



HAL
open science

Better by Default: Strong Pre-Tuned MLPs and Boosted Trees on Tabular Data

David Holzmüller, Léo Grinsztajn, Ingo Steinwart

► **To cite this version:**

David Holzmüller, Léo Grinsztajn, Ingo Steinwart. Better by Default: Strong Pre-Tuned MLPs and Boosted Trees on Tabular Data. 2024. hal-04641923v1

HAL Id: hal-04641923

<https://hal.science/hal-04641923v1>

Preprint submitted on 9 Jul 2024 (v1), last revised 5 Nov 2024 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Better by Default: Strong Pre-Tuned MLPs and Boosted Trees on Tabular Data

David Holzmüller*
SIERRA Team, Inria Paris
Ecole Normale Supérieure
PSL University

Léo Grinsztajn
SODA Team, Inria Saclay

Ingo Steinwart
University of Stuttgart
Faculty of Mathematics and Physics
Institute for Stochastics and Applications

Abstract

For classification and regression on tabular data, the dominance of gradient-boosted decision trees (GBDTs) has recently been challenged by often much slower deep learning methods with extensive hyperparameter tuning. We address this discrepancy by introducing (a) RealMLP, an improved multilayer perceptron (MLP), and (b) improved default parameters for GBDTs and RealMLP. We tune RealMLP and the default parameters on a meta-train benchmark with 71 classification and 47 regression datasets and compare them to hyperparameter-optimized versions on a disjoint meta-test benchmark with 48 classification and 42 regression datasets, as well as the GBDT-friendly benchmark by Grinsztajn et al. (2022). Our benchmark results show that RealMLP offers a better time-accuracy tradeoff than other neural nets and is competitive with GBDTs. Moreover, a combination of RealMLP and GBDTs with improved default parameters can achieve excellent results on medium-sized tabular datasets (1K–500K samples) without hyperparameter tuning.

1 Introduction

Perhaps the most common type of data in practical machine learning (ML) is tabular data, characterized by a fixed number of numerical or categorical features (columns), lacking the spatiotemporal structure of most other data types such as image or text data. The moderate dimensionality and lack of symmetries make tabular data accessible to a wide variety of machine learning methods. While tabular data is very diverse and no method is dominant on all datasets, gradient-boosted decision trees (GBDTs) have demonstrated excellent results on benchmarks [57, 19, 43], although their superiority has been challenged by a variety of deep learning methods [3].

GBDTs and neural networks (NNs) are often compared using extensive hyperparameter optimization. This can be especially expensive for NNs, as multilayer perceptrons (MLPs) and Transformer-based models are roughly one and two orders of magnitude slower than typical GBDTs, respectively [19, 43]. To address this issue, we investigate the potential of improved MLPs and better dataset-independent default parameters for MLPs and GBDTs. Specifically, we compare the library defaults (D) to our tuned defaults (TD) and (dataset-dependent) hyperparameter optimization (HPO).

Besides offering convenient baselines for quick exploration and benchmarking, good default parameters also play an important role in automated ML (AutoML). AutoGluon [11] demonstrated that

*Work done partially while still at University of Stuttgart.

stacking and ensembling models with fixed parameters outperforms other AutoML approaches based on hyperparameter optimization [14]. Without stacking and ensembling, McElfresh et al. [43] have argued that light hyperparameter tuning for GBDTs is often more effective than trying out NNs. Here, we argue that even without stacking and ensembling, when using well-chosen default parameters for GBDTs and our improved MLPs, *trying GBDTs and MLPs is often faster and more beneficial than (naively) optimizing the hyperparameters of a single method.*

1.1 Contribution

The problem of finding better default parameters can be seen as a meta-learning problem [63]. We employ a meta-train benchmark consisting of 118 datasets on which the default hyperparameters are optimized, and a disjoint meta-test benchmark consisting of 90 datasets on which they are evaluated. We consider separate default parameters for classification, optimized for classification error, and for regression, optimized for RMSE. Our benchmarks do not contain missing numerical values, and we restrict ourselves to sizes between 1K and 500K samples, cf. Section 2.

In Section 3, we introduce **RealMLP**, which improves on standard MLPs through a **bag of tricks** and **better default parameters**, tuned entirely on the meta-train benchmark. We introduce many **novel or nonstandard components**, such as preprocessing using robust scaling and smooth clipping, a new numerical embedding variant, a diagonal weight layer, new schedules, different initialization methods, etc. Our benchmark results demonstrate that it outperforms other comparably fast NNs from the literature and can be competitive with GBDTs. We make RealMLP and the other models available through a **scikit-learn interface**.

In Section 4, we provide **new default parameters**, tuned on the meta-train benchmark, for XGBoost [10], LightGBM [31], and CatBoost [51]. While they cannot match HPO on average, they outperform the library defaults on the meta-test benchmark.

In Section 5, we evaluate these and other models on the meta-test **benchmark** and the benchmark by Grinsztajn et al. [19]. We also investigate several possibilities for algorithm selection and ensembling, demonstrating that algorithm selection over default methods provides a better time-performance tradeoff than HPO, thanks to our new improved default parameters and MLP.

Our code (including scikit-learn interfaces) for meta-train and meta-test benchmarks is available at

github.com/dholzmueLLer/pytabkit

At https://github.com/LeoGrin/tabular-benchmark/tree/better_by_default, we provide code for the adapted Grinsztajn et al. [19] benchmark. Our code and benchmark data will be archived at <https://doi.org/10.18419/darus-4255>.

1.2 Related Work

Neural networks Borisov et al. [3] review deep learning on tabular data and identify three main classes of methods: Data transformation methods, specialized architectures, and regularization models. In particular, recent research has mainly focused on specialized architectures based on attention [1, 27, 16, 8], including attention between datapoints [53, 59, 37, 55, 18]. However, these methods are usually significantly slower than MLPs or even GBDTs [19, 43, 18]. Our research instead expands on improvements to MLPs for tabular data such as the SELU activation function [35], bias initialization methods [60], regularization methods [30], categorical embedding layers [20], and numerical embedding layers [17].

Benchmarks Shwartz-Ziv and Armon [57] benchmarked three deep learning methods and noticed that they performed better on the datasets from their own papers than on other datasets. We address this issue by using more datasets and evaluating our methods on datasets that they were not tuned on. Grinsztajn et al. [19], McElfresh et al. [43], and Ye et al. [68] propose larger benchmarks and find that GBDTs still outperform deep learning methods on average, analyzing why and when this is the case. Kohli et al. [36] also emphasize the need for large benchmarks. We evaluate our methods on the benchmark by Grinsztajn et al. [19] as well as datasets from the AutoML benchmark [14] and the OpenML-CTR23 regression benchmark [13].

Better defaults Probst et al. [50] study the tunability of ML methods, i.e., the difference in benchmark scores between the best fixed hyperparameters and tuned hyperparameters. While their

Table 1: Characteristics of the meta-train and meta-test sets.

	$\mathcal{B}_{\text{class}}^{\text{train}}$	$\mathcal{B}_{\text{class}}^{\text{test}}$	$\mathcal{B}_{\text{reg}}^{\text{train}}$	$\mathcal{B}_{\text{reg}}^{\text{test}}$
#datasets	71	48	47	42
#dataset groups	46	48	26	42
min #samples	1847	1000	3338	1030
max #samples	45222	500000	48204	500000
max #classes	26	355	0	0
max #features	561	10000	520	4991
max #categories	41	7019	38	359

approach involves finding better defaults, they do not evaluate them on a separate meta-test benchmark, only consider classification, and do not provide defaults for LightGBM, CatBoost, and NNs.

Meta-learning The problem of finding the best fixed hyperparameters is a meta-learning problem [4, 63]. Although we do not introduce or employ a fully automated method to find good defaults, we use a meta-learning benchmark setup to properly evaluate them. Wistuba et al. [65] and Pfisterer et al. [49] learn portfolios of configurations and van Rijn et al. [62] learn symbolic defaults, but neither of the three papers considers GBDTs or NNs. Salinas and Erickson [54] learn large portfolios of configurations on an extensive benchmark, without studying the best defaults for individual model families. Such portfolios are successfully applied in modern AutoML methods [11, 12]. At the other end of the meta-learning spectrum, TabPFN [24] meta-learns a (tuning-free) learning method on small synthetic datasets. Unlike TabPFN, we only meta-learn hyperparameters and can therefore use fewer but larger and more realistic meta-train datasets, resulting in methods that scale to larger datasets.

2 Methodology

To evaluate a fixed hyperparameter configuration \mathcal{H} , we need a collection \mathcal{B} of benchmark datasets and a scoring function that computes a benchmark score $\mathcal{S}(\mathcal{B}, \mathcal{H})$ by aggregating the errors attained by the method with hyperparameters \mathcal{H} on each dataset. However, when optimizing \mathcal{H} on \mathcal{B} , we might overfit to the benchmark and therefore ideally need a second benchmark \mathcal{B}' to get an unbiased score for \mathcal{H} . We refer to $\mathcal{B}, \mathcal{B}'$ as meta-train and meta-test benchmarks and subdivide them into classification and regression benchmarks $\mathcal{B}_{\text{class}}^{\text{train}}, \mathcal{B}_{\text{reg}}^{\text{train}}, \mathcal{B}_{\text{class}}^{\text{test}},$ and $\mathcal{B}_{\text{reg}}^{\text{test}}$. Since the meta-train benchmark contains groups of datasets that are variants of the same dataset, for example by using different columns as targets, we use weighting factors inversely proportional to the group size.

Table 1 shows some characteristics of the meta-train and meta-test benchmarks. The meta-test benchmark includes datasets that are more extreme in several dimensions, allowing us to test whether our default parameters generalize “out of distribution”. For all datasets, we remove rows with missing numerical values and encode missing categorical values as a separate category. For regression, we standardize the targets on all datasets, such that 1 is the RMSE of the best constant predictor.

2.1 Benchmark Data Selection

The meta-train set consists of medium-sized datasets from the UCI Repository [32], adapted from Steinwart [60]. The meta-test set consists of datasets from the AutoML Benchmark [14] and additional regression datasets from the OpenML-CTR23 benchmark [13]. We subsample some large datasets and remove datasets that are already contained in the meta-train set, are too small, or have categories with too large cardinality. More details on the datasets and preprocessing can be found in Appendix C.3.

2.2 Aggregate Benchmark Score

To optimize the default parameters, we need to define a single benchmark score. To this end, we evaluate a method on $N_{\text{splits}} = 10$ random training-validation-test splits (60%-20%-20%) on each dataset. As metrics on individual dataset splits, we use classification error (100% – accuracy) for

classification and RMSE for regression. There are various options to aggregate these errors into a single score. Some, such as average rank or mean normalized error, depend on which other methods are included in the evaluation, hindering an independent optimization. We would like to use the geometric mean error because arguably, an error reduction from 0.02 to 0.01 is more valuable than an error reduction from 0.42 to 0.41. However, since the geometric mean error is too sensitive to cases with zero error (especially for classification error), we instead use a *shifted geometric mean error*, where a small value $\varepsilon := 0.01$ is added to the errors err_{ij} before taking the geometric mean:

$$\text{SGM}_\varepsilon := \exp \left(\sum_{i=1}^{N_{\text{datasets}}} \frac{w_i}{N_{\text{splits}}} \sum_{j=1}^{N_{\text{splits}}} \log(\text{err}_{ij} + \varepsilon) \right).$$

Here, we use weights $w_i = 1/N_{\text{datasets}}$ on the meta-test set. On the meta-train set, we make the w_i dependent on the number of related datasets, cf. Appendix C.3.

3 Improved MLP

Here, we introduce RealMLP-TD, our improved MLP with tuned defaults, which we designed based on the meta-train benchmark. We also introduce a simplified version called RealMLP-TD-S.

Data preprocessing We first apply one-hot encoding to categorical columns with at most eight distinct values (not counting missing values). Binary categories are encoded to a single feature with values $\{-1, 1\}$. Missing values in categorical columns are encoded to zero. We then preprocess all numerical columns, including the one-hot encoded ones, independently as follows: Let $x_1, \dots, x_n \in \mathbb{R}$ be the values in column i , and let q_p be the p -quantile of (x_1, \dots, x_n) for $p \in [0, 1]$. Then,

$$x_{j,\text{processed}} := f(s_j \cdot (x_j - q_{1/2})), \quad f(x) := \frac{x}{\sqrt{1 + (\frac{x}{3})^2}},$$

$$s_j := \begin{cases} \frac{1}{q_{3/4} - q_{1/4}} & , \text{ if } q_{3/4} \neq q_{1/4} \\ \frac{2}{q_1 - q_0} & , \text{ if } q_{3/4} = q_{1/4} \text{ and } q_1 \neq q_0 \\ 0 & , \text{ otherwise.} \end{cases}$$

In scikit-learn [48], this corresponds to applying a `RobustScaler` (first case) or `MinMaxScaler` (second case), and then the function f , which smoothly clips its input to the range $(-3, 3)$. Smooth clipping functions like f have been used by, e.g., Holzmüller et al. [25] and Hafner et al. [21]. Intuitively, when features have large outliers, smooth clipping prevents the outliers from affecting the result too strongly, while robust scaling prevents the outliers from affecting the inlier scaling.

NN architecture Our architecture, visualized in Figure 1 (a), is a multilayer perceptron (MLP) with three hidden layers containing 256 neurons each, except for the following additions and modifications:

- We use categorical embedding layers [20] to embed the remaining categorical features with cardinality > 8 .
- For numerical features, excluding the one-hot encoded ones, we introduce PBLD (periodic bias linear densenet) embeddings, which concatenate the original value to the PL embeddings proposed by Gorishniy et al. [17] and use a different periodic embedding with biases, inspired by Rahimi and Recht [52]. PBLD embeddings apply separate small two-layer MLPs to each feature x_i as

$$\left(x_i, \mathbf{W}_{\text{emb}}^{(2,i)} \cos(2\pi \mathbf{w}_{\text{emb}}^{(1,i)} x_i + \mathbf{b}_{\text{emb}}^{(1,i)}) + \mathbf{b}_{\text{emb}}^{(2,i)} \right) \in \mathbb{R}^4.$$

For efficiency reasons, we use 4-dimensional embeddings with $\mathbf{w}_{\text{emb}}^{(1,i)}, \mathbf{b}_{\text{emb}}^{(1,i)} \in \mathbb{R}^{16}, \mathbf{b}_{\text{emb}}^{(2,i)} \in \mathbb{R}^3, \mathbf{W}_{\text{emb}}^{(2,i)} \in \mathbb{R}^{3 \times 16}$.

- To encourage (soft) feature selection, we introduce a scaling layer before the first linear layer, which is simply a matrix-vector product with a diagonal weight matrix. In other words, it computes $x_{i,\text{out}} = s_i \cdot x_{i,\text{in}}$, with a learnable scaling factor s_i for each feature i . We found it beneficial to use a larger learning rate for this layer.
- Our linear layers use the neural tangent parametrization (NTP) as proposed by Jacot et al. [28], i.e., they compute $\mathbf{z}^{(l+1)} = d_l^{-1/2} \mathbf{W}^{(l)} \mathbf{x}^{(l)} + \mathbf{b}^{(l)}$, where d_l is the dimension of the

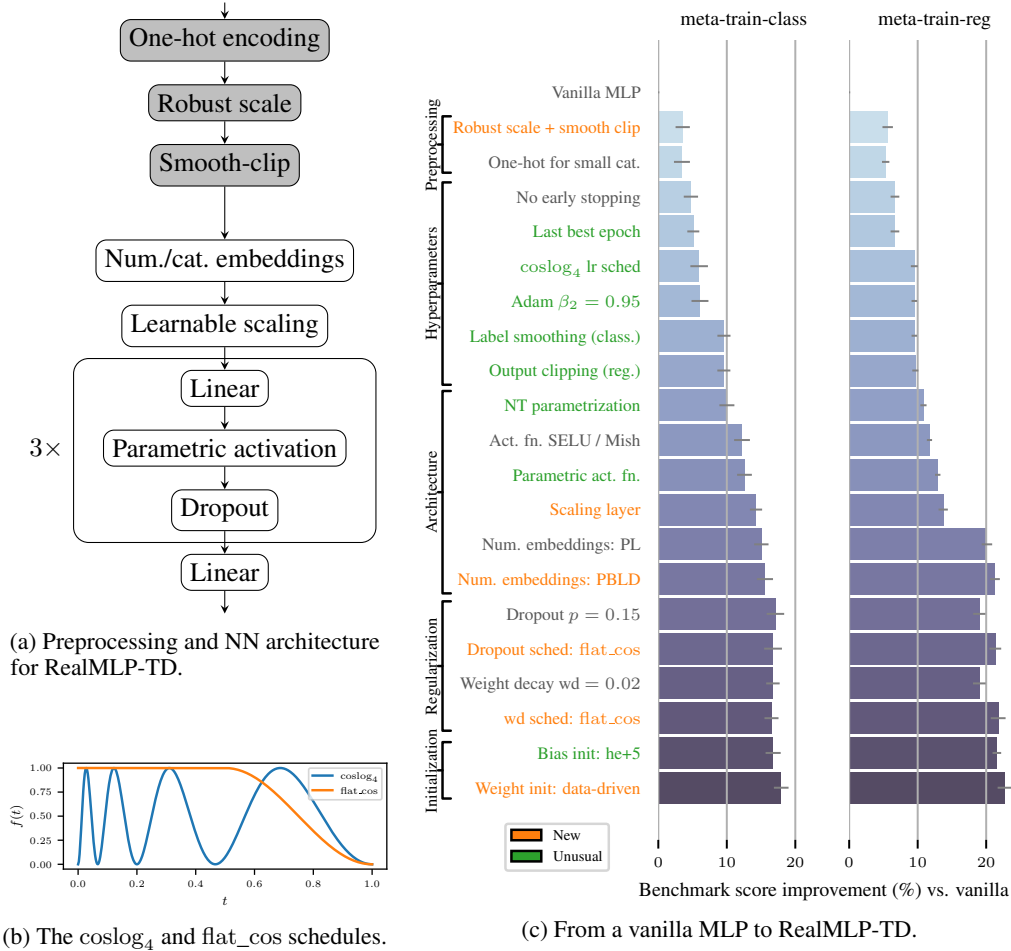


Figure 1: **Components of RealMLP-TD.** In (c), we add one component in each step, tuning the best default learning rate for each step separately. The vanilla MLP uses categorical embeddings, a quantile transform to preprocess numerical features, default PyTorch initialization, ReLU activation, early stopping, and is optimized with Adam with default parameters. For more details, see Appendix A.3. The error bars are approximate 95% confidence intervals for the limit $\#\text{splits} \rightarrow \infty$, see Appendix C.5.

layer input $\mathbf{x}^{(l)}$. The motivation behind the use of the NTP here is that it effectively modifies the learning rate for the weight matrices depending on the input dimension d_l , hopefully preventing too large steps whenever the number of columns is large. We did not observe improvements when using the Adam version of the maximal update parametrization [67].

- We use parametric activation functions inspired by PReLU [22]. In general, for an activation function σ , we define a parametric version with separate learnable α_i for each neuron i :

$$\sigma_{\alpha_i}(x_i) = (1 - \alpha_i)x_i + \alpha_i\sigma(x_i).$$

When $\alpha_i = 1$, we recover σ , and when $\alpha_i = 0$, the activation function is linear. As activation functions, we use SELU [35] for classification and Mish [45] for regression.

- We use dropout after each activation function. We do not use the Alpha-dropout variant originally proposed for SELU [35], as we were not able to obtain good results with it.
- For regression, at test time, we clip the MLP outputs to the observed range during training. (We observed that this is mainly helpful for suboptimal hyperparameters.)

Initialization We initialize the parameters s_i of the scaling layer to 1, making it an identity function at initialization. We initialize the parameters α_i of the parametric activation functions to 1, recovering the standard activation functions at initialization. We initialize weights and biases in a data-dependent fashion during a forward pass on the (possibly subsampled) training set. We rescale rows of standard-

normal-initialized weight matrices to scale the variance of the output pre-activations over the dataset to one. For the biases, we use the data-dependent he+5 initialization method [60].

Training We use the Adam optimizer [34] in the AdamW variant [40] for weight decay. We set its momentum hyperparameters to $\beta_1 = 0.9$ and $\beta_2 = 0.95$ instead of the default $\beta_2 = 0.999$. The idea to use a smaller value for β_2 is adopted from the fastai tabular MLP [26]. We optimize for 256 epochs with a batch size of 256. As a loss function for classification, we use softmax + cross-entropy with label smoothing [61] with parameter $\varepsilon = 0.1$. To make the binary and multi-class cases more similar, we also use this loss function for binary classification instead of using a single output neuron with sigmoid and log-loss. For regression, we use the MSE loss and affinely transform the targets to have zero mean and unit variance on the training and validation set.

Hyperparameters We allow parameter-specific scheduled hyperparameters computed in each iteration using a base value, optional parameter-specific factors, and a schedule, as

$$\text{base_value} \cdot \text{param_factor} \cdot \text{schedule} \left(\frac{\text{iteration}}{\#\text{iterations}} \right),$$

allowing us, for example, to use a high learning rate factor for scaling layer parameters. Because we do not tune the number of epochs separately on each dataset, we use a multi-cycle learning rate schedule, providing multiple valleys that are usually preferable for stopping the training, while allowing high learning rates in between. Our schedule is similar to Loshchilov and Hutter [39] and Smith [58], but with a simpler analytical expression:

$$\text{coslog}_k(t) := \frac{1}{2}(1 - \cos(2\pi \log_2(1 + (2^k - 1)t))).$$

We set $k = 4$ to obtain four cycles as shown in Figure 1 (b). To allow stopping at different levels of regularization, we schedule dropout and weight decay using the following schedule, cf. Figure 1 (b):²

$$\text{flat_cos}(t) := \frac{1}{2}(1 + \cos(\pi(\max\{1, 2t\} - 1))).$$

The detailed hyperparameters can be found in Table A.1.

