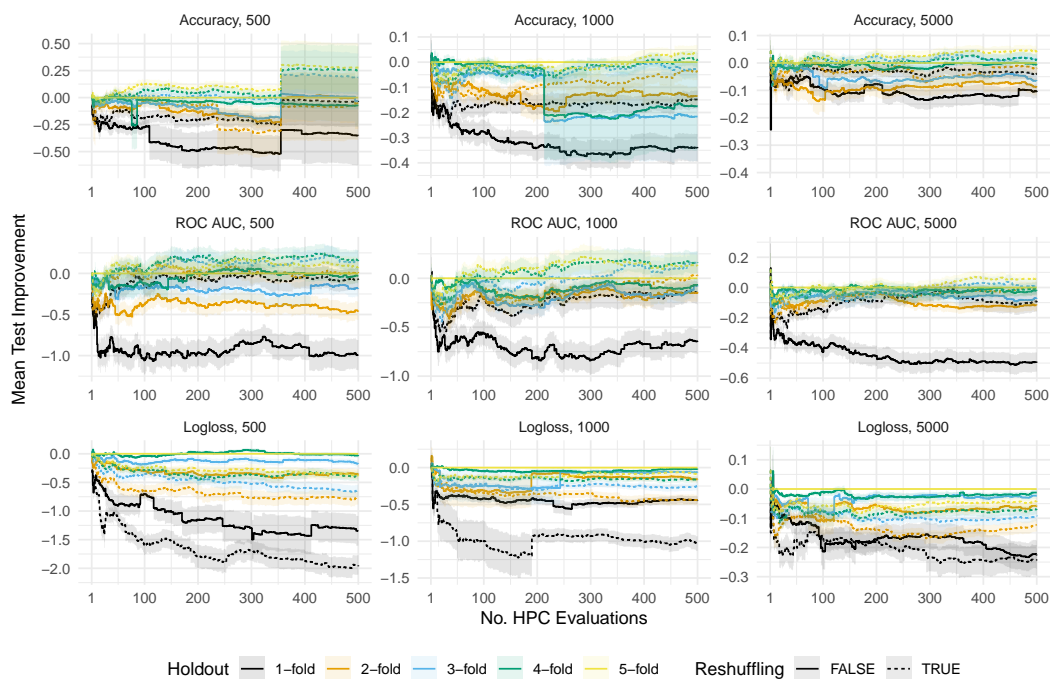


Figure 29: Random search. Average improvement (compared to standard 5-fold holdout) with respect to test performance of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.



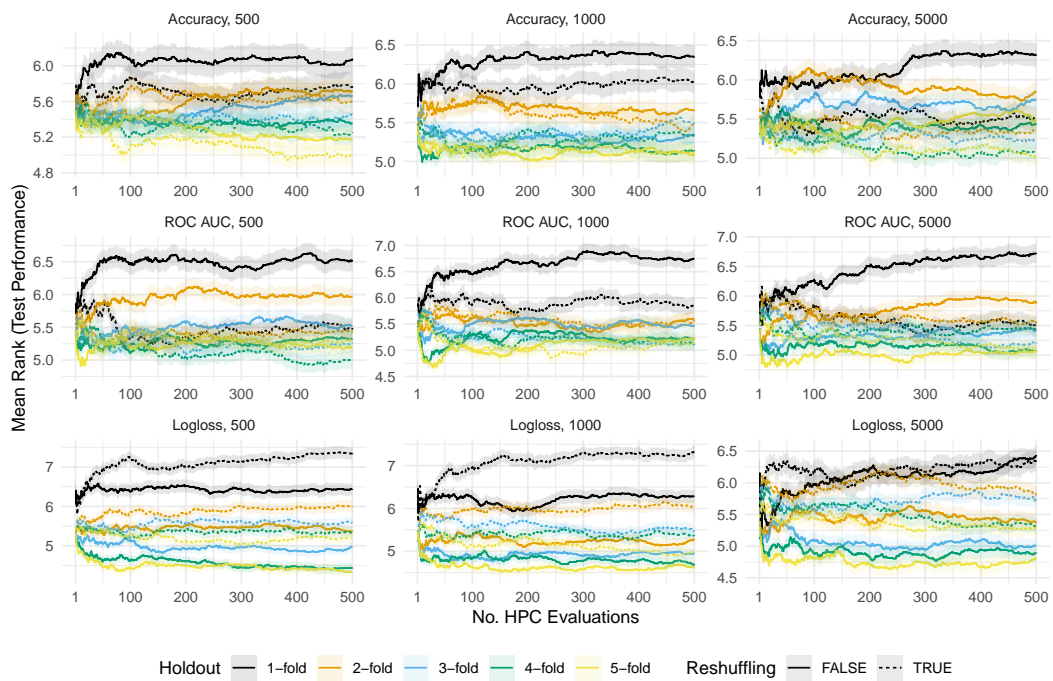


Figure 30: Random search. Average ranks (lower is better) with respect to test performance of the incumbent over tasks, learners and replications for different  $n$  (train-validation sizes, columns). Shaded areas represent standard errors.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We outline our three main contributions in the introduction (Section 1). We do not discuss generalization in the introduction, but rather in the discussion in Section 5.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The paper provides an analysis of reshuffling data in the context of estimating the generalization error for hyperparameter optimization. Our theoretical analysis explains why reshuffling works, and we experimentally verify the theoretical analysis. We discuss the limitations of our work in Section 5.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]

Justification: Full assumptions and proofs for our main results (Theorem 2.1 and Theorem 2.2) are given in Appendix C.1 and Appendix C.2, respectively. Derivations for the parameters in Table 1 are provided in Appendix E. The additional results for the grid density are stated and proven directly in Appendix D.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We provide thorough details on the experimental setup in Section 4.1 and Appendix F. Moreover, we provide code to reproduce our results under an open source license at [https://github.com/slds-lmu/paper\\_2024\\_resuffling](https://github.com/slds-lmu/paper_2024_resuffling).

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Regarding datasets, we rely on [OpenML.org](https://openml.org). We provide thorough details on the experimental setup in Section 4.1 and Appendix F. Moreover, we provide code to reproduce our results under an open source license at [https://github.com/slds-lmu/paper\\_2024\\_resuffling](https://github.com/slds-lmu/paper_2024_resuffling).

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide thorough details on the experimental setup in Section 4.1 and Appendix F. Moreover, we provide code to reproduce our results under an open source license at [https://github.com/slds-lmu/paper\\_2024\\_resuffling](https://github.com/slds-lmu/paper_2024_resuffling).

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We report the standard error in every analysis.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: We provide details in Appendix F.4.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: Our work provides a study on reshuffling data when estimating the generalization error in hyperparameter tuning. Therefore, our work is applicable wherever standard machine learning is applicable, and we do not see any ethical concerns in our method.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: The paper conducts fundamental research that is not tied to particular applications, let alone deployment.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The paper conducts fundamental research that is not tied to particular applications, let alone deployment. The paper does not develop models that have a high risk for misuse.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We used datasets from [OpenML.org](https://openml.org) and reference the dataset pages. Further information of the datasets, including their licenses, are available at [OpenML.org](https://openml.org).

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We provide code as a new asset and describe how we make our code available in Point 5 of the NeurIPS Paper Checklist.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does neither involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does neither involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.