Best-epoch selection Due to the multi-cycle learning rate schedule, we do not perform classical early stopping. Instead, we always train for the full 256 epochs and then revert the model to the epoch with the lowest validation error, which in this paper is based on classification error, or RMSE for regression. In case of a tie, we found it beneficial to use the last of the tied best epochs.

RealMLP-TD-S Since certain aspects of RealMLP-TD are somewhat complex to implement, we introduce a simplified (and faster) variant called RealMLP-TD-S. Among the simplifications (see Appendix A) are: omitting embedding layers, using non-parametric activations, using a simpler initialization method, and omitting dropout and weight decay.

4 Gradient-Boosted Decision Trees

To find better default hyperparameters for GBDTs, we employ a semi-automatic approach: We use hyperparameter optimization libraries like hyperopt [2] and SMAC3 [38] to explore a reasonably large hyperparameter space, evaluating the benchmark score of each configuration on the meta-train benchmarks, and then perform some small manual adjustments like rounding the best obtained hyperparameters. To balance efficiency and accuracy, we fix the number of estimators to 1000 and use the `hist` method for XGBoost. We only consider the libraries’ default tree-building strategies since it is one of their main differences. The tuned defaults (TD) for LightGBM (LGBM), XGBoost (XGB), and CatBoost can be found in Table C.1, C.2, and C.3, respectively.

While some of the obtained hyperparameter values might be sensitive to the tuning and benchmark setup, we observe some general trends. First, row subsampling is applied in all settings, while column subsampling is rarely applied. Second, trees are generally allowed to be deeper for regression than for classification. Third, the Bernoulli bootstrap in CatBoost is competitive with the Bayesian bootstrap while also being faster.

²inspired by a similar schedule in <https://github.com/lessw2020/Ranger-Deep-Learning-Optimizer>

5 Experiments

In the following, we evaluate different methods with library defaults (D), tuned defaults (TD), and hyperparameter optimization (HPO). Recall that TD uses fixed parameters optimized on the meta-train benchmarks, while HPO tunes hyperparameters on each dataset split independently. All methods except random forests select the best iteration/epoch on the validation set of the respective dataset split. All NN-based regression methods standardize the labels for training.

5.1 Methods

We use the following abbreviations, see Appendix C for more details:

- **MLP-D, ResNet-D**: MLP and ResNet from Gorishniy et al. [16] with default parameters from McElfresh et al. [43].
- **TabR-S-D**: TabR-S with default parameters [18], using context freeze for one slow dataset.
- **RF-SKL-D**: Random Forest [5] implementation of scikit-learn [48] with default parameters.
- **CatBoost-D, LGBM-D, and XGB-D**: CatBoost, LightGBM, and XGBoost defaults.
- **XGB-PBB-D**: XGBoost with optimized defaults (for AUC) from Probst et al. [50], using “hist” boosting.
- **RealMLP-TD and RealMLP-TD-S**: Our MLPs with tuned defaults from Section 3.
- **CatBoost-TD, LGBM-TD, and XGB-TD**: Tuned defaults for CatBoost, LightGBM, and XGBoost, respectively, as presented in Section 4.
- **RealMLP-HPO, MLP-HPO, CatBoost-HPO, LGBM-HPO, and XGB-HPO**: We apply 50 random search steps to each of the models, see Appendix C.2.
- **Best-HPO, Best-TD, Best-D**: Selecting the best model out of RealMLP, CatBoost, XGBoost, and LightGBM in the respective HPO, TD, or D variant based on the validation error on each dataset. For Best-D, we use MLP-D instead of RealMLP.
- **Ensemble-HPO, Ensemble-TD, Ensemble-D**: Similar to Best, but creating a weighted ensemble using the validation set as proposed by Caruana et al. [6], with 40 greedy selection steps as in Salinas and Erickson [54].

5.2 Results

Figure 2 shows the results of the aforementioned methods on the meta-train and meta-test benchmarks, along with their runtimes on a CPU. We also show results for the benchmark of Grinsztajn et al. [19] in Figure 3, specifically the version with medium-size datasets ($\leq 10K$ training samples). The latter include scikit-learn’s GradientBoostingTree (**GBT-SKL**) as well as two transformer-based models, **FT-Transformer** [16] and **SAINT** [59]. For the baselines, we use (slightly adapted) HPO search spaces from the original paper, see Appendix C.4.

How good are tuned defaults on new datasets? To answer this question, we compare the relative gaps between TD and HPO benchmark scores on the meta-test benchmarks to those on the meta-train benchmarks. The gap between RealMLP-HPO and RealMLP-TD is about equally large on the meta-train and meta-test benchmarks, indicating that the tuned defaults transfer very well to the meta-test benchmark. For GBDTs, tuned defaults are competitive with HPO on the meta-train set, but not as good on the meta-test set. Still, they are considerably better than the untuned defaults on the meta-test set. Note that we did not limit the TD parameters to the literature search spaces for the HPO models (cf. Appendix C.2); for example, XGB-TD uses a smaller value of `min_child_weight` for classification and CatBoost-TD uses deeper trees and Bernoulli boosting. The XGBoost defaults XGB-PBB-D from Probst et al. [50] outperform XGB-TD on $\mathcal{B}_{\text{class}}^{\text{test}}$, perhaps because their benchmark is more similar to $\mathcal{B}_{\text{class}}^{\text{test}}$ or because XGB-PBB-D uses more estimators (4168) and deeper trees.

Among GBDTs, CatBoost defaults are better and slower. Several papers have found CatBoost to perform favorably among GBDTs while being more computationally expensive to train [51, 43, 9, 33, 68]. We observe the same for our tuned defaults on the meta-test benchmark.

RealMLP performs favorably among NNs. We can see on the meta-test benchmark as well as the benchmark by Grinsztajn et al. [19] that RealMLP-TD performs significantly better than MLP-D and ResNet-D. RealMLP-TD is slower on average, mainly due to the use of numerical embeddings and

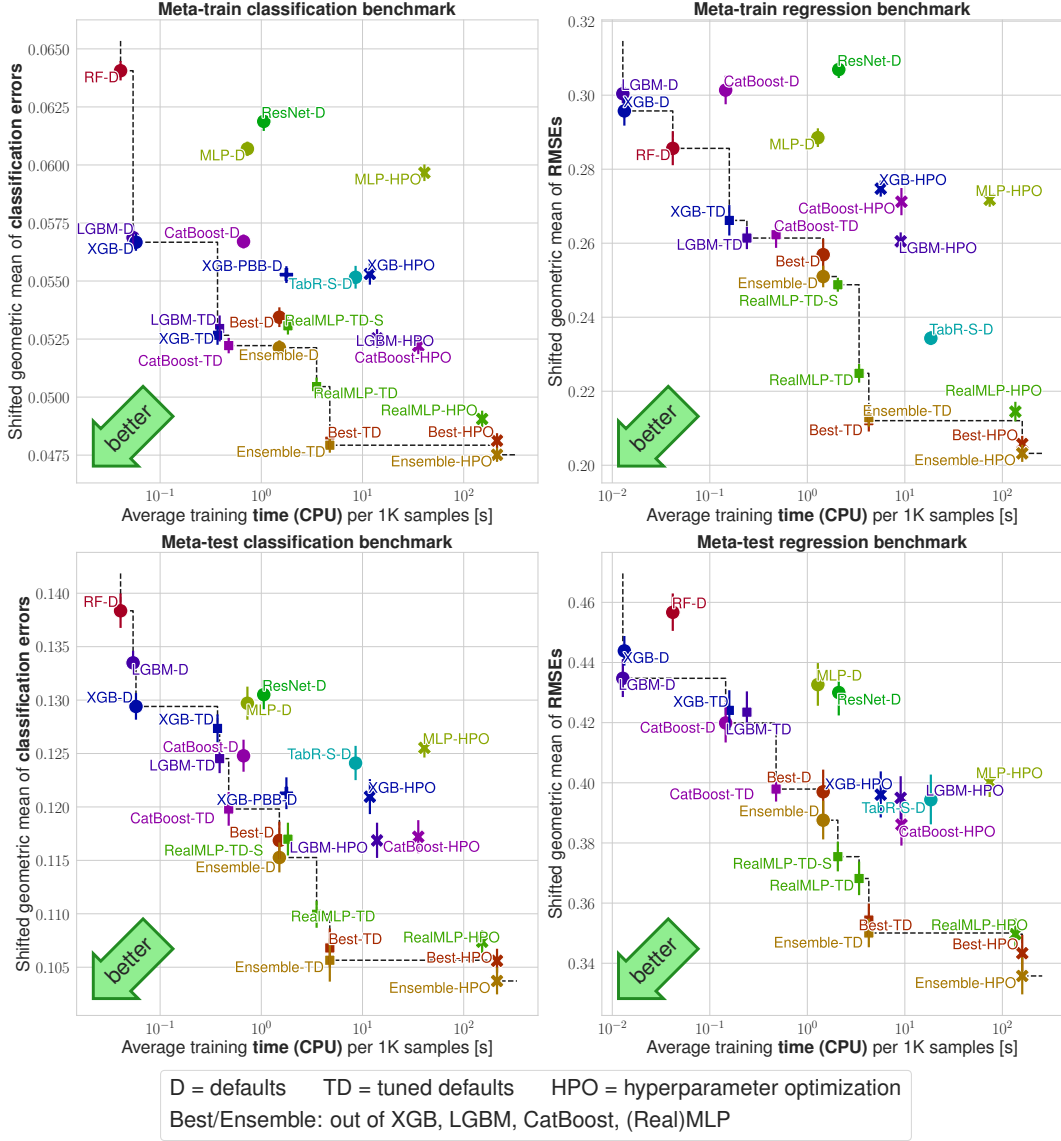


Figure 2: **Benchmark scores on $\mathcal{B}^{\text{train}}$, $\mathcal{B}^{\text{train}}_{\text{class}}$, $\mathcal{B}^{\text{test}}_{\text{class}}$, and $\mathcal{B}^{\text{test}}_{\text{reg}}$ vs. average training time.** The y -axis shows the shifted geometric mean (SGM $_{\epsilon}$) classification error (left) or RMSE (right) as explained in Section 2.2. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit #splits $\rightarrow \infty$, see Appendix C.5.

the lack of early stopping. However, the latter allows efficient cross-validation and ensemble training through vectorization. On the benchmark by Grinsztajn et al. [19], RealMLP-TD achieves equal or better results than other NNs, and in the HPO setting, it outperforms transformer-based models in terms of speed and benchmark scores. In addition, RealMLP-TD is faster than TabR-S-D while also achieving better meta-test benchmark scores. When looking at rank-based metrics in Appendix B, RealMLP-TD is worse than TabR-S-D on $\mathcal{B}^{\text{train}}_{\text{reg}}$ but better on $\mathcal{B}^{\text{train}}_{\text{class}}$ and $\mathcal{B}^{\text{test}}_{\text{class}}$.

NN improvements Figure 1 (c) shows how adding the proposed RealMLP components to a simple MLP improves the meta-train benchmark performance. However, these results depend on the order in which components are added, which is addressed by a separate ablation study in Appendix B. We also show in Appendix B.7 that our architectural improvements alone are beneficial when applied




Figure 3: **Results on the benchmarks of Grinsztajn et al. [19].** The y -axis (inverted) shows the normalized accuracy / R2 score used in the original paper (see Appendix C.4). The x -axis shows average training times per 1000 samples, using GPUs for NNs as in Grinsztajn et al. [19], see Appendix C.6.

to MLP-D directly, although non-architectural aspects are at least as important. In particular, our numerical preprocessing is easy to adopt and often beneficial for other NNs as well.

RealMLP is competitive with tree-based models. On the meta-test benchmark, in the TD and HPO setting, our MLP performs better than GBDTs in terms of shifted geometric mean error. When considering arithmetic mean error and average rank, RealMLP is comparable to or slightly better than GBDTs (Appendix B.9). The improvement in shifted geometric mean error reflects the fact that on several datasets, the NN has a lower error by roughly an order of magnitude. The win rate of RealMLP-TD vs. CatBoost-TD is around 51–62% and the win rate of RealMLP-HPO vs. CatBoost-HPO is around 55–60% (Appendix B.11). On the benchmark of Grinsztajn et al. [19], RealMLP performs worse for classification but comparably for regression.

Simply trying all default algorithms is very often better than (naive) single-algorithm HPO, and always faster. When comparing Best-TD to 50-step HPO on any individual algorithm, we notice that Best-TD is always faster, while also being better on the meta-train and meta-test benchmarks and only slightly worse than the best HPO models on the benchmark of Grinsztajn et al. [19]. We also note that ensemble selection [6] usually gives 1–3% improvement on the benchmark score compared to selecting the best model, and can potentially be further improved [7]. Our results indicate that our methods have the potential to be useful within AutoML methods, which may implement more sophisticated methods such as training-time-aware HPO or stacking. We note that this surprising performance directly results from the quality of our tuned default hyperparameters and RealMLP, as selecting or ensembling non-tuned defaults is matched by LGBM-HPO and outperformed by RealMLP-HPO (Figure 2).

Further insights In Appendix B, we present further experimental results, including an ablation study. We compare bagging and refitting for RealMLP-TD and LGBM-TD, finding that refitting multiple models is often better on average. We demonstrate that GBDTs benefit from high early stopping patiences for classification, especially when using accuracy as the stopping metric. When considering ROC AUC as a stopping metric, we show that stopping on cross-entropy is preferable to accuracy, and that label smoothing is harmful. We also investigate robust scaling and smooth clipping for other NNs, showing that it is often helpful.

Limitations In our benchmarks, we have considered medium-to-large tabular datasets with random train-test splits, using classification error and RMSE as metrics, with additional results for AUROC in Appendix B.5. It is unclear to which extent the obtained defaults can generalize to very small

datasets, distribution shifts, datasets with missing numerical values, and other metrics such as log-loss. Additionally, runtimes and the resulting tradeoffs may change with different parallelization, different (time-aware) HPO algorithms, or different hardware. For computational reasons, we only use a single training-validation split per train-test split. This means that HPO can overfit the validation set more easily than in a cross-validation setup. While we extensively benchmark different NN models from the literature, we do not attempt to equalize non-architectural aspects, and our work should therefore not be seen as a comparison of architectures. We compared to TabR-S-D as a recent promising method with good default parameters [18, 68]. However, due to a surge of recently published deep tabular models [e.g., 8, 9, 56, 41, 33, 66, 29], it is unclear what the current “best” deep tabular model is.

6 Conclusion

In this paper, we studied the potential of improved default parameters for GBDTs and an improved MLP, evaluated on a large separate meta-test benchmark as well as the benchmark by Grinsztajn et al. [19], and investigated the time-accuracy tradeoffs of various algorithm selection and ensembling scenarios. Our improved MLP mostly outperforms other NNs from the literature with moderate runtime and is competitive with GBDTs. Since many of the proposed improvements to NNs are orthogonal to the improvements in other papers, they offer exciting opportunities for combinations. While the “NNs vs GBDTs” debate remains interesting, our results demonstrate that with good default parameters, it is worth trying both algorithm families even with a moderate training time budget.

Acknowledgments and Disclosure of Funding

We thank Gaël Varoquaux, Frank Sehnke, Katharina Strecker, Ravid Shwartz-Ziv, Lennart Purucker, and Francis Bach for helpful discussions. We thank Katharina Strecker for help with code refactoring.

Funded by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany’s Excellence Strategy - EXC 2075 – 390740016. The authors thank the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for supporting David Holzmüller. LG acknowledges support in part by the French Agence Nationale de la Recherche under Grant ANR-20-CHIA-0026 (LearnI). Part of this work was performed on the computational resource bwUniCluster funded by the Ministry of Science, Research and the Arts Baden-Württemberg and the Universities of the State of Baden-Württemberg, Germany, within the framework program bwHPC. Part of this work was performed using HPC resources from GENCI-IDRIS (Grant 2023-AD011012804R1 and 2024-AD011012804R2).

Contribution statement DH and IS conceived the project. DH implemented and experimentally validated the newly proposed methods and wrote the initial paper draft. DH and LG contributed to benchmarking, plotting, and implementing baseline methods. LG and IS helped revise the draft. IS supervised the project and contributed dataset downloading code.

References

- [1] Sercan O. Arik and Tomas Pfister. TabNet: Attentive interpretable tabular learning. In *AAAI Conference on Artificial Intelligence*, 2021.
- [2] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In *International Conference on Machine Learning*, 2013.
- [3] Vadim Borisov, Tobias Leemann, Kathrin Seßler, Johannes Haug, Martin Pawelczyk, and Gjergji Kasneci. Deep neural networks and tabular data: A survey. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [4] Pavel Brazdil, Christophe Giraud Carrier, Carlos Soares, and Ricardo Vilalta. *Metalearning: Applications to Data Mining*. Springer Science & Business Media, 2008.
- [5] Leo Breiman. Random forests. *Machine learning*, 45:5–32, 2001.

- [6] Rich Caruana, Alexandru Niculescu-Mizil, Geoff Crew, and Alex Ksikes. Ensemble selection from libraries of models. In *International Conference on Machine Learning*, 2004.
- [7] Rich Caruana, Art Munson, and Alexandru Niculescu-Mizil. Getting the most out of ensemble selection. In *International Conference on Data Mining*, pages 828–833. IEEE, 2006.
- [8] Jintai Chen, Jiahuan Yan, Danny Ziyi Chen, and Jian Wu. ExcelFormer: A neural network surpassing GBDTs on tabular data. *arXiv:2301.02819*, 2023.
- [9] Kuan-Yu Chen, Ping-Han Chiang, Hsin-Rung Chou, Ting-Wei Chen, and Tien-Hao Chang. Trompt: Towards a better deep neural network for tabular data. In *International Conference on Machine Learning*, 2023.
- [10] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *International Conference on Knowledge Discovery and Data Mining*, 2016.
- [11] Nick Erickson, Jonas Mueller, Alexander Shirkov, Hang Zhang, Pedro Larroy, Mu Li, and Alexander Smola. AutoGluon-Tabular: Robust and accurate AutoML for structured data. In *7th ICML Workshop on Automated Machine Learning*, 2020.
- [12] Matthias Feurer, Katharina Eggenberger, Stefan Falkner, Marius Lindauer, and Frank Hutter. Auto-sklearn 2.0: Hands-free automl via meta-learning. *The Journal of Machine Learning Research*, 23(261), 2022.
- [13] Sebastian Felix Fischer, Matthias Feurer, and Bernd Bischl. OpenML-CTR23—A curated tabular regression benchmarking suite. In *AutoML Conference 2023 (Workshop)*, 2023.
- [14] Pieter Gijsbers, Marcos LP Bueno, Stefan Coors, Erin LeDell, Sébastien Poirier, Janek Thomas, Bernd Bischl, and Joaquin Vanschoren. AMLB: An AutoML benchmark. *arXiv:2207.12560*, 2022.
- [15] Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- [16] Yury Gorishniy, Ivan Rubachev, Valentin Khruikov, and Artem Babenko. Revisiting deep learning models for tabular data. *Neural Information Processing Systems*, 2021.
- [17] Yury Gorishniy, Ivan Rubachev, and Artem Babenko. On embeddings for numerical features in tabular deep learning. *Neural Information Processing Systems*, 2022.
- [18] Yury Gorishniy, Ivan Rubachev, Nikolay Kartashev, Daniil Shlenskii, Akim Kotelnikov, and Artem Babenko. TabR: Tabular deep learning meets nearest neighbors in 2023. *arXiv:2307.14338*, 2023.
- [19] Léo Grinsztajn, Edouard Oyallon, and Gaël Varoquaux. Why do tree-based models still outperform deep learning on typical tabular data? *Neural Information Processing Systems*, 2022.
- [20] Cheng Guo and Felix Berkhahn. Entity embeddings of categorical variables. *arXiv:1604.06737*, 2016.
- [21] Danijar Hafner, Jurgis Pasukonis, Jimmy Ba, and Timothy Lillicrap. Mastering diverse domains through world models. *arXiv:2301.04104*, 2023.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *IEEE International Conference on Computer Vision*, pages 1026–1034, 2015.
- [23] Steffen Herbold. Autorank: A Python package for automated ranking of classifiers. *Journal of Open Source Software*, 5(48):2173, 2020. doi: 10.21105/joss.02173. URL <https://doi.org/10.21105/joss.02173>. Publisher: The Open Journal.
- [24] Noah Hollmann, Samuel Müller, Katharina Eggenberger, and Frank Hutter. TabPFN: A transformer that solves small tabular classification problems in a second. In *International Conference on Learning Representations*, 2022.

- [25] David Holzmüller, Viktor Zaverkin, Johannes Kästner, and Ingo Steinwart. A framework and benchmark for deep batch active learning for regression. *Journal of Machine Learning Research*, 24(164), 2023.
- [26] Jeremy Howard and Sylvain Gugger. Fastai: A layered API for deep learning. *Information*, 11(2):108, 2020.
- [27] Xin Huang, Ashish Khetan, Milan Cvitkovic, and Zohar Karnin. TabTransformer: Tabular data modeling using contextual embeddings. *arXiv:2012.06678*, 2020.
- [28] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Neural Information Processing Systems*, 2018.
- [29] Manu Joseph and Harsh Raj. GANDALF: Gated Adaptive Network for Deep Automated Learning of Features. *arXiv:2207.08548*, 2024.
- [30] Arlind Kadra, Marius Lindauer, Frank Hutter, and Josif Grabocka. Well-tuned simple nets excel on tabular datasets. In *Neural Information Processing Systems*, 2021.
- [31] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. LightGBM: A highly efficient gradient boosting decision tree. In *Neural Information Processing Systems*, 2017.
- [32] Markelle Kelly, Rachel Longjohn, and Kolby Nottingham. The UCI Machine Learning Repository. URL <https://archive.ics.uci.edu>.
- [33] Myung Jun Kim, Léo Grinsztajn, and Gaël Varoquaux. CARTE: pretraining and transfer for tabular learning. *arXiv:2402.16785*, 2024.
- [34] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [35] Günter Klambauer, Thomas Unterthiner, Andreas Mayr, and Sepp Hochreiter. Self-normalizing neural networks. In *Neural Information Processing Systems*, 2017.
- [36] Ravin Kohli, Matthias Feurer, Katharina Eggensperger, Bernd Bischl, and Frank Hutter. Towards Quantifying the Effect of Datasets for Benchmarking: A Look at Tabular Machine Learning. In *ICLR 2024 Data-centric Machine Learning Research Workshop*, 2024.
- [37] Jannik Kossen, Neil Band, Clare Lyle, Aidan N. Gomez, Thomas Rainforth, and Yarin Gal. Self-attention between datapoints: Going beyond individual input-output pairs in deep learning. In *Neural Information Processing Systems*, 2021.
- [38] Marius Lindauer, Katharina Eggensperger, Matthias Feurer, André Biedenkapp, Difan Deng, Carolin Benjamins, Tim Ruhkopf, René Sass, and Frank Hutter. SMAC3: A versatile Bayesian optimization package for hyperparameter optimization. *Journal of Machine Learning Research*, 23(54), 2022.
- [39] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic gradient descent with warm restarts. In *International Conference on Learning Representations*, 2017.
- [40] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018.
- [41] Sascha Marton, Stefan Lüdtkke, Christian Bartelt, and Heiner Stuckenschmidt. GRANDE: Gradient-based decision tree ensembles for tabular data. In *International Conference on Learning Representations*, 2024.
- [42] Calvin McCarter. The kernel density integral transformation. *Transactions on Machine Learning Research*, 2023.
- [43] Duncan McElfresh, Sujay Khandagale, Jonathan Valverde, Ganesh Ramakrishnan, Micah Goldblum, and Colin White. When do neural nets outperform boosted trees on tabular data? *arXiv:2305.02997*, 2023.

- [44] Dmytro Mishkin and Jiri Matas. All you need is a good init. In *International Conference on Learning Representations*, 2016.
- [45] Diganta Misra. Mish: A self regularized non-monotonic activation function. In *British Machine Vision Conference*, 2020.
- [46] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, and Michael I. Jordan. Ray: A distributed framework for emerging AI applications. In *USENIX Symposium on Operating Systems Design and Implementation*, 2018.
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, and Luca Antiga. PyTorch: An imperative style, high-performance deep learning library. *Neural Information Processing Systems*, 32, 2019.
- [48] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, and Vincent Dubourg. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12(85), 2011.
- [49] Florian Pfisterer, Jan N. van Rijn, Philipp Probst, Andreas Müller, and Bernd Bischl. Learning multiple defaults for machine learning algorithms. *arXiv:1811.09409*, 2021.
- [50] Philipp Probst, Anne-Laure Boulesteix, and Bernd Bischl. Tunability: Importance of hyper-parameters of machine learning algorithms. *Journal of Machine Learning Research*, 20(53), 2019.
- [51] Liudmila Prokhorenkova, Gleb Gusev, Aleksandr Vorobev, Anna Veronika Dorogush, and Andrey Gulin. CatBoost: Unbiased boosting with categorical features. In *Neural Information Processing Systems*, 2018.
- [52] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. In *Neural Information Processing Systems*, 2007.
- [53] Hubert Ramsauer, Bernhard Schöfl, Johannes Lehner, Philipp Seidl, Michael Widrich, Lukas Gruber, Markus Holzleitner, Thomas Adler, David Kreil, and Michael K. Kopp. Hopfield networks is all you need. In *International Conference on Learning Representations*, 2020.
- [54] David Salinas and Nick Erickson. TabRepo: A Large Scale Repository of Tabular Model Evaluations and its AutoML Applications. *arXiv:2311.02971*, 2023.
- [55] Bernhard Schöfl, Lukas Gruber, Angela Bitto-Nemling, and Sepp Hochreiter. Modern Hopfield networks as memory for iterative learning on tabular data. In *NeurIPS Workshop on Associative Memory & Hopfield Networks in 2023*, 2023.
- [56] Junhong Shen, Liam Li, Lucio M. Dery, Corey Staten, Mikhail Khodak, Graham Neubig, and Ameet Talwalkar. Cross-Modal Fine-Tuning: Align then Refine. *arXiv:2302.05738*, 2023.
- [57] Ravid Shwartz-Ziv and Amitai Armon. Tabular data: Deep learning is not all you need. *Information Fusion*, 81:84–90, 2022.
- [58] Leslie N. Smith. Cyclical learning rates for training neural networks. In *Winter Conference on Applications of Computer Vision*, 2017.
- [59] Gowthami Somepalli, Micah Goldblum, Avi Schwarzschild, C. Bayan Bruss, and Tom Goldstein. SAINT: Improved neural networks for tabular data via row attention and contrastive pre-training. In *NeurIPS 2022 Table Representation Learning Workshop*, 2022.
- [60] Ingo Steinwart. A sober look at neural network initializations. *arXiv:1903.11482*, 2019.
- [61] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Computer Vision and Pattern Recognition*, 2016.

- [62] Jan N. van Rijn, Florian Pfisterer, Janek Thomas, Andreas Muller, Bernd Bischl, and Joaquin Vanschoren. Meta learning for defaults: Symbolic defaults. In *NeurIPS 2018 Workshop on Meta-Learning*, 2018.
- [63] Joaquin Vanschoren. Meta-learning: A survey. *arXiv:1810.03548*, 2018.
- [64] Joaquin Vanschoren, Jan N. van Rijn, Bernd Bischl, and Luis Torgo. OpenML: Networked science in machine learning. *ACM SIGKDD Explorations Newsletter*, 15(2):49–60, 2014. Publisher: ACM New York, NY, USA.
- [65] Martin Wistuba, Nicolas Schilling, and Lars Schmidt-Thieme. Learning hyperparameter optimization initializations. In *International Conference on Data Science and Advanced Analytics*, pages 1–10, 2015.
- [66] Chenwei Xu, Yu-Chao Huang, Jerry Yao-Chieh Hu, Weijian Li, Ammar Gilani, Hsi-Sheng Goan, and Han Liu. BiShop: Bi-directional cellular learning for tabular data with generalized sparse modern Hopfield model. *arXiv:2404.03830*, 2024.
- [67] Ge Yang, Edward Hu, Igor Babuschkin, Szymon Sidor, Xiaodong Liu, David Farhi, Nick Ryder, Jakub Pachocki, Weizhu Chen, and Jianfeng Gao. Tuning large neural networks via zero-shot hyperparameter transfer. In *Neural Information Processing Systems*, 2021.
- [68] Han-Jia Ye, Si-Yang Liu, Hao-Run Cai, Qi-Le Zhou, and De-Chuan Zhan. A closer look at deep learning on tabular data. *arXiv:2407.00956*, 2024.

Appendices

Appendix Contents.

A Further Details on RealMLP-TD and RealMLP-TD-S	16
A.1 RealMLP-TD Details	16
A.2 RealMLP-TD-S Details	16
A.3 Details on Cumulative Ablation	16
A.4 Discussion	18
B More Experiments	18
B.1 MLP Ablations	18
B.2 MLP Preprocessing	19
B.3 Bagging, Refitting, and Ensembling	20
B.4 Early stopping for GBDTs	21
B.5 Results for AUROC	22
B.6 Results Without Missing-Value Datasets	23
B.7 Results for Varying Architecture and Preprocessing	23
B.8 Comparing HPO Methods	26
B.9 More Time-Error Plots	27
B.10 Critical Difference Diagrams	27
B.11 Win-rate Plots	28
C Benchmark Details	35
C.1 Default Configurations	35
C.2 Hyperparameter Optimization	35
C.3 Dataset Selection and Preprocessing	35
C.4 Grinsztajn et al. benchmark	40
C.5 Confidence Intervals	42
C.6 Time Measurements	43
C.7 Compute Resources	44
C.8 Used Libraries	44
D Results for Individual Datasets	45

A Further Details on RealMLP-TD and RealMLP-TD-S

The detailed hyperparameter settings for RealMLP-TD and RealMLP-TD-S are listed in Table A.1.

A.1 RealMLP-TD Details

Initialization We initialize categorical embedding parameters from $\mathcal{N}(0, 1)$. We initialize the components of $\mathbf{w}_{\text{emb}}^{(1,i)}$ from $\mathcal{N}(0, 0.1^2)$ and of $\mathbf{b}_{\text{emb}}^{(1,i)}$ from $\mathcal{U}[-\pi, \pi]$. The other numerical embedding parameters are initialized according to PyTorch’s default initialization, that is, from the uniform distribution $\mathcal{U}[-1/\sqrt{16}, 1/\sqrt{16}]$. For weights and biases of the linear layers, we use a data-dependent initialization. The initialization is performed on the fly during a first forward pass of the network on the training set (which can be subsampled adaptively not to use more than 1 GB of RAM). We realize this by providing `fit_transform()` methods similar to a pipeline in scikit-learn. For the weight matrices, we use a custom two-step procedure: First, we initialize all entries from $\mathcal{N}(0, 1)$. Then, we rescale each row of the weight matrix such that the outputs $\frac{1}{\sqrt{d_l}} \mathbf{W}^{(l)} \mathbf{x}_j^{(l)}$ have variance 1 over the dataset (i.e. when considering the sample index $j \in \{1, \dots, n\}$ as a uniformly distributed random variable). This is somewhat similar to the LSUV initialization method [44]. For the biases, we use the data-dependent he+5 initialization method [60].

Training We implement weight decay as in PyTorch using $\theta \leftarrow \theta - \text{lr} \cdot \text{wd} \cdot \theta$, which includes the learning rate unlike the original version [40].

A.2 RealMLP-TD-S Details

For RealMLP-TD-S, we make the following changes compared to RealMLP-TD:

- We apply one-hot encoding to all categorical variables and do not apply categorical embeddings.
- We do not apply numerical embeddings.
- We use the standard non-parametric versions of the SELU and Mish activation functions.
- We do not use dropout and weight decay.
- We use simpler weight and bias initializations: We initialize weights and biases from $\mathcal{N}(0, 1)$, except in the last layer, where we initialize them to zero.
- We do not clip the outputs, even in the regression case.
- We apply a different base learning rate in the regression case.

A.3 Details on Cumulative Ablation

Here, we provide more details on the vanilla MLP and the ablation steps from Figure 1 (c). For each step, we choose the best default learning rate out of a learning rate grid, using $\{0.0004, 0.0007, 0.001, 0.0015, 0.0025, 0.004, 0.007, 0.01, 0.015\}$ for NNs using standard parametrization and $\{0.01, 0.02, 0.03, 0.04, 0.07, 0.1, 0.2, 0.3, 0.4\}$ for NNs using neural tangent parametrization.

- Vanilla MLP: We use three hidden layers with 256 hidden neurons in each layer, just like RealMLP-TD, and the ReLU activation function. Each linear layer uses standard parametrization and the PyTorch default initialization, which is uniform from $[-1/\sqrt{\text{fan_in}}, 1/\sqrt{\text{fan_in}}]$ for both weights and biases, where `fan_in` is the input dimension. Categorical features are embedded using embedding layers, using eight-dimensional embeddings for each feature. Numerical features are transformed using a scikit-learn `QuantileTransformer` to approximately normal-distributed features. Optimization is performed using Adam with constant learning rate and default parameters $\beta_1 = 0.9, \beta_2 = 0.999, \varepsilon = 10^{-8}$ for at most 256 epochs with batch size 256, with constant learning rate. If the best validation error (classification error or RMSE) does not improve for 40 epochs, training is stopped. In each case, the model is reverted to the parameters of the epoch with the best validation score, using the first best epoch in case of a tie.
- Robust scale + smooth clip: We replace the `QuantileTransformer` with robust scaling and smooth clipping.

Table A.1: Overview of hyperparameters for RealMLP-TD and RealMLP-TD-S.

Hyperparameter	RealMLP-TD		RealMLP-TD-S	
	classification	regression	classification	regression
Num. embedding type	PBLD	PBLD	None	None
Num. embedding periodic init std.	0.1	0.1	—	—
Num. embedding hidden dimension	16	16	—	—
Num. embedding dimension	4	4	—	—
Max one-hot size (without missing)	8	8	∞	∞
Num. preprocessing		robust scale + smooth clip		
Categorical embedding dimension	8	8	—	—
Categorical embedding initialization	$\mathcal{N}(0, 1)$	$\mathcal{N}(0, 1)$	—	—
Use scaling layer		yes		
Scaling layer initialization		1.0 (constant)		
Number of linear layers		4		
Hidden layer sizes		[256, 256, 256]		
Activation function	SELU	Mish	SELU	Mish
Use parametric activation function	yes	yes	no	no
Parametric activation function initialization	1.0	1.0	—	—
Linear layer parametrization		NTP		
Last linear layer weight initialization	data-driven	data-driven	zero	zero
Other linear layer weight initialization	data-driven	data-driven	std normal	std normal
Last linear layer bias initialization	he+5	he+5	zero	zero
Other linear layer bias initialization	he+5	he+5	std normal	std normal
Optimizer		AdamW		
Batch size		256		
Number of epochs		256		
Adam β_1		0.9		
Adam β_2		0.95		
Adam ε		1e-8		
Learning rate (base value)	0.04	0.2	0.04	0.07
Learning rate schedule		coslog ₄		
Learning rate (num. emb. factor)	0.1	0.1	—	—
Learning rate (scaling layer factor)		6		
Learning rate (bias factor)		0.1		
Learning rate (param. act. factor)	0.1	0.1	—	—
Dropout probability (base value)	0.15	0.15	0.0	0.0
Dropout schedule	flat_cos	flat_cos	—	—
Weight decay (base value)	0.02	0.02	0.0	0.0
Weight decay schedule	flat_cos	flat_cos	—	—
Weight decay (bias factor)	0.0	0.0	—	—
Loss function	cross-entropy	MSE	cross-entropy	MSE
Label smoothing ε	0.1	—	0.1	—
Standardize targets during training	—	yes	—	yes
Output min-max clipping	—	yes	—	no
Best epoch selection metric	class. error	MSE	class. error	MSE
Best epoch selection method		last best validation error		

- One-hot for small cat.: As in RealMLP-TD, we use one-hot encoding for categories with at most eight values, not counting missing values.
- No early stopping: We always train the full 256 epochs.
- Last best epoch: In case of a tie, we use the last of the best epochs.
- coslog₄ lr sched: We use the coslog₄ learning rate schedule instead of a constant one.
- Adam $\beta_2 = 0.95$: We set $\beta_2 = 0.95$.
- Label smoothing (class.): We enable label smoothing with $\varepsilon = 0.1$ in the classification case.
- Output clipping (reg.): For regression, outputs are clipped to the min-max range observed during training.
- NT parametrization: We use the neural tangent parametrization for linear layers, setting the bias learning rate factor to 0.1.
- Act. fn. SELU / Mish: We change the activation function from ReLU to SELU (classification) or Mish (regression).

- Parametric act. fn.: We use parametric versions of the activation functions, with a learning rate factor of 0.1 for the parameters.
- Scaling layer: We use a scaling layer with a learning rate factor of 6 before the first linear layer.
- Num. embeddings: PL: We apply the PL embeddings [17] to numerical features.
- Num. embeddings: PBLD: We apply our PBLD embeddings instead.
- Dropout $p = 0.15$: We apply dropout with probability 0.15.
- Dropout sched: flat_cos: We apply the flat_cos schedule to the dropout probability.
- Weight decay wd = 0.02: We apply weight decay (as in AdamW, PyTorch version) with value 0.02.
- wd sched: flat_cos: We apply the flat_cos schedule to weight decay.
- Bias init: he+5: We apply the he+5 bias initialization method from Steinwart [60].
- Weight init: data-driven: We apply our data-driven weight initialization method.

A.4 Discussion

Here, we discuss some of the design decisions behind RealMLP-TD and possible trade-offs. First, our implementation allows us to train RealMLP-TD in a vectorized fashion on multiple train-validation-test splits at the same time. On the one hand, this can lead to speedups on GPUs when training multiple models in parallel, including on the benchmarks. On the other hand, it can hinder the implementation of certain methods like patience-based early stopping or loss-based learning rate schedules. While our ablations in Appendix B.1 show the advantage of our multi-cycle schedule over decreasing learning rate schedules, the latter ones could potentially enable a faster average training time through low-patience early stopping. An interesting follow-up question could be whether the multi-cycle schedule still works well with larger-patience early stopping.

Regarding categorical embeddings, our meta-train benchmark does not contain many high-cardinality categorical variables, and we were not able to conclude whether categorical embeddings are helpful or harmful compared to one-hot encoding (see Appendix B.1). Our motivation to include categorical embeddings stems from Guo and Berkhahn [20] as well as their potential to be more efficient for high-cardinality categorical variables. However, in practice, we find pure one-hot encoding to be faster on most datasets. Regarding the embedding size, we found that 4 already gave good results for numerical embeddings and decided to use 8 for categorical variables.

Additionally, other speed-accuracy tradeoffs are possible. Especially for regression, we observed that more epochs and larger hidden layers can be helpful. When faster networks are desired, the omission of numerical and categorical embedding layers as well as parametric activations from RealMLP-TD can be helpful, while the other omissions in RealMLP-TD-S do not considerably affect the training time. Of course, using larger batch sizes can also be helpful for larger datasets.

One caveat for classification is that cross-entropy with label smoothing is not a proper scoring rule, that is, in the infinite-sample limit, it is not minimized by the true probabilities $P(y|x)$ [15]. Hence, label smoothing might not be suitable when other classification error metrics are used, as demonstrated in Appendix B.5 for AUROC.

B More Experiments

B.1 MLP Ablations

To assess the importance of different improvements in RealMLP-TD, we perform an ablation study. We perform the ablation study only on the *meta-train* benchmarks, first because they are considerably faster to run, and second because we tune the default parameters only on the meta-train benchmarks. Since the hyperparameters of RealMLP-TD have been tuned on the meta-train benchmarks, the ablation scores are not unbiased but represent some of the considerations that have been made when tuning the defaults. For each ablation, we multiply the default learning rate by learning rate factors from the grid $\{0.1, 0.15, 0.25, 0.35, 0.5, 0.7, 1.0, 1.4, 2.0, 3.0, 4.0\}$ and pick the best one. Table B.1 shows the results of the ablation study in terms of the relative increase of the benchmark score for each ablation.

In general, we observe that ablations often lead to much larger changes for regression than for classification. Perhaps this is because RMSE is more sensitive compared to classification error.

Another factor could be that the classification benchmark contains more datasets than the regression benchmark. For the specific ablations, we observe a few things:

- For the **numerical embeddings**, we see that PBLD outperforms PL, PLR, and no numerical embeddings. Contrary to Gorishniy et al. [17], PL embeddings perform better than PLR embeddings in our setting. While the configurations with PLR and no numerical embeddings appear extremely bad for regression, we observed that they can perform more benignly with lower weight decay values.
- Using the Adam default value of $\beta_2 = 0.999$ instead of our default $\beta_2 = 0.95$ leads to considerably worse performance, especially for regression. As for numerical embeddings, we observed that the difference is less pronounced at lower weight decay values.
- Using a cosine decay **learning rate schedule** instead of our multi-cycle schedule leads to small deteriorations. A constant learning rate schedule performs even worse, especially for regression.
- Not employing **label smoothing** for classification is detrimental by around 1.8%.
- The **learnable scaling** layer yields improvements around 1.2% on both benchmarks.
- The use of **parametric activations** results in a considerable 4.8% improvement for regression but is insignificant for classification. We observed that parametric activations can sometimes alleviate optimization difficulties with weight decay.
- The differences between **activation functions** are rather small. For classification, Mish is competitive with SELU in this ablation but we found it to be worse in some other hyperparameter settings, so we keep SELU as the default. For regression, Mish performs best.
- For **dropout** and **weight decay**, we observe that they yield comparable but not always significant benefits for classification and regression. Scheduling dropout and weight decay parameters with the flat_cos schedule is helpful for regression, but not for classification in this setting.
- When comparing the **standard parametrization** (SP) to the neural tangent parametrization (NTP), we disable weight decay for a fair comparison. Moreover, for SP, we set the learning rate factors for weight and bias layers to $1/16 = 1/\sqrt{256}$. This is because, for the weights in NTP, the effective updates by Adam are damped by this factor in all hidden layers except the first one. Compared to NTP without weight decay, SP without weight decay performs insignificantly worse on both benchmarks. It is unclear to us why the parametrization, which has a considerable influence on how the effective learning speed of the first linear layer scales with the number of features, is apparently of little importance.
- When comparing the data-dependent **initialization** of RealMLP-TD to a vanilla initialization with standard normal weights and zero biases, we see that the data-dependent initialization gains around 1% on both benchmarks.
- For selecting the best epoch, we consider selecting the **first best epoch** instead of the last best epoch in case of a tie. This is only relevant for classification metrics like classification error, where ties are somewhat likely to occur, especially on small and “easy” datasets. We observe a non-significant 0.4% deterioration in the benchmark score.
- We do not observe a significant difference when using **one-hot encoding** for all categorical variables, since our benchmarks contain only very few datasets with large-cardinality categorical variables.

B.2 MLP Preprocessing

In Table B.2, we compare different preprocessing methods for numerical features. Since we want to compare these methods in a relatively conventional setting, we apply them to RealMLP-TD-S (without numerical embeddings) and before one-hot encoding. We compare the following methods:

- Robust scaling and smooth clipping, our method used in RealMLP-TD and RealMLP-TD-S and described in Section 3.
- Robust scaling without smooth clipping.
- Standardization, i.e. subtracting the mean and dividing by the standard deviation. If the standard deviation of a feature is zero, we set the feature to zero.
- Standardization followed by smooth clipping.
- The quantile transformation from scikit-learn [48] with normal output distribution, which is popular in recent works [19, 17, 18, 43].

Table B.1: **Ablation experiments for RealMLP-TD.** We re-tune the learning rate (picking the one with the best SGM_ε benchmark score) for each ablation separately. For each ablation, we specify the increase in the benchmark score (SGM_ε) relative to RealMLP-TD, with approximate 95% confidence intervals (Appendix C.5), and the best learning rate factor found. In the cases where values are missing, the corresponding option is already the default.

Ablation	meta-train-class		meta-train-reg	
	Error increase in %	best lr factor	Error increase in %	best lr factor
MLP-TD (without ablation)	0.0 [0.0, 0.0]	1.0	0.0 [0.0, 0.0]	1.0
Num. embeddings: PL	0.7 [-0.0, 1.4]	1.0	0.5 [-0.5, 1.6]	1.0
Num. embeddings: PLR	4.2 [2.8, 5.7]	1.0	19.0 [13.7, 24.5]	0.25
Num. embeddings: None	2.3 [1.7, 2.9]	1.0	20.6 [19.4, 21.8]	0.25
Adam $\beta_2 = 0.999$ instead of $\beta_2 = 0.95$	2.0 [1.6, 2.4]	2.0	22.8 [21.3, 24.4]	0.35
Learning rate schedule = cosine decay	1.1 [0.6, 1.5]	1.0	0.4 [-0.5, 1.2]	3.0
Learning rate schedule = constant	1.8 [0.9, 2.8]	0.25	13.4 [11.9, 15.0]	0.15
No label smoothing	1.8 [1.2, 2.5]	4.0		
No learnable scaling	1.4 [0.7, 2.1]	2.0	1.0 [-0.0, 2.0]	2.0
Non-parametric activation	0.5 [-0.2, 1.2]	3.0	4.8 [3.4, 6.2]	0.35
Activation=Mish	-0.0 [-0.6, 0.6]	3.0		
Activation=ReLU	0.5 [-0.1, 1.2]	2.0	0.7 [-0.1, 1.6]	1.0
Activation=SELU			2.3 [1.1, 3.6]	1.0
No dropout	0.8 [0.2, 1.3]	3.0	0.8 [-0.5, 2.1]	1.4
Dropout prob. 0.15 (constant)	-0.1 [-1.0, 0.8]	1.4	3.6 [3.0, 4.2]	1.0
No weight decay	0.8 [-0.2, 1.8]	0.5	0.9 [-0.1, 1.9]	0.5
Weight decay = 0.02 (constant)	-0.3 [-0.7, 0.1]	3.0	3.1 [1.7, 4.4]	1.4
Standard param + no weight decay	1.1 [0.2, 2.1]	0.5	1.3 [0.7, 1.8]	0.7
No data-dependent init	0.9 [0.1, 1.8]	3.0	1.2 [0.2, 2.2]	1.4
First best epoch instead of last best	0.4 [-0.1, 1.0]	4.0	0.0 [-0.0, 0.0]	1.0
Only one-hot encoding	-0.0 [-0.1, 0.0]	1.0	0.0 [-0.0, 0.0]	1.0

- A variant of the quantile transform, which we call the RTDL version, used by Gorishniy et al. [16] and Gorishniy et al. [18]. This version uses a dataset size dependent number of quantiles and adds some noise before fitting the transformation.
- The recent kernel density integral transform [42], which interpolates between the quantile transformation and min-max scaling, with default parameter $\alpha = 1$.

Table B.2 shows that on the meta-train benchmark, robust scaling and smooth clipping performs best for both classification and regression.

B.3 Bagging, Refitting, and Ensembling

In our benchmark, for each training-test split, we only train one model on one training-validation split for efficiency reasons. However, ensembling and cross-validation techniques usually allow additional improvements to models. Here, we study multiple variants for RealMLP-TD and LGBM-TD. Let \mathcal{D} be the available data for training and validation, split into five equal-size subsets $\mathcal{D}_1, \dots, \mathcal{D}_5$. (When $|\mathcal{D}|$ is not divisible by five, $\mathcal{D}_1 \cup \dots \cup \mathcal{D}_5 \subsetneq \mathcal{D}$ since we need equal-size validation sets for vectorized NNs.) Let $f_{\mathcal{D},t}(X)$ be the predictions on inputs X of the model trained on training set \mathcal{D} after $t \in \{1, \dots, T\}$ epochs (for NNs) or iterations (for LGBM). For classification, we consider the class probabilities as predictions. Let $L_{\mathcal{D}'}(f_{\mathcal{D},t})$ be the loss of $f_{\mathcal{D},t}$ on dataset \mathcal{D}' . Then, we compare the test errors of an ensemble of $M = 1$ or $M = 5$ models, trained using bagging or refitting, with individual or joint stopping (best-epoch selection), which is formally given as follows:

$$y_{\text{pred}} := \frac{1}{M} \sum_{i=1}^M f_{\mathcal{D}_i, t_i^*}(X_{\text{test}}), \quad (M \text{ models})$$

Table B.2: **Effects of different preprocessing methods for numerical features for RealMLP-TD-S.** We report the relative increase in the shifted geometric mean benchmark scores compared to the standard method used in RealMLP-TD and RealMLP-TD-S, which is robust scaling and smooth clipping. We also report approximate 95% confidence intervals. To have a more common setting, we do not apply the preprocessing methods to one-hot encoded categorical features. In each column, the best score is highlighted in bold, and errors whose confidence interval contains the best score are underlined.

Method	Error increase relative to robust scale + smooth clip in %	
	meta-train-class	meta-train-reg
Robust scale + smooth clip	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]
Robust scale	<u>0.5</u> [-0.4, 1.4]	9.5 [4.4, 14.8]
Standardize + smooth clip	1.6 [0.9, 2.2]	1.2 [0.6, 1.8]
Standardize	2.1 [1.2, 3.0]	8.7 [3.9, 13.9]
Quantile transform (output dist. = normal)	2.3 [1.5, 3.2]	6.3 [5.5, 7.0]
Quantile transform (RTDL version)	2.6 [1.5, 3.7]	2.6 [0.4, 4.8]
KDI transform ($\alpha = 1$, output dist. = normal)	4.9 [3.8, 6.0]	4.4 [2.6, 6.2]

Table B.3: **Improvements for LGBM-TD by bagging or (ensembled) refitting.** We perform 5-fold cross-validation, stratified for classification, and 5-fold refitting. We compare compare bagging vs. refitting, one model vs. five models, and individual stopping vs. joint stopping. The table shows the relative reduction in shifted geometric mean benchmark scores, including approximated 95% confidence intervals (Appendix C.5). In each column, the best score is highlighted in bold, and errors whose confidence interval contains the best score are underlined.

Method	Error reduction relative to 1 fold in %			
	meta-train-class	meta-test-class	meta-train-reg	meta-test-reg
LGBM-TD (bagging, 1 model, indiv. stopping)	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]
LGBM-TD (bagging, 1 model, joint stopping)	0.0 [-0.2, 0.2]	-0.6 [-1.1, -0.2]	-0.0 [-0.1, 0.0]	0.2 [-0.1, 0.6]
LGBM-TD (bagging, 5 models, indiv. stopping)	3.7 [3.1, 4.4]	4.4 [3.8, 5.0]	3.9 [2.2, 5.6]	4.2 [3.6, 4.8]
LGBM-TD (bagging, 5 models, joint stopping)	3.5 [2.8, 4.2]	3.6 [3.0, 4.2]	<u>3.9</u> [2.2, 5.5]	4.2 [3.7, 4.8]
LGBM-TD (refitting, 1 model, indiv. stopping)	5.2 [4.5, 5.8]	1.7 [-0.5, 3.9]	<u>2.4</u> [-0.1, 4.9]	4.2 [3.6, 4.8]
LGBM-TD (refitting, 1 model, joint stopping)	5.3 [4.8, 5.9]	4.7 [3.9, 5.5]	<u>2.4</u> [-0.0, 4.7]	4.2 [3.7, 4.8]
LGBM-TD (refitting, 5 models, indiv. stopping)	6.0 [5.4, 6.5]	6.3 [5.2, 7.5]	<u>3.8</u> [1.6, 6.0]	5.7 [5.1, 6.3]
LGBM-TD (refitting, 5 models, joint stopping)	<u>5.8</u> [5.2, 6.4]	<u>6.2</u> [5.4, 7.1]	<u>3.9</u> [1.6, 6.1]	<u>5.6</u> [5.1, 6.1]

$$\tilde{\mathcal{D}}_i := \begin{cases} \mathcal{D} \setminus \mathcal{D}_i & \text{(bagging)} \\ \mathcal{D} & \text{(refitting)}, \end{cases}$$

$$t_i^* := \begin{cases} \operatorname{argmin}_{t \in \{1, \dots, T\}} L_{\mathcal{D}_i}(f_{\mathcal{D} \setminus \mathcal{D}_i, t}) & \text{(indiv. stopping)} \\ \operatorname{argmin}_{t \in \{1, \dots, T\}} \sum_{j=1}^5 L_{\mathcal{D}_j}(f_{\mathcal{D} \setminus \mathcal{D}_j, t}) & \text{(joint stopping)}. \end{cases}$$

Here, each model is trained with a different random seed. For LGBM, since we use an early stopping patience of 300 for each of the individual models, the argmin in the definition of t_i^* can only go up to the minimum stopping iteration T across the considered models.

The results of our experiments can be found in Table B.3 for LGBM-TD and in Table B.4 for RealMLP-TD. As expected, five models are considerably better than one. We find that refitting is mostly better than bagging, although a disadvantage of refitted models is that no validation scores are available, and it is unclear how HPO would affect this comparison. Comparing individual stopping to joint stopping, we find that individual stopping has a slight advantage in five-model bagging, while joint stopping performs better for single-model refitting. In the other two scenarios, joint stopping appears slightly better for RealMLP-TD and slightly worse for LGBM-TD. We also observe that the benefit of using five models instead of one appears to be larger for RealMLP-TD than for LGBM-TD.

B.4 Early stopping for GBDTs

In Figure B.1 and Figure B.2, we study the influence of different early stopping patiences and metrics on the resulting benchmark performance of XGB-TD, LGBM-TD, and CatBoost-TD. While the regression results only deteriorate slightly for low patiences of 10 or 20 iterations, classification results are much more hurt by low patiences. In the classification setting, we evaluate the use of

Table B.4: **Improvements for RealMLP-TD by bagging or (ensembled) refitting.** We perform 5-fold cross-validation, stratified for classification, and 5-fold refitting. We compare bagging vs. refitting, one model vs. five models, and individual stopping vs. joint stopping. The table shows the relative reduction in shifted geometric mean benchmark scores, including approximated 95% confidence intervals (Appendix C.5). In each column, the best score is highlighted in bold, and errors whose confidence interval contains the best score are underlined.

Method	Error reduction relative to 1 fold in %			
	meta-train-class	meta-test-class	meta-train-reg	meta-test-reg
RealMLP-TD (bagging, 1 model, indiv. stopping)	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]
RealMLP-TD (bagging, 1 model, joint stopping)	1.1 [0.0, 2.2]	0.7 [0.2, 1.2]	0.5 [0.2, 0.8]	0.8 [-0.5, 2.1]
RealMLP-TD (bagging, 5 models, indiv. stopping)	6.2 [5.2, 7.2]	8.0 [7.3, 8.7]	6.6 [6.0, 7.1]	4.9 [4.3, 5.6]
RealMLP-TD (bagging, 5 models, joint stopping)	6.2 [5.3, 7.1]	7.6 [6.7, 8.5]	6.5 [5.9, 7.1]	4.6 [3.7, 5.4]
RealMLP-TD (refitting, 1 model, indiv. stopping)	2.2 [1.2, 3.3]	3.5 [2.4, 4.7]	2.6 [1.6, 3.5]	1.0 [0.0, 2.0]
RealMLP-TD (refitting, 1 model, joint stopping)	4.8 [3.9, 5.6]	5.0 [4.2, 5.8]	4.3 [3.3, 5.4]	2.4 [1.2, 3.5]
RealMLP-TD (refitting, 5 models, indiv. stopping)	7.1 [6.4, 7.7]	9.1 [8.3, 9.8]	8.3 [7.6, 9.0]	5.1 [4.2, 6.0]
RealMLP-TD (refitting, 5 models, joint stopping)	7.7 [6.9, 8.6]	8.9 [8.3, 9.4]	8.5 [7.8, 9.2]	5.5 [4.6, 6.4]

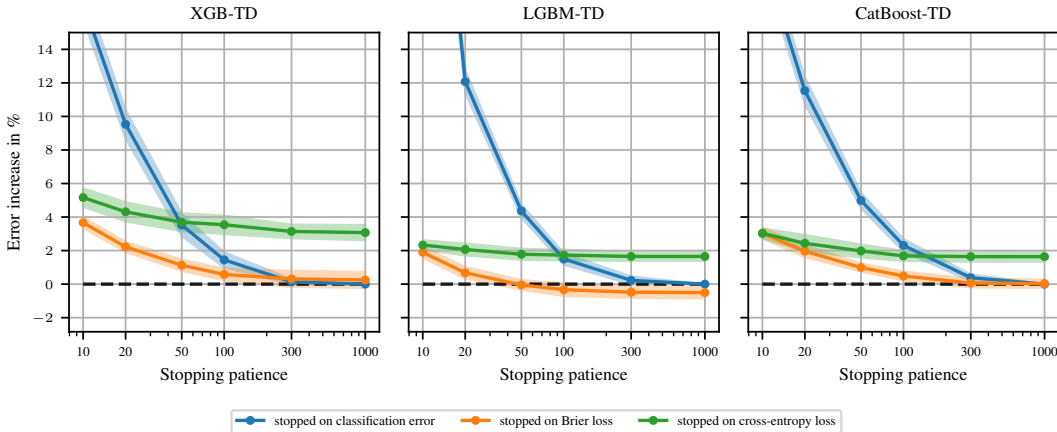


Figure B.1: **Effect of stopping patiences and metrics on the performance of GBDTs on $\mathcal{B}_{\text{class}}^{\text{train}}$.** We run the XGB-TD, LGBM-TD, and CatBoost-TD with different early stopping patiences (`early_stopping_rounds`). We compare three different metrics used for stopping and best-epoch selection: classification error, Brier loss, and cross-entropy loss. The y -axis reports the relative increase in the benchmark score relative to stopping on classification error with patience 1000 (i.e., never stopping early). The shaded areas are approximate 95% confidence intervals, cf. Appendix C.5.

different losses for early stopping and for best-epoch selection: classification error, Brier score, and cross-entropy loss. In each case, cross-entropy loss is used as the training loss, and classification error is used for evaluating the models on the test sets in the computation of the benchmark score. We observe that models stopped on classification error strongly deteriorate at low patiences ($\lesssim 100$), while our default patience of 300 achieves close-to-optimal results. Models stopped on cross-entropy loss deteriorate much less at low patiences, but achieve roughly 2% worse benchmark score at high patiences. Stopping on Brier loss achieves very good high-patience performance and is still only slightly more sensitive to the patience than stopping on cross-entropy loss. An interesting follow-up question would be if HPO can attenuate the differences between different settings.

B.5 Results for AUROC

For classification, there are many different metrics to capture model performance. In the main paper, we use classification error to evaluate models. All TD configurations were tuned for classification error, early stopping and best-epoch selection were performed for classification error, and HPO was performed for classification error. Here, we evaluate models on the area under the ROC curve, also known as AUROC, AUC ROC, or AUC. For the multi-class case, we use the one-vs-rest formulation of AUC, which is faster to evaluate than one-vs-one. Higher AUC values are better and the optimal value is 1. Since we are interested in the shifted geometric mean error, we use $1 - \text{AUC}$ instead.




Figure B.2: **Effect of stopping patiences on the performance of GBDTs on $\mathcal{B}_{\text{reg}}^{\text{train}}$** . We run the TD configurations of XGB, LGBM, and CatBoost with different early stopping patiences (early_stopping_rounds). As in the remainder of the paper, we use RMSE for early stopping and best-epoch selection. The y -axis reports the relative increase in the benchmark score relative to stopping on classification error with patience 1000 (i.e., never stopping early). The shaded areas are approximate 95% confidence intervals, cf. Appendix C.5.

We compare two settings:

- (1) A variant of the original setting where early stopping and the selection of the best epoch/iteration is based on accuracy but HPO is performed on $1 - \text{AUC}$. (Thanks to using random search, we do not have to re-run the HPO for this.)
- (2) A setting where we use the cross-entropy loss for stopping and selecting the best epoch/iteration. While it would be possible to stop on AUC directly, this can be significantly slower since AUC is slower to evaluate. We do not perform HPO in this setting since it is expensive to run.

In both settings, we also evaluate RealMLP without label smoothing (no ls). Figure B.4 shows the results optimized for accuracy and Figure B.3 shows the results optimized for cross-entropy. We make a few observations:

- Stopping for cross-entropy generally performs better than stopping for classification error.
- Label smoothing harms RealMLP for AUC, perhaps because the stopping metric does not use label smoothing, or because it encourages near-constant logits in areas where the model is relatively certain.
- Tuned defaults are mostly still better than the library defaults, except for XGBoost on $\mathcal{B}_{\text{class}}^{\text{test}}$.
- RealMLP without label smoothing is still competitive with GBDTs on the meta-test benchmark but does not perform better than GBDTs unlike what we observed for classification error.

B.6 Results Without Missing-Value Datasets

To assess whether the results are influenced by our choices in missing value handling and exclusion, Figure B.5 presents results on all meta-test datasets that originally did not contain missing values. Only six meta-test datasets originally contain missing values: Three from $\mathcal{B}_{\text{class}}^{\text{test}}$ (kick, okcupid-stem, and porto-seguro) and three from $\mathcal{B}_{\text{reg}}^{\text{test}}$ (fps_benchmark, house_prices_nominal, SAT11-HAND-runtime-regression). While RealMLP deteriorates slightly, especially due to the exclusion of fps_benchmark, qualitative takeaways remain similar.

B.7 Results for Varying Architecture and Preprocessing

Table B.5 shows the effects of including the preprocessing and architecture of RealMLP within other models. In particular,

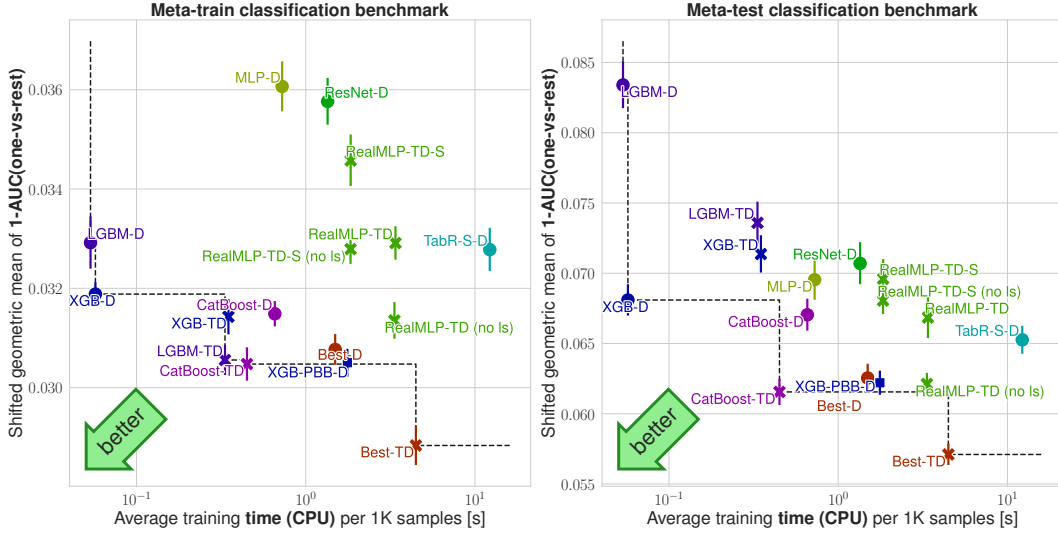


Figure B.3: **Benchmark scores on $\mathcal{B}_{\text{class}}^{\text{train}}$ and $\mathcal{B}_{\text{class}}^{\text{test}}$ vs. average training time for AUC, optimized for cross-entropy.** BestModel-TD uses RealMLP-TD without label smoothing. The y -axis shows the shifted geometric mean (SGM $_{\epsilon}$) $1 - \text{AUC}$ as explained in Section 2.2. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit #splits $\rightarrow \infty$, see Appendix C.5.

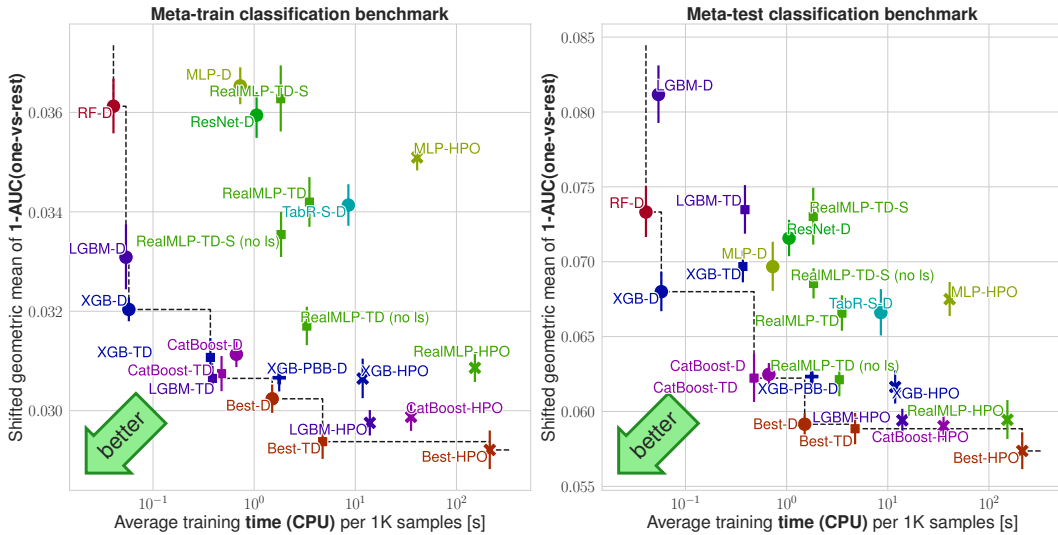


Figure B.4: **Benchmark scores on $\mathcal{B}_{\text{class}}^{\text{train}}$ and $\mathcal{B}_{\text{class}}^{\text{test}}$ vs. average training time for AUC.** Stopping and best-epoch selection are performed on accuracy, while HPO is performed on AUC. The y -axis shows the shifted geometric mean (SGM $_{\epsilon}$) $1 - \text{AUC}$ as explained in Section 2.2. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit #splits $\rightarrow \infty$, see Appendix C.5.




Figure B.5: **Benchmark scores on $\mathcal{B}_{\text{class}}^{\text{test}}$ and $\mathcal{B}_{\text{reg}}^{\text{test}}$ without missing value datasets vs. average training time.** The y -axis shows the shifted geometric mean (SGM $_{\epsilon}$) classification error (left) or RMSE (right) as explained in Section 2.2. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit #splits $\rightarrow \infty$, see Appendix C.5.

Table B.5: **Comparison of preprocessing and architecture for different models.** We include variants with robust scaling and smooth clipping (RS+SC), as well as other modified aspects, cf. Appendix B.7. We report the relative decrease in the shifted geometric mean benchmark scores compared to MLP-D. We also report approximate 95% confidence intervals, cf. Appendix C.5.

Method	Error reduction relative to MLP-D in %			
	meta-train-class	meta-train-reg	meta-test-class	meta-test-reg
MLP-D	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]	-0.0 [-0.0, -0.0]
MLP-D (RS+SC)	1.5 [0.7, 2.4]	-1.6 [-1.9, -1.2]	-0.7 [-1.6, 0.2]	4.3 [3.4, 5.2]
MLP-D (RS+SC, no wd, meta-tuned lr)	2.5 [1.8, 3.3]	-1.0 [-1.5, -0.5]	-1.6 [-2.7, -0.6]	4.3 [3.3, 5.3]
MLP-D (RS+SC, no wd, meta-tuned lr, PL embeddings)	4.6 [4.0, 5.2]	-1.5 [-1.9, -1.0]	-10.9 [-12.3, -9.4]	5.4 [4.0, 6.9]
MLP-D (RS+SC, no wd, meta-tuned lr, RealMLP architecture)	7.7 [6.9, 8.5]	10.4 [9.4, 11.3]	3.2 [2.0, 4.4]	9.6 [8.6, 10.6]
RealMLP-TD-S	12.6 [11.9, 13.2]	13.8 [13.2, 14.4]	9.8 [8.4, 11.2]	13.2 [12.1, 14.3]
RealMLP-TD	16.9 [16.1, 17.6]	22.1 [21.2, 22.9]	15.2 [14.0, 16.5]	14.9 [14.0, 15.8]
TabR-S-D	9.1 [8.2, 10.1]	18.8 [18.3, 19.3]	4.3 [3.0, 5.6]	8.8 [7.8, 9.8]
TabR-S-D (RS+SC)	12.4 [11.6, 13.1]	21.9 [21.1, 22.7]	6.9 [5.6, 8.3]	11.6 [10.4, 12.8]
ResNet-D	-1.9 [-3.0, -0.9]	-6.4 [-7.0, -5.8]	-0.6 [-1.3, 0.1]	0.6 [-0.4, 1.6]
ResNet-D (RS+SC)	2.0 [1.3, 2.8]	-5.9 [-6.6, -5.2]	-1.5 [-2.7, -0.4]	2.3 [1.4, 3.3]

- we include robust scaling and smooth clipping (RS+SC) into MLP-D, ResNet-D, and TabR-S-D. By default, these models use a version of the quantile transform (see Appendix B.2). We observe that it is beneficial on average but not on every benchmark for MLP-D and ResNet-D, while providing notable gains on every benchmark for TabR-S-D.
- we study the benefits of our architectural changes, cf. Figure 1 (c), when applied directly to the setting of MLP-D. To this end, we approximately reproduce MLP-D in our codebase without weight decay (since the optimal value changes when including the NTP) and with marginally different early stopping thresholding logic. We also determine the best default learning rate on the meta-train benchmark, similar to Appendix A.3. Our reproduction achieves benchmark scores within 1% of the benchmark scores of the MLP-D (RS+SC) version. Adding the PL embeddings from Gorishniy et al. [17] with our default settings sometimes gives good results but is significantly worse on $\mathcal{B}_{\text{class}}^{\text{test}}$, indicating that they need more tuning. In contrast, incorporating the RealMLP architectural changes (including their associated learning rate factors) improves scores on all benchmarks by around 5% or more, although they alone do not match the results of TabR-S-D. However, the non-architectural changes in RealMLP-TD make an even larger difference.




Figure B.6: **Benchmark scores of selected methods on $\mathcal{B}_{\text{class}}^{\text{train}}$, $\mathcal{B}_{\text{reg}}^{\text{train}}$, $\mathcal{B}_{\text{class}}^{\text{test}}$ and $\mathcal{B}_{\text{reg}}^{\text{test}}$ vs. average training time.** The y -axis shows the shifted geometric mean (SGM_ϵ) classification error (left) or RMSE (right) as explained in Section 2.2. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit $\#\text{splits} \rightarrow \infty$, see Appendix C.5.

B.8 Comparing HPO Methods

In Figure B.6, we compare two different HPO methods for GBDTs:

- Random search (HPO), as used in the main paper, with 50 steps.
- Tree parzen estimator (HPO-TPE) as implemented in hyperopt Bergstra et al. [2], with 50 steps. The first 20 of these steps use random search.

While TPE often performs slightly better, the differences in benchmark scores are relatively small.




Figure B.7: **Benchmark scores (arithmetic mean) vs. average training time.** The y -axis shows the *arithmetic mean* classification error (left) or RMSE (right). The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit $\#\text{splits} \rightarrow \infty$, see Appendix C.5.

B.9 More Time-Error Plots

Here, we provide more time-vs-error plots. Figure B.7 shows results for the arithmetic mean error, and Figure B.8 shows results for the arithmetic mean rank, which is also shown in Figure B.9 for the Grinsztajn et al. [19] benchmark. In Figure B.10 we also reproduce the main plots from Grinsztajn et al. [19] with added methods.

B.10 Critical Difference Diagrams

Figure B.11 analyzes the external validity of differences in average ranks between methods, i.e., whether they will generalize to new datasets from a distribution. While establishing external validity requires a large number of datasets, our benchmarks consistently show the improvements of RealMLP-TD over MLP-D to be externally valid.

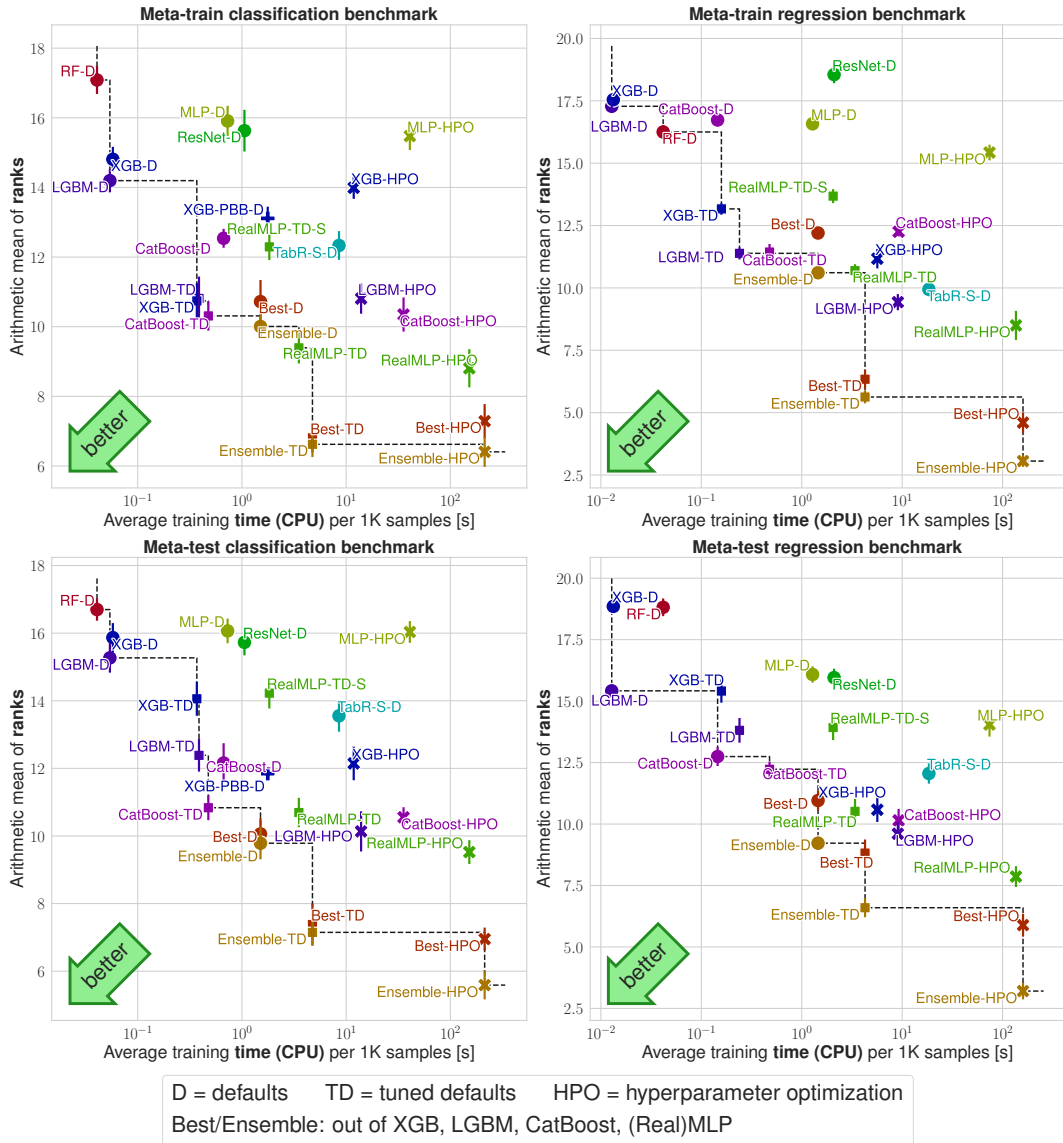


Figure B.8: **Benchmark scores (ranks) vs. average training time.** The y -axis shows the *arithmetic mean* rank, averaged over all splits and datasets. The x -axis shows average training times per 1000 samples (measured on $\mathcal{B}^{\text{train}}$ for efficiency reasons), see Appendix C.6. The error bars are approximate 95% confidence intervals for the limit $\#\text{splits} \rightarrow \infty$, see Appendix C.5.

B.11 Win-rate Plots

For pairs of methods, we analyze the percentage of (dataset, split) combinations on which the first method has a lower error than the second method. We plot these win-rates in matrix plots: Figure B.12 shows the results on $\mathcal{B}_{\text{class}}^{\text{train}}$, Figure B.13 shows the results on $\mathcal{B}_{\text{class}}^{\text{test}}$, Figure B.14 shows the results on $\mathcal{B}_{\text{reg}}^{\text{train}}$, and Figure B.15 shows the results on $\mathcal{B}_{\text{reg}}^{\text{test}}$.




Figure B.9: **Results on the benchmarks of Grinsztajn et al. [19].** The y -axis shows the arithmetic mean rank of each model. The x -axis shows average training times per 1000 samples, see Appendix C.6.




Figure B.10: **Results on the benchmarks of Grinsztajn et al. [19], for classification (left) and regression (right).** The plot is similar to the one in the main part of Grinsztajn et al. [19], with our algorithms added. The y -axis shows the result of the best (on val, but evaluated on test) hyperparameter combination up to n steps of random step (x -axis). As in the original paper, we normalize each score between the max and the 10% quantile (classification) or 50% (regression), and truncate scores below 0 for regression.




Figure B.11: **Critical difference diagrams on the meta-train and meta-test benchmarks.** The plots show the average rank of methods on each benchmark. Horizontal bars indicate groups of algorithms that are not statistically significantly different at a 95% confidence level according to a Friedman test and post-hoc Nemenyi test implemented in `autorank` [23].




Figure B.12: Percentages of wins of row algorithms vs column algorithms on $\mathcal{B}_{\text{class}}^{\text{train}}$. Wins are averaged over all datasets and splits. Ties count as half-wins. Methods are sorted by average winrate (i.e., the average of the values in the row). When averaging, we use dataset-dependent weighting as explained in Section C.3.1.




Figure B.13: Percentages of wins of row algorithms vs column algorithms on $\mathcal{B}_{\text{class}}^{\text{test}}$. Wins are averaged over all datasets and splits. Ties count as half-wins. Methods are sorted by average winrate (i.e., the average of the values in the row).




Figure B.14: Percentages of wins of row algorithms vs column algorithms on $\mathcal{B}_{\text{reg}}^{\text{train}}$. Wins are averaged over all datasets and splits. Ties count as half-wins. Methods are sorted by average winrate (i.e., the average of the values in the row). When averaging, we use dataset-dependent weighting as explained in Section C.3.1.




Figure B.15: Percentages of wins of row algorithms vs column algorithms on $\mathcal{B}_{\text{reg}}^{\text{test}}$. Wins are averaged over all datasets and splits. Ties count as half-wins. Methods are sorted by average winrate (i.e., the average of the values in the row).

C Benchmark Details

C.1 Default Configurations

The parameters for RealMLP-TD and RealMLP-TD-S have already been given in Table A.1. Table C.1 shows the hyperparameters of LGBM-TD and LGBM-D. Table C.2 shows the hyperparameters of XGB-TD and XGB-D. Table C.3 shows the hyperparameters of CatBoost-TD and CatBoost-D. The parameters for LGBM-D, XGB-D, and CatBoost-D have been taken from the respective libraries at the time of writing and are given here for completeness. The parameters for MLP-D are given in Table C.4 and the parameters for ResNet-D are given in Table C.5. By “RTDL quantile transform”, we refer to the version adding noise before fitting the quantile transform. The parameters for TabR-S-D are given in Table C.6. On the *dionis* dataset, we activate context freeze for TabR-S-D after four epochs, since training can take up to 30 GPU-hours even with context freeze.

For XGB-PBB-D, we use the default parameters from Probst et al. [50], with the following modifications: We use `hist` gradient boosting since it is the new default in XGBoost 2.0. Moreover, since we have high-cardinality categories, we limit one-hot encoding to categories with less than 20 distinct values (not counting missing values), and use XGBoost’s native categorical feature handling for the remaining categorical features.

C.2 Hyperparameter Optimization

For LGBM-HPO, XGB-HPO, and CatBoost-HPO, we run 50 steps of random search using the search spaces presented in Table C.7, C.8, and C.9, respectively. These search spaces are adapted from the literature, using `n_estimators=1000` in each case.

For RealMLP-HPO, we run 50 steps of random search. We provide a custom search space specified in Table C.10. For MLP-HPO, we also run 50 steps of random search with the search space in Table C.11, adapted from [16].

C.3 Dataset Selection and Preprocessing

C.3.1 Meta-train Benchmarks

For the meta-train benchmarks, we adapt code from Steinwart [60] to collect all datasets from the UCI repository that follow certain criteria:

- Between 2,500 and 50,000 samples.
- Number of features at most 1,000.
- Labeled as classification or regression task.
- Description made it straightforward to convert the original dataset into a numeric .csv format.
- Uploaded before 2019-05-08.

Table C.1: **Hyperparameters for LGBM-TD and LGBM-D.** Italic hyperparameters have not been tuned.

Hyperparameter	LGBM-TD		LGBM-D
	classif.	reg.	
<code>num_leaves</code>	50	100	31
<code>learning_rate</code>	0.04	0.05	0.1
<code>subsample</code>	0.75	0.7	1.0
<code>colsample_bytree</code>	1.0	1.0	1.0
<code>min_data_in_leaf</code>	40	3	20
<code>min_sum_hessian_in_leaf</code>	1e-7	1e-7	1e-3
<i><code>n_estimators</code></i>	<i>1000</i>	<i>1000</i>	<i>100</i>
<i><code>bagging_freq</code></i>	<i>1</i>	<i>1</i>	<i>1</i>
<i><code>max_bin</code></i>	<i>255</i>	<i>255</i>	<i>255</i>
<i><code>early_stopping_rounds</code></i>	<i>300</i>	<i>300</i>	<i>1000</i>

Table C.2: **Hyperparameters for XGB-TD and XGB-D.** *Italic hyperparameters have not been tuned for XGB-TD.*

Hyperparameter	XGB-TD		XGB-D
	classif.	reg.	
max_depth	6	9	6
learning_rate	0.08	0.05	0.3
subsample	0.65	0.7	1.0
colsample_bytree	1.0	1.0	1.0
colsample_bylevel	0.9	1.0	1.0
min_child_weight	5e-6	2.0	1.0
lambda	0.0	0.0	1.0
<i>tree_method</i>	<i>hist</i>	<i>hist</i>	<i>hist</i>
<i>n_estimators</i>	<i>1000</i>	<i>1000</i>	<i>100</i>
<i>max_bin</i>	<i>256</i>	<i>256</i>	<i>256</i>
<i>early_stopping_rounds</i>	<i>300</i>	<i>300</i>	<i>1000</i>

Table C.3: **Hyperparameters for CatBoost-TD and CatBoost-D.** *Italic hyperparameters have not been tuned for CatBoost-TD.*

Hyperparameter	CatBoost-TD		CatBoost-D
	classif.	reg.	
boosting_type	Plain	Plain	Plain
bootstrap_type	Bernoulli	Bernoulli	Bayesian
max_depth	7	9	6
learning_rate	0.08	0.09	automatic
subsample	0.9	0.9	—
bagging_temperature	—	—	1.0
l2_leaf_reg	1e-5	1e-5	3.0
random_strength	0.8	0.0	1.0
one_hot_max_size	15	20	2
leaf_estimation_iterations	1	20	None
<i>n_estimators</i>	<i>1000</i>	<i>1000</i>	<i>1000</i>
<i>max_bin</i>	<i>254</i>	<i>254</i>	<i>256</i>
<i>od_wait</i>	<i>300</i>	<i>300</i>	<i>None</i>
<i>od_type</i>	<i>Iter</i>	<i>Iter</i>	<i>Iter</i>

We remove rows with missing values and keep only those datasets that still have at least 2,500 samples.³ Some datasets are labeled both as regression and classification datasets, in which case we use them for both. We standardize the targets for regression datasets, such that a constant predictor can achieve an RMSE of 1. Some datasets contain different versions (e.g., different target columns), in which case we use all of them. To avoid biasing the results towards one dataset, we compute benchmark scores using weights proportional to $1/\#\text{versions}$. In total, we obtain 71 classification datasets (including versions) out of 46 original datasets, and 47 regression datasets (including versions) out of 26 original datasets. Tables C.12 and C.13 summarize key characteristics of these datasets. We count datasets with the same prefix (before the first underscore) as being versions of the same dataset for weighting, except for the two “facebook” datasets in $\mathcal{B}_{\text{reg}}^{\text{train}}$, which we count as distinct. For regression, we standardize the targets to have mean zero and variance 1.

During earlier development of the MLP, the meta-train benchmark used to include an epileptic seizure recognition dataset, which has since been removed from the UCI repository, hence we do not report results on it.

³We noticed later that the ozone_level_1hr and ozone_level_8hr datasets contain less than 2,500 samples, but we decided to keep them since we already used them for tuning the hyperparameters.

Table C.4: Hyperparameters for MLP-D, adapted from McElfresh et al. [43].

Hyperparameter	Space
lr scheduler	None
n_layers	3
d_layers	[128, 256, 128]
Dropout prob.	0.1
lr	1e-3
Optimizer	AdamW
d_embedding	8
batch_size	128
max_epochs	1000
early stopping patience	20
Preprocessing	RTDL quantile transform
Activation function	ReLU
Initialization	PyTorch default
Weight decay	0.01

Table C.5: Hyperparameters for ResNet-D, adapted from McElfresh et al. [43].

Hyperparameter	Space
lr scheduler	None
Activation	ReLU
Normalization	BatchNorm
n_layers	2
d_layers	[128, 128]
d_hidden_factor	2
hidden_dropout	0.25
residual_dropout	0.1
lr	1e-3
weight_decay	0.01
Optimizer	AdamW
d_embedding	8
batch_size	128
max_epochs	1000
early stopping patience	20
Preprocessing	RTDL quantile transform

C.3.2 Meta-test Benchmarks

The meta-test benchmarks consist of datasets from the AutoML Benchmark [14] and additional regression datasets from the OpenML-CTR23 benchmark [13], obtained from OpenML [64].

We make the following modifications:

- We use `brazilian_houses` from OpenML-CTR23 and exclude `Brazilian_houses` from the AutoML regression benchmark, since the latter contains three additional features that should not be used for predicting the target.
- We use another version of the `sarcos` dataset where the original test set is not included, since the original test set consists of duplicates of training samples.
- We excluded the following datasets because versions of them were already contained in the meta-training set:
 - For classification: `kr-vs-kp`, `wilt`, `ozone-level-8hr`, `first-order-theorem-proving`, `GesturePhaseSegmentationProcessed`, `PhishingWebsites`, `wine-quality-white`, `nomao`, `bank-marketing`, `adult`
 - For regression: `wine_quality`, `abalone`, `OnlineNewsPopularity`, `Brazilian_houses`, `physicochemical_protein`, `naval_propulsion_plant`, `superconductivity`, `white_wine`, `red_wine`, `grid_stability`

Table C.6: Hyperparameters for TabR-S-D, taken from Gorishniy et al. [18].

Hyperparameter	Space
num_embeddings	None
d_main	265
context_dropout	0.38920071545944357
d_multiplier	2.0
encoder_n_blocks	0
predictor_n_blocks	1
mixer_normalization	auto
dropout0	0.38852797479169876
dropout1	0.0
normalization	LayerNorm
activation	ReLU
batch_size	auto
patience	16
n_epochs	100,000
context_size	96
optimizer	AdamW
lr	0.0003121273641315169
weight_decay	1.2260352006404615e-06
Preprocessing	RTDL quantile transform

Table C.7: Hyperparameter search space for LGBM-HPO, adapted from Prokhorenkova et al. [51] with 1000 estimators instead of 5000.

Hyperparameter	Space
n_estimators	1000
bagging_freq	1
early_stopping_rounds	300
num_leaves	LogUniformInt[1, e^7]
learning_rate	LogUniform[e^{-7} , 1]
subsample	Uniform[0.5, 1]
feature_fraction	Uniform[0.5, 1]
min_data_in_leaf	LogUniformInt[1, e^6]
min_sum_hessian_in_leaf	LogUniform[e^{-16} , e^5]
lambda_l1	Random{0, LogUniform[e^{-16} , e^2]}
lambda_l2	Random{0, LogUniform[e^{-16} , e^2]}

We preprocess the datasets as follows:

- We remove rows with missing continuous values
- We subsample large datasets to contain at most 500,000 samples. Since the dionis dataset was particularly slow to train with GBDT models due to its 355 classes, we subsampled it to 100,000 samples.
- We standardize the targets for regression datasets, such that a constant predictor can achieve an RMSE of 1.
- We encode missing categorical values as a separate category.
- For regression, we standardize the targets to have mean zero and variance 1.

After preprocessing, we

- exclude datasets with less than 1,000 samples, these were
 - for classification: albert, APSFailure, arcene, Australian, blood-transfusion-service-center, eucalyptus, KDDCup09_appetency, KDDCup09-Upselling, micro-mass, vehicle
 - for regression: boston, cars, colleges, energy_efficiency, forest_fires, Moneyball, QSAR_fish_toxicity, sensory, student_performance_por, tecator, us_crime

Table C.8: Hyperparameter search space for XGB-HPO, adapted from Grinsztajn et al. [19]. We use the `hist` method, which is the new default in XGBoost 2.0 and supports native handling of categorical values, while the old `auto` method selection is not available in XGBoost 2.0. We also increase `early_stopping_rounds` to 300.

Hyperparameter	Space
<code>tree_method</code>	<code>hist</code>
<code>n_estimators</code>	1000
<code>early_stopping_rounds</code>	300
<code>max_depth</code>	UniformInt[1, 11]
<code>learning_rate</code>	LogUniform[1e-5, 0.7]
<code>subsample</code>	Uniform[0.5, 1]
<code>colsample_bytree</code>	Uniform[0.5, 1]
<code>colsample_bylevel</code>	Uniform[0.5, 1]
<code>min_child_weight</code>	LogUniformInt[1, 100]
<code>alpha</code>	LogUniform[1e-8, 1e-2]
<code>lambda</code>	LogUniform[1, 4]
<code>gamma</code>	LogUniform[1e-8, 7.0]

Table C.9: Hyperparameter search space for CatBoost-HPO, adapted from Shwartz-Ziv and Armon [57], who did not specify the number of estimators.

Hyperparameter	Space
<code>boosting_type</code>	Plain
<code>bootstrap_type</code>	Bayesian
<code>n_estimators</code>	1000
<code>max_depth</code>	6
<code>od_wait</code>	300
<code>od_type</code>	Iter
<code>learning_rate</code>	LogUniform[e^{-5} , 1]
<code>bagging_temperature</code>	Uniform[0, 1]
<code>l2_leaf_reg</code>	LogUniform[0, 10]
<code>random_strength</code>	UniformInt[1, 20]
<code>one_hot_max_size</code>	UniformInt[0, 25]
<code>leaf_estimation_iterations</code>	UniformInt[1, 20]

- exclude datasets that have more than 10,000 features after one-hot encoding. These were `Amazon_employee_access`, `Click_prediction_small`, and `sf-police-incidents` (all classification).

Table C.10: Hyperparameter search space for RealMLP-HPO. The remaining hyperparameters are set as in RealMLP-TD.

Hyperparameter	classif.	reg.
Num. embedding type	Choice([None, PBLD, PL, PLR])	same
Use scaling layer	Choice([True, False], p=[0.6, 0.4])	same
Learning rate	LogUniform([2e-2, 3e-1])	same
Dropout prob.	Choice([0.0, 0.15, 0.3], p=[0.3, 0.5, 0.2])	same
Activation fct.	Choice([ReLU, SELU, Mish])	same
Hidden layer sizes	Choice([[256, 256, 256], [64, 64, 64, 64, 64], [512]], p=[0.6, 0.2, 0.2])	same
Weight decay	Choice([0.0, 2e-2])	same
$w_{\text{emb}}^{(1,i)}$ init std.	LogUniform([0.05, 0.5])	
Label smoothing ϵ	Choice([0.0, 0.1], p=[0.3, 0.7])	no label smoothing

Table C.11: Hyperparameter search space for MLP-HPO, adapted from Gorishniy et al. [16]. We reduced the embedding dimension upper bound, the maximum number of epochs, and the number of layers to have a more acceptable runtime on the meta-test benchmarks. We also used a minimum batch size of 256. As in the original paper, the size of the first and the last layers are tuned and set separately, while the size for “in-between” layers is the same for all of them.

Hyperparameter	Space	
	#samples \leq 100,000	#samples $>$ 100,000
n_layers	UniformInt[1, 8]	UniformInt[1, 11]
d	UniformInt[64, 512]	UniformInt[64, 1024]
d_hidden_factor	UniformInt[1, 4]	
hidden_dropout	Uniform[0, 0.5]	
residual_dropout	Choice(0, Uniform[0, 0.5])	
lr	LogUniform[1e-5, 1e-2]	
weight decay	Choice(0, LogUniform[1e-6, 1e-3])	
d_embedding	UniformInt[8, 31]	
batch_size	256 if train_size < 30K else 512 if train_size < 100K else 1024	
lr_scheduler	None	
Optimizer	AdamW	
max #epochs	400	
early stopping patience	16	
Preprocessing	RTDL quantile transform	

C.4 Grinsztajn et al. benchmark

The datasets are taken from the benchmarks described in Grinsztajn et al. [19]. When a dataset is used both in benchmarks with and without categorical features, we use the version with categorical features. We preprocess the datasets following the same steps as in Grinsztajn et al. [19]:

- For neural networks, we quantile-transform the features to have a Gaussian distribution. For TabR [18], we use the modified quantile transform from the TabR paper. For RealMLP, we use the preprocessing described in Section 3, namely robust scaling and smooth clipping.
- For neural networks, we add as a hyperparameter the possibility to normalize the target variable for the model fit and transform it back for evaluation (via scikit-learn’s `TransformedTargetRegressor` and `StandardScaler`, which differs from the `QuantileTransformer` from the original paper, as we found it to work better). The same standardization is also applied to all default-parameter versions of neural networks.
- For models that do not handle categorical variables natively, we encode categorical features using `OneHotEncoder` from scikit-learn.
- Train size is restricted to 10,000 samples and test and validation size to 50,000 samples.

Note that the datasets from the original benchmark are already slightly preprocessed, e.g., heavy-tailed targets are standardized and missing values are removed. More details can be found in the original paper.

Results normalization For Figure 3, as in the original paper, we normalize the R2 or accuracy score for each dataset before averaging them. We use an affine normalization between 0 and 1, 1 corresponding to the score of the best model for each dataset, and 0 corresponding to the score of the worst model (for classification) and the 10th percentile of the scores (for regression). We use slightly different percentiles compared to the original paper as we normalize across the scores of the tuned and default models, and not all steps of the random search, which reduces the number of outliers. Other aggregation metrics are shown in Appendix B.9.

Other details We rerun classification results for neural networks compared to the original results to early stop on accuracy rather than on cross-entropy, to make results more comparable with the rest of this paper.

Table C.12: Datasets in the meta-train classification benchmark.

Name	#samples	#num. features	#cat. features	largest #categories	#classes
abalone	4177	8	0		3
adult	45222	7	7	41	2
anuran_calls_families	7127	22	0		3
anuran_calls_genus	6073	22	0		5
anuran_calls_species	5696	22	0		7
avila	20867	10	0		12
bank_marketing	41579	12	5	11	2
bank_marketing_additional	39457	19	3	11	2
chess	3196	1	31	3	2
chess_krvk	28056	3	3	8	18
crowd_sourced_mapping	10494	28	0		4
default_credit_card	30000	23	1	2	2
eeg_eye_state	14980	14	0		2
electrical_grid_stability_simulated	10000	12	0		2
facebook_live_sellers_thailand_status	6622	9	0		2
firm_teacher_clave	10800	0	16	2	4
first_order_theorem_proving	6118	51	0		2
gas_sensor_drift_class	13910	128	0		6
gesture_phase_segmentation_raw	9900	19	0		5
gesture_phase_segmentation_va3	9873	32	0		5
htru2	17898	8	0		2
human_activity_smartphone	10299	561	0		6
indoor_loc_building	21048	470	50	2	3
indoor_loc_relative	21048	470	50	2	3
insurance_benchmark	9822	80	4	5	2
landsat_satimage	6435	36	0		6
letter_recognition	20000	16	0		26
madelon	2600	500	0		2
magic_gamma_telescope	19020	10	0		2
mushroom	8124	0	21	12	2
musk	6598	166	0		2
nomao	34465	118	2	2	2
nursery	12960	7	1	2	4
occupancy_detection	20560	7	0		2
online_shoppers_attention	12330	16	2	3	2
optical_recognition_handwritten_digits	5620	59	3	2	10
ozone_level_1hr	1848	72	0		2
ozone_level_8hr	1847	72	0		2
page_blocks	5473	10	0		5
pen_recognition_handwritten_characters	10992	16	0		10
phishing	11055	8	22	2	2
polish_companies_bankruptcy_1year	7027	64	0		2
polish_companies_bankruptcy_2year	10173	64	0		2
polish_companies_bankruptcy_3year	10503	64	0		2
polish_companies_bankruptcy_4year	9792	64	0		2
polish_companies_bankruptcy_5year	5910	64	0		2
seismic_bumps	2584	12	3	2	2
skill_craft	3338	18	0		7
smartphone_human_activity	5744	561	0		6
smartphone_human_activity_postural	10411	561	0		6
spambase	4601	57	0		2
superconductivity_class	21263	81	0		2
thyroid_all_bp	3621	6	17	5	2
thyroid_all_hyper	3621	6	17	5	2
thyroid_all_hypo	3621	6	17	5	3
thyroid_all_rep	3621	6	17	5	2
thyroid_ann	7200	6	11	3	3
thyroid_dis	3621	6	17	5	2
thyroid_hypo	2700	7	14	3	2
thyroid_sick	3621	6	17	5	2
thyroid_sick_eu	3163	8	18	2	2
turkiye_student_evaluation	5820	32	0		3
wall_follow_robot_2	5456	2	0		4
wall_follow_robot_24	5456	24	0		4
wall_follow_robot_4	5456	4	0		4
waveform	5000	21	0		3
waveform_noise	5000	40	0		3
wilt	4839	5	0		2
wine_quality_all	6497	11	1	2	7
wine_quality_type	6497	11	0		2
wine_quality_white	4898	11	0		7

Table C.13: Datasets in the meta-train regression benchmark.

Name	#samples	#num. features	#cat. features	largest #categories
air_quality_bc	8991	10	0	
air_quality_co2	7674	10	0	
air_quality_no2	7715	10	0	
air_quality_nox	7718	10	0	
appliances_energy	19735	29	0	
beijing_pm25	41757	12	0	
bike_sharing_casual	17379	9	3	2
bike_sharing_total	17379	9	3	2
carbon_nanotubes_u	10721	5	0	
carbon_nanotubes_v	10721	5	0	
carbon_nanotubes_w	10721	5	0	
chess_krvk	28056	3	3	8
cycle_power_plant	9568	4	0	
electrical_grid_stability_simulated	10000	12	0	
facebook_comment_volume	40949	38	2	7
facebook_live_sellers_thailand_shares	7050	9	0	
five_cities_beijing_pm25	19062	14	0	
five_cities_chengdu_pm25	21074	14	0	
five_cities_guangzhou_pm25	20074	14	0	
five_cities_shanghai_pm25	21436	14	0	
five_cities_shenyang_pm25	19038	14	0	
gas_sensor_drift_class	13910	128	0	
gas_sensor_drift_conc	13910	128	0	
indoor_loc_alt	21048	470	50	2
indoor_loc_lat	21048	470	50	2
indoor_loc_long	21048	470	50	2
insurance_benchmark	9822	80	4	5
metro_interstate_traffic_volume_long	48204	6	2	38
metro_interstate_traffic_volume_short	48204	6	2	11
naval_propulsion_comp	11934	14	0	
naval_propulsion_turb	11934	14	0	
nursery	12960	7	1	2
online_news_popularity	39644	44	3	7
parking_birmingham	35717	5	0	
parkinson_motor	5875	18	1	2
parkinson_total	5875	18	1	2
protein_tertiary_structure	45730	9	0	
skill_craft	3338	18	0	
sml2010_dining	4137	17	0	
sml2010_room	4137	17	0	
superconductivity	21263	81	0	
travel_review_ratings	5456	23	0	
wall_follow_robot_2	5456	2	0	
wall_follow_robot_24	5456	24	0	
wall_follow_robot_4	5456	4	0	
wine_quality_all	6497	11	1	2
wine_quality_white	4898	11	0	

C.5 Confidence Intervals

Here, we specify how our confidence intervals are computed. Let X_{ij} denote the score (error/rank) of a method on dataset i and split j , with $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$. Then, the benchmark score \mathcal{S} can be written as

$$\mathcal{S} = g \left(\sum_{i=1}^n \frac{w_i}{m} \sum_{j=1}^m f(X_{ij}) \right), \quad (1)$$

where $f = g = \text{id}$ for the arithmetic mean. For the shifted geometric mean, we instead have $g = \exp$ and $f(x) = \log(x + \varepsilon)$, $\varepsilon = 0.01$. We interpret the benchmark datasets as fixed, but the splits as random. For each dataset i , X_{i1}, \dots, X_{im} are i.i.d. random variables. We first take the dataset averages

$$Z_j := \sum_{i=1}^n w_i f(X_{ij}).$$

The random variables X_{1j}, \dots, X_{nj} are independent but not identically distributed. Still, for lack of a better option, we assume that the Z_j are normally distributed with unknown mean and variance. We know that the Z_j are i.i.d., hence we use the confidence intervals from the Student's t -distribution for normally distributed random variables with unknown mean and variance. This gives us a confidence interval $[a, b]$ for $\frac{1}{m} \sum_{j=1}^m Z_j$. Since g is increasing, we hence obtain a confidence interval $[g(a), g(b)]$ for $\mathcal{S} = g \left(\frac{1}{m} \sum_{j=1}^m Z_j \right)$.

Table C.14: Datasets in the meta-test classification benchmark.

Name	#samples	#num. features	#cat. features	largest #categories	#classes	OpenML task ID
Bioresponse	3751	1776	0		2	359967
Diabetes130US	101766	13	36	789	3	211986
Fashion-MNIST	70000	784	0		10	359976
Higgs	500000	28	0		2	360114
Internet-Advertisements	3279	3	1555	2	2	359966
KDDCup99	500000	32	9	65	21	360112
MiniBooNE	130064	50	0		2	359990
Satellite	5100	36	0		2	359975
ada	4147	48	0		2	190411
airlines	500000	3	4	293	2	189354
amazon-commerce-reviews	1500	10000	0		50	10090
car	1728	0	6	4	4	359960
christine	5418	1599	37		2	359973
churn	5000	16	4	10	2	359968
cmc	1473	2	7	4	3	359959
cnae-9	1080	856	0		9	359957
connect-4	67557	0	42	3	3	359977
covertype	500000	10	44	2	7	7593
credit-g	1000	7	13	10	2	168757
dilbert	10000	2000	0		5	168909
dionis	100000	60	0		355	189355
dna	3186	0	180	2	3	359964
fabert	8237	800	0		7	168910
gina	3153	970	0		2	189922
guillermo	2000	4296	0		2	359988
helena	65196	27	0		100	359984
jannis	83733	54	0		4	211979
jasmine	2984	8	136	2	2	168911
jungle_chess_2pcs_raw_endgame_complete	44819	6	0		3	359981
kc1	2109	21	0		2	359962
kick	72600	14	18	1054	2	359991
madeline	3140	259	0		2	190392
mfeat-factors	2000	216	0		10	359961
numera128.6	96320	21	0		2	167120
okcupid-stem	50788	2	17	7019	3	359993
pc4	1458	37	0		2	359958
philippine	5832	308	0		2	190410
phoneme	5404	5	0		2	168350
porto-seguro	453046	26	31	102	2	360113
qsar-biodeg	1055	41	0		2	359956
riccardo	20000	4296	0		2	359989
robert	10000	7200	0		10	359986
segment	2310	16	0		7	359963
shuttle	58000	9	0		7	359987
steel-plates-fault	1941	27	0		7	168784
sylvine	5124	20	0		2	359972
volkert	58310	180	0		10	359985
yeast	1484	8	0		10	2073

Comparison of two methods We often compute the error increase in % in the benchmark score of method A compared to method B with the shifted geometric mean, given by

$$100 \cdot \left(\frac{\mathcal{S}^{(A)}}{\mathcal{S}^{(B)}} - 1 \right).$$

Here, we leverage that the shifted geometric mean uses $g = \exp$ to write

$$\frac{\mathcal{S}^{(A)}}{\mathcal{S}^{(B)}} = g \left(\sum_{i=1}^n \frac{w_i}{m} \sum_{j=1}^m (f(X_{ij}^{(A)}) - f(X_{ij}^{(B)})) \right),$$

which is of the same form as Eq. (1). Hence, we obtain confidence intervals for this quantity using the same method.

C.6 Time Measurements

For our meta-train and meta-test benchmarks, we report training times measured as follows: We run all methods on a single compute node with a 32-core AMD Ryzen Threadripper Pro 3975 WX CPU, using 32 threads for GBDTs and the PyTorch default settings for MLPs. No method is run on GPUs. We run methods sequentially on one split on each dataset of the meta-train-class and meta-train-reg benchmarks. For RealMLP-HPO, we only run two random search steps and extrapolate the runtime to 50 steps. Runtimes for combinations of models (Best and Ensemble) are computed as the sum of the individual runtimes. We compute the runtime per 1000 samples on each dataset and then average them. For simplicity, we do not use the dataset-dependent weighting employed otherwise on the

Table C.15: Datasets in the meta-test regression benchmark.

Name	#samples	#num. features	#cat. features	largest #categories	OpenML task ID
Airlines_DepDelay_10M	500000	6	3	359	359929
Allstate_Claims_Severity	188318	14	116	326	233212
Buzzinsocialmedia_Twitter	500000	77	0		233213
MIP-2016-regression	1090	143	1	5	360945
Mercedes_Benz_Greener_Manufacturing	4209	368	8	47	233215
QSAR-TID-10980	5766	1024	0		360933
QSAR-TID-11	5742	1024	0		360932
SAT11-HAND-runtime-regression	1725	115	1	15	359948
Santander_transaction_value	4459	4991	0		233214
Yolanda	400000	100	0		317614
airfoil_self_noise	1503	5	0		361235
auction_verification	2043	5	2	6	361236
black_friday	166821	5	4	7	359937
brazilian_houses	10692	5	4	35	361267
california_housing	20640	8	0		361255
concrete_compressive_strength	1030	8	0		361237
cps88wages	28155	2	4	4	361261
cpu_activity	8192	21	0		361256
diamonds	53940	6	3	8	361257
elevators	16599	18	0		359936
fifa	19178	27	1	163	361272
fps_benchmark	2592	29	14	24	361268
geographical_origin_of_music	1059	116	0		361243
health_insurance	22272	4	7	6	361269
house_16H	22784	16	0		359952
house_prices_nominal	1121	36	43	25	359951
house_sales	21613	20	1	70	359949
kin8nm	8192	8	0		361258
kings_county	21613	17	4	70	361266
miami_housing	13932	15	0		361260
nyc-taxi-green-dec-2016	500000	9	9	259	359943
pol	15000	48	0		359946
pumadyn32nh	8192	32	0		361259
quake	2178	3	0		359930
sarcos	44484	21	0		361011
socmob	1156	1	4	17	361264
solar_flare	1066	2	8	6	361244
space_ga	3107	6	0		361623
topo_2_1	8885	266	0		359939
video_transcoding	68784	16	2	4	361252
wave_energy	72000	48	0		361253
yprop_4_1	8885	251	0		359940

meta-train benchmark. For the Grinsztajn et al. [19] benchmark, we follow the original paper and run neural networks on a GPU and the other models on 1 core of AMD EPYC 7742 64-Core Processor, and we average the time across all random steps (for each random step, the time is averaged across splits). To compute the runtime of neural networks, we restrict ourselves to steps ran on the same GPU model (NVIDIA A100-40GB), which means that we exclude datasets for which we have less than 15 steps of each model on this GPU (leaving us with 11 datasets for classification and 15 for regression). We then compute the average runtime per 1000 samples on each dataset and average them.

C.7 Compute Resources

While we did not measure compute resources precisely, our experiments required at least around 2000 hours on RTX 3090 GPUs and other GPUs, as well as roughly 10,000 hours on HPC CPU nodes (32–64 cores).

C.8 Used Libraries

Our implementation uses various libraries, out of which we would like to particularly acknowledge PyTorch [47], Scikit-learn [48], Ray [46], XGBoost [10], LightGBM [31], and CatBoost [51]. For using XGBoost, LightGBM, and CatBoost, we adapted wrapping code from the CatBoost quality benchmarks [51]. For the Grinsztajn et al. [19] benchmark, we adapt the code from the original paper, which is available at the new branch *better_by_default* in the repository https://github.com/LeoGrin/tabular-benchmark/tree/better_by_default.




Figure D.1: **Best-TD vs CatBoost-HPO on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).

D Results for Individual Datasets

Here, we provide and compare the results of central methods per dataset. First, Figures D.1 – D.7 show scatterplot comparisons for different models.

Table D.1 and Table D.2 show results on $\mathcal{B}_{\text{class}}^{\text{train}}$. Table D.3 and Table D.4 show results on $\mathcal{B}_{\text{reg}}^{\text{train}}$. Table D.5 and Table D.6 show results on $\mathcal{B}_{\text{class}}^{\text{test}}$. Table D.7 and Table D.7 show results on $\mathcal{B}_{\text{reg}}^{\text{test}}$.

Table D.1: Classification error of *untuned* methods on datasets in $\mathcal{B}_{\text{class}}^{\text{train}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-TD	TabR-S-D	MLP-D	CatBoost-TD	LGBM-TD	XGB-TD	RF-D
abalone	<u>0.447</u> ±0.014	0.440 ±0.010	<u>0.448</u> ±0.014	0.458±0.009	0.455±0.012	<u>0.451</u> ±0.013	0.457±0.010
adult	0.140±0.004	0.142±0.004	0.144±0.003	0.130 ±0.003	<u>0.131</u> ±0.003	<u>0.131</u> ±0.003	0.146±0.003
anuran_calls_families	0.006 ±0.001	0.008±0.002	0.009±0.002	<u>0.007</u> ±0.002	0.009±0.002	<u>0.008</u> ±0.003	0.012±0.003
anuran_calls_genus	0.007 ±0.002	<u>0.008</u> ±0.002	0.010±0.002	0.008±0.002	0.009±0.002	0.009±0.002	0.012±0.003
anuran_calls_species	0.006 ±0.001	0.010±0.002	0.009±0.002	<u>0.007</u> ±0.001	<u>0.008</u> ±0.002	0.008±0.002	0.010±0.001
avila	0.000 ±0.000	0.001±0.000	0.016±0.002	<u>0.001</u> ±0.000	0.001±0.001	0.001±0.000	0.011±0.001
bank_marketing	<u>0.089</u> ±0.002	<u>0.088</u> ±0.002	0.091±0.001	0.088 ±0.002	<u>0.090</u> ±0.002	0.090±0.002	0.091±0.002
bank_marketing_additional	<u>0.085</u> ±0.003	<u>0.086</u> ±0.003	0.086±0.002	0.084±0.003	0.084 ±0.002	<u>0.085</u> ±0.002	0.086±0.002
chess	<u>0.005</u> ±0.002	0.015±0.003	<u>0.008</u> ±0.003	0.008±0.003	0.011±0.002	0.005 ±0.001	0.015±0.003
chess_krvk	0.081 ±0.004	0.128±0.004	0.141±0.009	0.153±0.002	0.147±0.003	0.146±0.003	0.293±0.017
crowd_sourced_mapping	0.032±0.004	0.028 ±0.001	0.034±0.004	0.035±0.003	0.031±0.003	0.033±0.003	0.058±0.004
default_credit_card	0.179±0.004	0.182±0.003	<u>0.181</u> ±0.004	0.179±0.004	0.178 ±0.004	<u>0.180</u> ±0.004	0.183±0.003
eeg_eye_state	0.016 ±0.002	0.107±0.024	0.120±0.017	0.051±0.002	0.052±0.001	0.051±0.002	0.083±0.003
electrical_grid_stability_simulated	0.032 ±0.004	0.048±0.004	0.057±0.003	0.051±0.004	0.053±0.004	0.057±0.002	0.086±0.005
facebook_live_sellers_thailand_status	<u>0.134</u> ±0.007	0.139±0.007	0.139±0.007	0.131 ±0.005	<u>0.135</u> ±0.006	<u>0.133</u> ±0.008	0.141±0.004
firm_teacher_clave	0.129 ±0.006	<u>0.130</u> ±0.005	0.134±0.005	0.150±0.006	0.149±0.003	0.149±0.005	0.191±0.006
first_order_theorem_proving	0.179±0.006	0.182±0.007	0.181±0.006	0.158 ±0.008	<u>0.160</u> ±0.008	<u>0.160</u> ±0.006	<u>0.162</u> ±0.007
gas_sensor_drift_class	0.005±0.001	0.004 ±0.001	<u>0.005</u> ±0.001	0.006±0.001	0.006±0.001	0.006±0.001	0.008±0.001
gesture_phase_segmentation_raw	0.087±0.006	0.079±0.006	0.106±0.007	0.071±0.005	<u>0.067</u> ±0.005	0.066 ±0.004	0.074±0.005
gesture_phase_segmentation_va3	0.323±0.007	0.293 ±0.006	0.368±0.008	0.321±0.007	0.312±0.006	0.316±0.009	0.355±0.006
htru2	<u>0.020</u> ±0.002	<u>0.020</u> ±0.001	<u>0.020</u> ±0.002	<u>0.021</u> ±0.002	0.019 ±0.002	<u>0.020</u> ±0.002	<u>0.020</u> ±0.002
human_activity_smartphone	0.008±0.001	0.010±0.001	0.015±0.002	0.008±0.002	0.008±0.002	0.008 ±0.001	0.024±0.002
indoor_loc_building	<u>0.002</u> ±0.000	0.002±0.000	<u>0.002</u> ±0.000	0.002±0.000	0.003±0.001	<u>0.002</u> ±0.000	0.002 ±0.000
indoor_loc_relative	0.070±0.002	0.090±0.002	0.099±0.004	0.077±0.003	0.059 ±0.003	0.066±0.007	0.077±0.003
insurance_benchmark	0.061±0.005	<u>0.060</u> ±0.004	<u>0.060</u> ±0.004	0.061±0.005	0.059 ±0.004	<u>0.061</u> ±0.004	0.072±0.004
landsat_satimage	0.077 ±0.006	0.094±0.006	0.090±0.005	0.087±0.005	<u>0.080</u> ±0.004	<u>0.080</u> ±0.004	0.090±0.004
letter_recognition	0.019 ±0.001	<u>0.020</u> ±0.001	0.030±0.002	0.034±0.003	0.032±0.002	0.033±0.002	0.045±0.002
madelon	0.340±0.016	0.443±0.011	0.435±0.019	0.139 ±0.010	0.211±0.010	0.210±0.014	0.320±0.011
magic_gamma_telescope	0.115±0.005	0.109 ±0.004	0.122±0.003	0.116±0.004	0.118±0.004	0.118±0.003	0.122±0.003
mushroom	0.000 ±0.000	0.000 ±0.000	<u>0.000</u> ±0.000	0.000 ±0.000	<u>0.000</u> ±0.000	0.000 ±0.000	0.000 ±0.000
musk	0.003 ±0.002	<u>0.005</u> ±0.002	0.011±0.003	0.011±0.002	0.012±0.002	0.012±0.002	0.028±0.002
nomao	0.022±0.002	0.022±0.001	0.026±0.002	<u>0.018</u> ±0.001	0.017 ±0.001	<u>0.017</u> ±0.002	0.020±0.002
nursery	0.020±0.001	0.011 ±0.002	0.024±0.002	0.021±0.001	0.026±0.002	0.024±0.002	0.034±0.001
occupancy_detection	0.006 ±0.001	0.008±0.001	0.008±0.001	0.007±0.001	<u>0.007</u> ±0.001	0.007±0.000	<u>0.006</u> ±0.001
online_shoppers_attention	0.098±0.004	0.098±0.004	0.099±0.004	0.097±0.003	0.096 ±0.005	0.097±0.004	0.096±0.003
optical_recognition_handwritten_digits	0.011 ±0.003	0.015±0.003	0.018±0.003	0.017±0.003	0.015±0.003	0.016±0.003	<u>0.020</u> ±0.004
ozone_level_1hr	<u>0.036</u> ±0.007	<u>0.035</u> ±0.008	0.035 ±0.008	<u>0.036</u> ±0.009	<u>0.035</u> ±0.007	<u>0.035</u> ±0.007	<u>0.035</u> ±0.007
ozone_level_8hr	<u>0.071</u> ±0.010	<u>0.072</u> ±0.012	<u>0.072</u> ±0.010	<u>0.071</u> ±0.011	<u>0.070</u> ±0.012	0.067 ±0.011	<u>0.070</u> ±0.009
page_blocks	<u>0.028</u> ±0.004	<u>0.027</u> ±0.004	<u>0.028</u> ±0.004	0.025 ±0.003	<u>0.025</u> ±0.003	<u>0.025</u> ±0.004	<u>0.026</u> ±0.003
pen_recognition_handwritten_characters	0.004 ±0.001	0.007±0.001	0.008±0.002	0.007±0.001	0.007±0.001	0.006±0.001	0.011±0.001
phishing	<u>0.031</u> ±0.002	0.029 ±0.003	0.034±0.002	0.031±0.002	0.032±0.002	<u>0.031</u> ±0.002	0.032±0.003
polish_companies_bankruptcy_1year	0.018 ±0.003	0.026±0.003	0.028±0.002	0.021±0.002	0.022±0.003	0.021±0.002	0.027±0.003
polish_companies_bankruptcy_2year	0.017 ±0.002	0.041±0.002	0.041±0.002	0.025±0.003	0.025±0.003	0.026±0.003	0.035±0.002
polish_companies_bankruptcy_3year	0.023 ±0.002	0.041±0.004	0.046±0.003	0.031±0.003	0.033±0.003	0.033±0.003	0.040±0.003
polish_companies_bankruptcy_4year	0.029 ±0.003	0.051±0.002	0.053±0.002	0.031±0.002	0.036±0.002	0.035±0.002	0.048±0.002
polish_companies_bankruptcy_5year	0.037±0.002	0.064±0.004	0.066±0.002	0.031 ±0.004	<u>0.033</u> ±0.003	0.036±0.004	0.050±0.005
seismic_bumps	0.068±0.009	0.066±0.009	0.065 ±0.008	0.070±0.007	0.066±0.009	0.067±0.007	0.068±0.008
skill_craft	0.589 ±0.010	0.601±0.012	0.597±0.015	<u>0.600</u> ±0.014	0.607±0.013	0.610±0.014	<u>0.593</u> ±0.014
smartphone_human_activity	0.045±0.004	0.042±0.004	0.071±0.004	0.035 ±0.003	<u>0.035</u> ±0.005	<u>0.036</u> ±0.004	0.079±0.004
smartphone_human_activity_postural	0.009±0.002	0.010±0.001	0.013±0.002	0.006 ±0.001	0.007±0.001	0.007±0.001	0.025±0.002
spambase	<u>0.052</u> ±0.008	<u>0.052</u> ±0.007	<u>0.054</u> ±0.006	0.048±0.008	<u>0.049</u> ±0.008	0.048 ±0.008	<u>0.052</u> ±0.008
superconductivity_class	<u>0.058</u> ±0.003	<u>0.059</u> ±0.002	0.063±0.003	0.057 ±0.003	0.058±0.003	0.058±0.003	0.059±0.002
thyroid_all_bp	0.026±0.004	0.028±0.003	0.029±0.005	<u>0.024</u> ±0.004	0.025±0.003	0.022 ±0.003	0.027±0.003
thyroid_all_hyper	<u>0.015</u> ±0.002	<u>0.017</u> ±0.003	0.018±0.003	<u>0.014</u> ±0.003	<u>0.014</u> ±0.003	0.014 ±0.003	<u>0.015</u> ±0.003
thyroid_all_hypo	0.008±0.002	0.018±0.004	0.020±0.002	0.003 ±0.002	0.005±0.002	0.005±0.002	0.007±0.002
thyroid_all_rep	0.009±0.002	0.009±0.002	0.016±0.005	<u>0.005</u> ±0.002	<u>0.006</u> ±0.002	0.005 ±0.002	0.009±0.002
thyroid_ann	0.008±0.002	0.012±0.004	0.010±0.002	0.003 ±0.001	0.003±0.001	0.005±0.002	0.003±0.001
thyroid_dis	0.013±0.002	0.014±0.002	0.016±0.003	0.010 ±0.002	<u>0.011</u> ±0.002	<u>0.010</u> ±0.003	0.013±0.002
thyroid_hypo	0.014±0.003	0.014±0.003	0.016±0.003	0.009 ±0.004	<u>0.010</u> ±0.002	<u>0.012</u> ±0.003	0.012±0.003
thyroid_sick	0.016±0.003	0.022±0.004	0.029±0.003	0.011 ±0.003	0.018±0.004	<u>0.011</u> ±0.003	0.017±0.005
thyroid_sick_eu	0.000 ±0.000	<u>0.001</u> ±0.001	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000	0.000 ±0.000
turkiye_student_evaluation	0.016±0.002	<u>0.017</u> ±0.003	0.033±0.005	0.016±0.002	0.016 ±0.002	<u>0.017</u> ±0.003	0.113±0.003
wall_follow_robot_2	<u>0.002</u> ±0.001	0.009±0.002	0.008±0.002	0.002±0.001	0.004±0.002	0.002±0.001	0.001 ±0.001
wall_follow_robot_24	0.029±0.004	0.040±0.007	0.042±0.005	0.002 ±0.001	<u>0.005</u> ±0.002	0.004±0.001	0.007±0.001
wall_follow_robot_4	<u>0.002</u> ±0.001	0.012±0.003	0.011±0.003	0.002±0.001	0.004±0.002	0.003±0.001	0.001 ±0.001
waveform	0.141 ±0.004	0.147±0.008	0.145±0.006	0.146±0.005	0.149±0.006	0.147±0.005	0.148±0.005
waveform_noise	0.139 ±0.006	0.161±0.012	0.143±0.007	0.145±0.007	0.146±0.009	0.145±0.007	0.150±0.008
wilt	<u>0.012</u> ±0.003	<u>0.011</u> ±0.002	0.011 ±0.002	0.014±0.002	0.014±0.003	0.014±0.003	0.016±0.002
wine_quality_all	0.377±0.008	0.352±0.010	0.412±0.015	<u>0.346</u> ±0.008	0.351±0.007	0.349±0.007	0.338 ±0.009
wine_quality_type	<u>0.004</u> ±0.001	<u>0.004</u> ±0.001	<u>0.004</u> ±0.001	<u>0.004</u> ±0.001	0.005±0.001	0.003 ±0.001	0.006±0.001
wine_quality_white	0.371±0.009	0.345±0.008	0.405±0.018	<u>0.334</u> ±0.008	0.343±0.009	0.338±0.008	0.330 ±0.008

Table D.2: Classification error of *tuned* methods on datasets in $\mathcal{B}_{\text{class}}^{\text{train}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-HPO	MLP-HPO	CatBoost-HPO	LGBM-HPO	XGB-HPO
abalone	0.444±0.011	0.443 ±0.013	0.466±0.017	0.463±0.013	0.459±0.016
adult	0.139±0.003	0.146±0.004	<u>0.130</u> ±0.003	0.129 ±0.003	<u>0.131</u> ±0.003
anuran_calls_families	0.006 ±0.001	0.009±0.001	<u>0.008</u> ±0.002	0.009±0.003	0.011±0.002
anuran_calls_genus	0.006 ±0.002	0.011±0.002	0.009±0.002	0.008±0.002	0.012±0.002
anuran_calls_species	0.007 ±0.002	0.009±0.002	<u>0.008</u> ±0.001	<u>0.009</u> ±0.002	0.010±0.002
avila	0.001 ±0.000	0.016±0.005	<u>0.001</u> ±0.000	0.001±0.000	0.002±0.000
bank_marketing	<u>0.088</u> ±0.002	0.093±0.003	<u>0.088</u> ±0.002	0.088 ±0.002	0.090±0.002
bank_marketing_additional	<u>0.084</u> ±0.003	0.087±0.002	<u>0.083</u> ±0.002	0.083 ±0.002	<u>0.084</u> ±0.002
chess	<u>0.008</u> ±0.004	0.010±0.004	0.006 ±0.002	<u>0.007</u> ±0.003	0.011±0.003
chess_krvk	0.070 ±0.004	0.125±0.010	0.137±0.005	0.149±0.005	0.194±0.008
crowd_sourced_mapping	0.031 ±0.005	<u>0.035</u> ±0.004	<u>0.034</u> ±0.003	<u>0.034</u> ±0.004	0.039±0.003
default_credit_card	<u>0.180</u> ±0.004	<u>0.180</u> ±0.003	<u>0.179</u> ±0.003	0.179 ±0.004	<u>0.180</u> ±0.004
eeg_eye_state	0.015 ±0.001	0.120±0.011	0.054±0.003	0.050±0.003	0.065±0.006
electrical_grid_stability_simulated	0.029 ±0.003	0.054±0.004	0.048±0.004	0.054±0.003	0.057±0.005
facebook_live_sellers_thailand_status	<u>0.133</u> ±0.007	0.139±0.006	0.132 ±0.006	<u>0.136</u> ±0.006	<u>0.138</u> ±0.007
firm_teacher_clave	0.125 ±0.006	0.133±0.007	0.144±0.007	0.143±0.006	0.141±0.004
first_order_theorem_proving	0.182±0.009	0.180±0.007	<u>0.161</u> ±0.008	0.160 ±0.006	<u>0.164</u> ±0.009
gas_sensor_drift_class	<u>0.005</u> ±0.001	0.005 ±0.001	<u>0.006</u> ±0.001	<u>0.006</u> ±0.001	0.007±0.001
gesture_phase_segmentation_raw	0.086±0.004	0.093±0.007	0.071±0.004	0.066 ±0.004	<u>0.069</u> ±0.004
gesture_phase_segmentation_va3	0.332±0.009	0.373±0.015	0.323±0.008	0.307 ±0.005	0.331±0.010
htru2	<u>0.020</u> ±0.001	0.020 ±0.002	<u>0.020</u> ±0.002	<u>0.020</u> ±0.002	<u>0.020</u> ±0.002
human_activity_smartphone	<u>0.008</u> ±0.001	0.015±0.002	0.009±0.002	0.007 ±0.001	0.011±0.002
indoor_loc_building	<u>0.002</u> ±0.000	<u>0.002</u> ±0.000	0.002 ±0.000	<u>0.002</u> ±0.000	<u>0.002</u> ±0.000
indoor_loc_relative	0.060±0.004	0.097±0.004	0.070±0.003	0.056 ±0.002	0.063±0.004
insurance_benchmark	<u>0.061</u> ±0.004	0.060 ±0.004	0.062±0.005	<u>0.061</u> ±0.004	<u>0.061</u> ±0.004
landsat_satimage	0.079 ±0.005	0.090±0.004	0.084±0.004	<u>0.080</u> ±0.003	0.088±0.007
letter_recognition	0.018 ±0.001	0.030±0.002	0.033±0.002	0.034±0.002	0.044±0.002
madelon	0.194±0.019	0.434±0.019	0.158 ±0.015	0.176±0.009	0.182±0.009
magic_gamma_telescope	0.115 ±0.003	0.120±0.003	<u>0.117</u> ±0.004	<u>0.118</u> ±0.004	0.119±0.004
mushroom	0.000 ±0.000	<u>0.000</u> ±0.000	<u>0.000</u> ±0.000	0.000 ±0.000	0.000 ±0.000
musk	0.003 ±0.002	0.011±0.002	0.010±0.003	0.010±0.003	0.017±0.003
nomao	0.021±0.001	0.025±0.002	<u>0.018</u> ±0.001	0.017 ±0.001	<u>0.018</u> ±0.001
nursery	0.019 ±0.002	0.022±0.001	0.021±0.002	<u>0.021</u> ±0.002	0.023±0.003
occupancy_detection	<u>0.007</u> ±0.001	0.009±0.001	<u>0.007</u> ±0.001	0.006 ±0.001	<u>0.007</u> ±0.001
online_shoppers_attention	<u>0.098</u> ±0.006	<u>0.098</u> ±0.004	0.098 ±0.004	<u>0.098</u> ±0.005	<u>0.098</u> ±0.004
optical_recognition_handwritten_digits	0.010 ±0.002	0.017±0.003	0.016±0.003	0.015±0.004	<u>0.020</u> ±0.004
ozone_level_1hr	<u>0.035</u> ±0.008	<u>0.038</u> ±0.009	<u>0.035</u> ±0.008	<u>0.038</u> ±0.008	0.035 ±0.008
ozone_level_8hr	<u>0.071</u> ±0.011	0.068 ±0.011	<u>0.073</u> ±0.011	<u>0.077</u> ±0.014	<u>0.073</u> ±0.007
page_blocks	<u>0.025</u> ±0.003	<u>0.026</u> ±0.004	0.024 ±0.004	<u>0.026</u> ±0.003	<u>0.025</u> ±0.004
pen_recognition_handwritten_characters	0.004 ±0.001	0.007±0.002	0.006±0.001	0.006±0.001	0.009±0.001
phishing	<u>0.030</u> ±0.002	<u>0.033</u> ±0.002	<u>0.031</u> ±0.002	0.030 ±0.002	0.033±0.002
polish_companies_bankruptcy_1year	0.018 ±0.003	0.028±0.002	0.021±0.002	0.021±0.001	0.022±0.002
polish_companies_bankruptcy_2year	0.016 ±0.002	0.041±0.002	0.024±0.003	0.024±0.003	0.024±0.003
polish_companies_bankruptcy_3year	0.024 ±0.002	0.046±0.004	0.030±0.002	0.031±0.004	0.032±0.004
polish_companies_bankruptcy_4year	0.029 ±0.002	0.053±0.003	0.031±0.002	0.032±0.002	0.034±0.002
polish_companies_bankruptcy_5year	0.038±0.003	0.064±0.004	0.030 ±0.004	<u>0.033</u> ±0.004	0.036±0.004
seismic_bumps	0.071±0.010	0.065 ±0.009	0.069±0.009	0.071±0.009	0.070±0.008
skill_craft	0.584 ±0.012	0.607±0.022	<u>0.587</u> ±0.017	0.609±0.015	0.602±0.015
smartphone_human_activity	0.039 ±0.006	0.070±0.007	<u>0.040</u> ±0.003	<u>0.040</u> ±0.005	0.048±0.004
smartphone_human_activity_postural	<u>0.007</u> ±0.002	0.013±0.002	<u>0.007</u> ±0.001	0.006 ±0.001	0.011±0.002
spambase	0.054±0.007	<u>0.053</u> ±0.008	0.046 ±0.006	<u>0.051</u> ±0.008	0.055±0.006
superconductivity_class	0.059±0.003	0.061±0.003	<u>0.058</u> ±0.001	<u>0.058</u> ±0.003	0.057 ±0.002
thyroid_all_bp	0.024 ±0.004	0.030±0.004	<u>0.025</u> ±0.004	<u>0.025</u> ±0.004	<u>0.027</u> ±0.003
thyroid_all_hyper	0.014 ±0.002	0.018±0.003	<u>0.014</u> ±0.003	<u>0.015</u> ±0.002	<u>0.015</u> ±0.002
thyroid_all_hypo	0.007±0.002	0.021±0.003	0.004 ±0.002	<u>0.006</u> ±0.002	0.005±0.002
thyroid_all_rep	0.010±0.003	0.017±0.005	0.005 ±0.003	<u>0.007</u> ±0.003	<u>0.007</u> ±0.002
thyroid_ann	0.005±0.001	0.012±0.001	0.004±0.001	<u>0.003</u> ±0.001	0.002 ±0.001
thyroid_dis	<u>0.013</u> ±0.004	0.017±0.003	0.010 ±0.002	0.013±0.002	<u>0.013</u> ±0.003
thyroid_hypo	0.014±0.002	0.015±0.003	<u>0.011</u> ±0.003	<u>0.011</u> ±0.004	0.009 ±0.003
thyroid_sick	0.011 ±0.003	0.026±0.005	<u>0.013</u> ±0.003	<u>0.015</u> ±0.004	0.017±0.004
thyroid_sick_eu	0.000 ±0.000	<u>0.000</u> ±0.001	0.000 ±0.000	0.000 ±0.000	<u>0.000</u> ±0.000
turkiye_student_evaluation	0.016 ±0.002	0.022±0.003	<u>0.016</u> ±0.003	0.019±0.002	<u>0.017</u> ±0.003
wall_follow_robot_2	<u>0.002</u> ±0.002	0.005±0.002	0.002±0.001	0.003±0.001	0.001 ±0.001
wall_follow_robot_24	0.011±0.004	0.043±0.003	0.003 ±0.001	0.005±0.001	<u>0.004</u> ±0.002
wall_follow_robot_4	<u>0.002</u> ±0.001	0.009±0.003	<u>0.003</u> ±0.002	<u>0.003</u> ±0.001	0.002 ±0.001
waveform	0.136 ±0.005	0.141±0.004	0.143±0.003	0.152±0.007	0.148±0.005
waveform_noise	0.141±0.008	0.143±0.010	0.139 ±0.008	0.145±0.008	0.146±0.007
wilt	<u>0.013</u> ±0.003	0.011 ±0.002	<u>0.014</u> ±0.003	0.014±0.003	0.015±0.003
wine_quality_all	0.367±0.011	0.384±0.012	<u>0.351</u> ±0.009	0.344 ±0.009	<u>0.350</u> ±0.011
wine_quality_type	0.005±0.002	<u>0.004</u> ±0.002	<u>0.004</u> ±0.002	0.004 ±0.001	<u>0.006</u> ±0.002
wine_quality_white	0.367±0.008	0.374±0.014	<u>0.344</u> ±0.008	0.338 ±0.011	<u>0.352</u> ±0.016




Figure D.2: **Best-TD vs Best-HPO on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).




Figure D.3: **RealMLP-TD vs CatBoost-TD on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).




Figure D.4: **RealMLP-HPO vs CatBoost-HPO on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).




Figure D.5: **Ensemble-TD vs Best-TD on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).

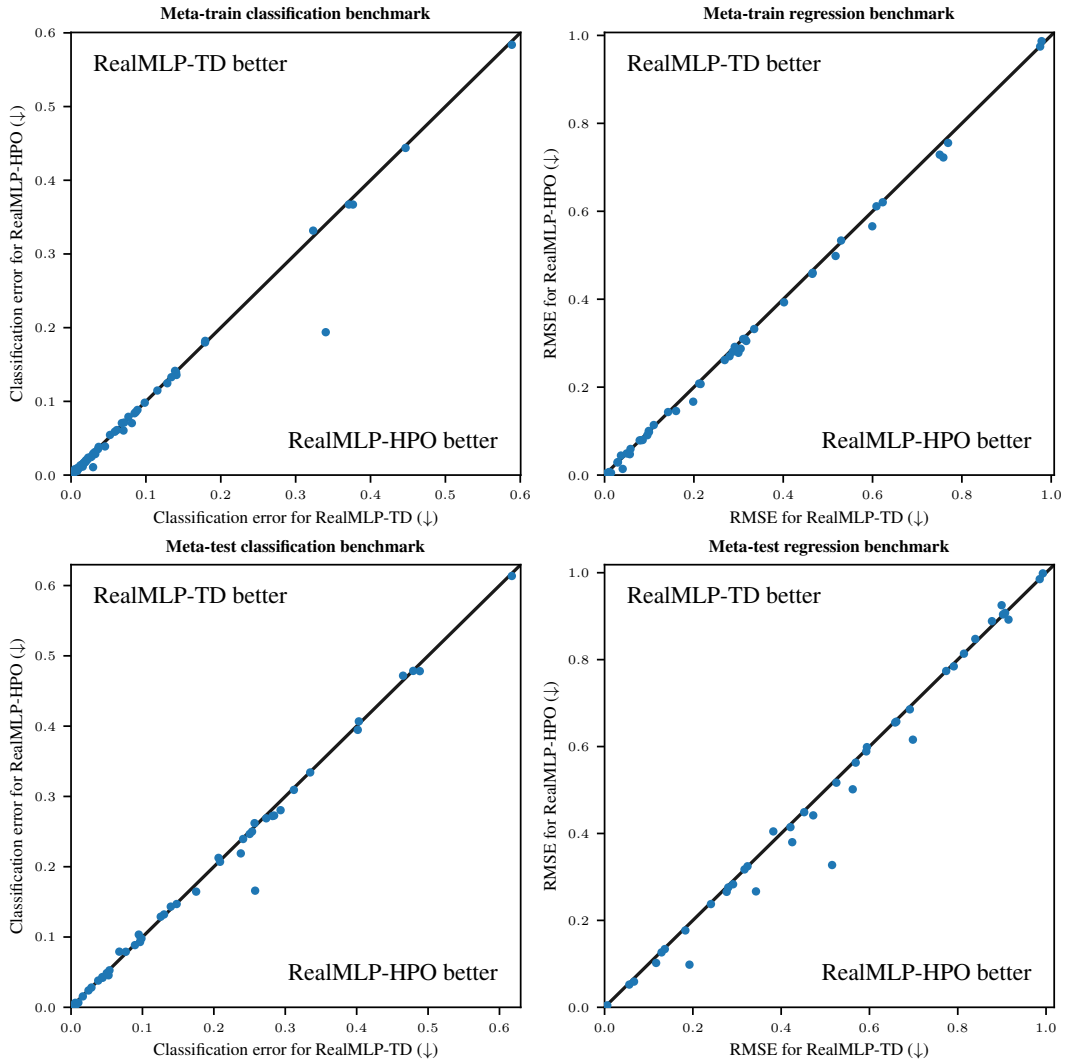


Figure D.6: **RealMLP-TD vs RealMLP-HPO on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).

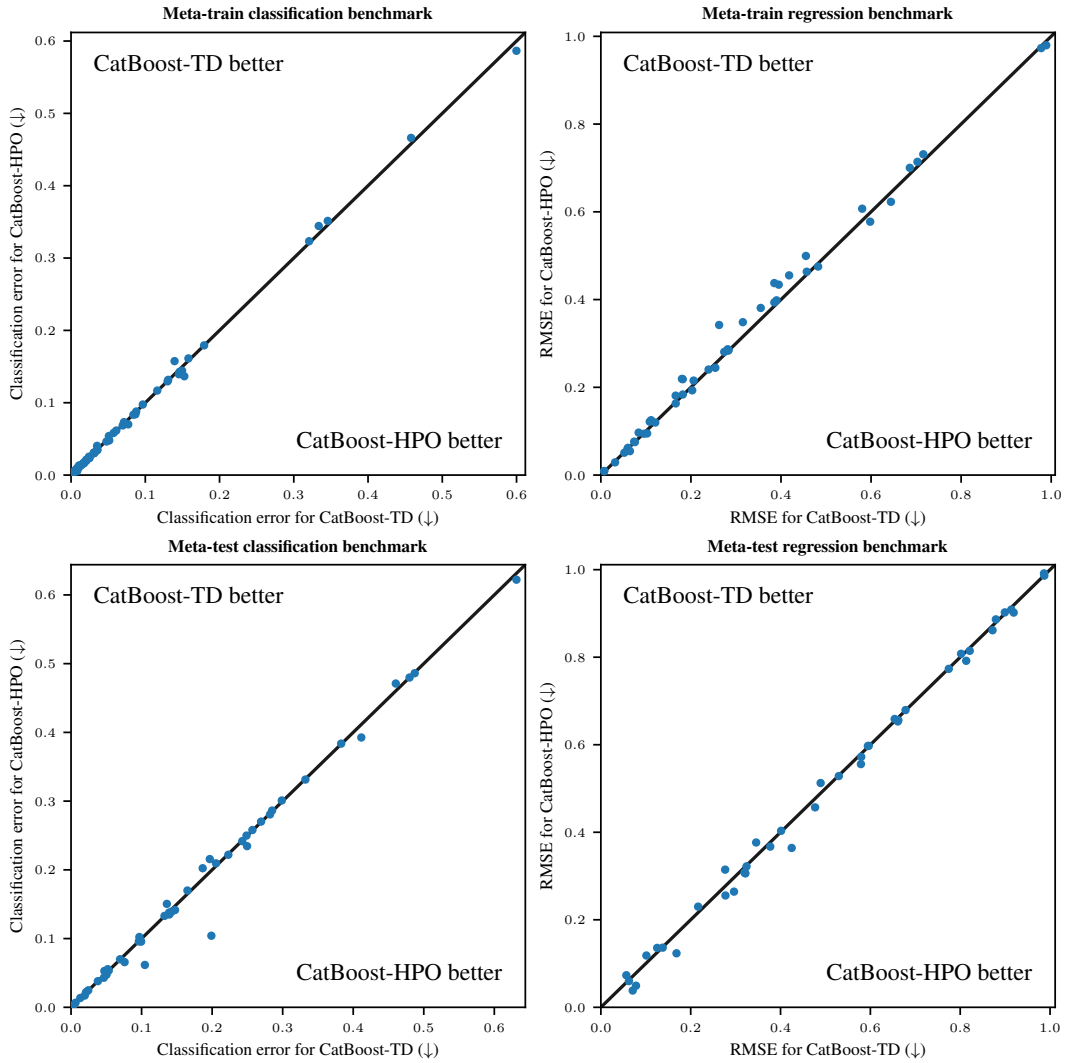


Figure D.7: **CatBoost-TD vs CatBoost-HPO on individual datasets.** Each point represents the errors of both models on a dataset, averaged across 10 train-valid-test splits. The black line represents equal errors ($x = y$).

Table D.3: RMSE of *untuned* methods on datasets in $\mathcal{B}_{\text{reg}}^{\text{train}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-TD	TabR-S-D	MLP-D	CatBoost-TD	LGBM-TD	XGB-TD	RF-D
air_quality_bc	0.005 ±0.000	0.025±0.004	0.043±0.003	0.032±0.007	0.030±0.005	0.033±0.007	0.013±0.004
air_quality_co2	0.305±0.019	<u>0.245</u> ±0.010	0.296±0.018	0.239 ±0.015	0.263±0.018	0.277±0.019	0.294±0.017
air_quality_no2	0.314±0.005	0.281 ±0.005	0.334±0.009	<u>0.284</u> ±0.007	0.294±0.011	0.293±0.008	0.322±0.009
air_quality_nox	0.287±0.016	<u>0.261</u> ±0.018	0.288±0.013	<u>0.254</u> ±0.020	0.252 ±0.018	<u>0.256</u> ±0.017	<u>0.257</u> ±0.011
appliances_energy	0.759±0.023	0.650 ±0.015	0.795±0.018	0.686±0.015	0.678±0.016	0.676±0.015	0.705±0.015
beijing_pm25	0.310±0.006	0.279 ±0.005	0.441±0.012	0.395±0.007	0.378±0.007	0.379±0.006	0.421±0.006
bike_sharing_casual	0.280±0.007	0.269 ±0.006	0.287±0.006	<u>0.274</u> ±0.005	0.276±0.006	0.282±0.006	0.304±0.007
bike_sharing_total	0.212±0.006	0.220±0.005	0.223±0.005	0.206 ±0.006	<u>0.210</u> ±0.006	0.214±0.006	0.240±0.007
carbon_nanotubes_u	0.010±0.000	0.013±0.001	0.026±0.003	0.007 ±0.000	0.009±0.000	0.010±0.000	0.011±0.000
carbon_nanotubes_v	0.010±0.000	0.014±0.001	0.028±0.005	0.007 ±0.000	0.009±0.000	0.010±0.000	0.011±0.000
carbon_nanotubes_w	0.050 ±0.013	<u>0.053</u> ±0.011	0.060±0.010	<u>0.052</u> ±0.012	<u>0.056</u> ±0.009	<u>0.058</u> ±0.009	0.060±0.007
chess_krvk	0.096 ±0.005	0.137±0.010	0.123±0.006	0.263±0.005	0.227±0.005	0.238±0.005	0.441±0.034
cycle_power_plant	0.215±0.006	0.170 ±0.005	0.223±0.003	0.182±0.004	0.184±0.004	0.188±0.005	0.201±0.004
electrical_grid_stability_simulated	0.142 ±0.003	0.177±0.003	0.186±0.004	0.203±0.005	0.216±0.002	0.250±0.003	0.330±0.004
facebook_comment_volume	<u>0.609</u> ±0.064	<u>0.622</u> ±0.055	0.625±0.040	<u>0.598</u> ±0.057	0.583 ±0.058	<u>0.588</u> ±0.056	<u>0.585</u> ±0.056
facebook_live_sellers_thailand_shares	0.530±0.059	<u>0.468</u> ±0.055	0.469±0.060	<u>0.457</u> ±0.056	0.451 ±0.044	<u>0.454</u> ±0.056	<u>0.463</u> ±0.062
five_cities_beijing_pm25	0.300±0.021	0.255 ±0.004	0.420±0.006	0.355±0.006	0.346±0.008	0.359±0.008	0.412±0.008
five_cities_chengdu_pm25	0.269±0.010	0.214 ±0.005	0.378±0.008	0.315±0.007	0.304±0.007	0.301±0.008	0.326±0.009
five_cities_guangzhou_pm25	0.402±0.016	0.331 ±0.014	0.519±0.013	0.455±0.011	0.454±0.013	0.458±0.014	0.489±0.012
five_cities_shanghai_pm25	0.317±0.015	0.253 ±0.008	0.445±0.012	0.385±0.008	0.385±0.009	0.397±0.011	0.449±0.012
five_cities_shenyang_pm25	0.335±0.020	0.299 ±0.014	0.524±0.027	0.418±0.019	0.423±0.020	0.434±0.021	0.523±0.019
gas_sensor_drift_class	<u>0.082</u> ±0.012	0.073 ±0.011	0.079±0.008	0.121±0.006	0.128±0.008	0.132±0.008	0.139±0.009
gas_sensor_drift_conc	0.160±0.014	0.145 ±0.011	0.169±0.020	0.166±0.016	0.170±0.016	0.171±0.014	0.171±0.015
indoor_loc_alt	0.099 ±0.004	0.182±0.004	0.188±0.005	0.166±0.003	0.148±0.002	0.162±0.003	0.171±0.003
indoor_loc_lat	0.079 ±0.004	0.109±0.004	0.112±0.004	0.109±0.004	0.097±0.004	0.108±0.004	0.104±0.004
indoor_loc_long	0.058 ±0.003	0.082±0.002	0.079±0.002	0.084±0.003	0.070±0.003	0.085±0.003	0.073±0.003
insurance_benchmark	0.975 ±0.031	<u>0.981</u> ±0.029	<u>0.980</u> ±0.030	<u>0.978</u> ±0.032	<u>0.983</u> ±0.031	<u>0.983</u> ±0.033	1.075±0.031
metro_interstate_traffic_volume_long	0.466±0.004	0.305 ±0.004	0.467±0.004	0.390±0.003	0.385±0.004	0.443±0.016	0.437±0.003
metro_interstate_traffic_volume_short	0.465±0.003	0.299 ±0.004	0.466±0.004	0.385±0.003	0.376±0.004	0.386±0.004	0.435±0.004
naval_propulsion_comp	0.014 ±0.003	0.029±0.005	0.059±0.003	0.060±0.002	0.064±0.002	0.064±0.003	0.079±0.005
naval_propulsion_turb	<u>0.041</u> ±0.027	0.039 ±0.003	0.086±0.005	0.096±0.005	0.097±0.005	0.096±0.004	0.114±0.005
nursery	0.085 ±0.003	<u>0.087</u> ±0.008	0.100±0.003	0.112±0.003	0.107±0.003	0.103±0.003	0.116±0.003
online_news_popularity	<u>0.979</u> ±0.118	<u>0.980</u> ±0.117	0.978 ±0.117	<u>0.989</u> ±0.116	<u>0.987</u> ±0.117	<u>0.988</u> ±0.116	<u>1.021</u> ±0.110
parking_birmingham	0.291±0.004	0.298±0.004	0.303±0.004	0.283 ±0.004	0.288±0.004	0.293±0.004	0.333±0.004
parkinson_motor	<u>0.099</u> ±0.010	0.095 ±0.008	0.407±0.012	0.182±0.011	0.168±0.010	0.163±0.006	0.194±0.011
parkinson_total	<u>0.110</u> ±0.010	0.105 ±0.010	0.423±0.024	0.180±0.010	0.163±0.010	0.158±0.010	0.181±0.008
protein_tertiary_structure	0.599±0.005	0.501 ±0.003	0.578±0.004	0.580±0.003	0.575±0.003	0.575±0.002	0.592±0.003
skill_craft	0.623 ±0.010	0.667±0.007	0.657±0.009	0.644±0.009	0.642±0.005	0.645±0.009	0.641±0.008
sml2010_dining	0.030 ±0.003	0.040±0.002	0.085±0.007	0.074±0.002	0.091±0.002	0.101±0.003	0.132±0.003
sml2010_room	0.029 ±0.002	0.041±0.003	0.081±0.004	0.076±0.003	0.089±0.004	0.098±0.004	0.129±0.005
superconductivity	0.293±0.008	0.294±0.008	0.300±0.006	0.281 ±0.005	<u>0.282</u> ±0.005	<u>0.281</u> ±0.005	0.287±0.005
travel_review_ratings	0.517±0.017	0.523±0.012	0.530±0.012	<u>0.482</u> ±0.013	0.480 ±0.014	<u>0.486</u> ±0.015	<u>0.485</u> ±0.011
wall_follow_robot_2	<u>0.037</u> ±0.010	0.109±0.014	0.090±0.018	0.059±0.021	0.069±0.026	0.054±0.026	0.027 ±0.020
wall_follow_robot_24	0.198±0.019	0.306±0.028	0.302±0.018	<u>0.103</u> ±0.014	0.090 ±0.021	<u>0.093</u> ±0.019	<u>0.095</u> ±0.019
wall_follow_robot_4	0.057±0.026	0.136±0.022	0.140±0.017	0.065±0.024	0.067±0.029	0.053±0.025	0.027 ±0.020
wine_quality_all	0.769±0.012	0.736±0.017	0.782±0.016	<u>0.716</u> ±0.018	0.714 ±0.018	<u>0.717</u> ±0.019	<u>0.722</u> ±0.019
wine_quality_white	0.750±0.021	0.720±0.015	0.766±0.012	<u>0.703</u> ±0.012	0.702 ±0.011	<u>0.703</u> ±0.011	<u>0.706</u> ±0.012

Table D.4: RMSE of *tuned* methods on datasets in $\mathcal{B}_{\text{reg}}^{\text{train}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-HPO	MLP-HPO	CatBoost-HPO	LGBM-HPO	XGB-HPO
air_quality_be	0.004 ±0.000	0.026±0.003	0.029±0.004	0.029±0.004	0.026±0.006
air_quality_co2	0.288±0.017	0.292±0.015	0.241 ±0.014	<u>0.246</u> ±0.016	<u>0.245</u> ±0.016
air_quality_no2	0.310±0.005	0.334±0.007	0.286 ±0.006	<u>0.291</u> ±0.009	<u>0.286</u> ±0.006
air_quality_nox	0.281±0.015	0.282±0.013	<u>0.245</u> ±0.016	<u>0.252</u> ±0.013	0.239 ±0.010
appliances_energy	0.722±0.024	0.785±0.018	0.700±0.013	0.673 ±0.017	<u>0.677</u> ±0.014
beijing_pm25	0.309 ±0.005	0.423±0.008	0.434±0.008	0.371±0.011	0.386±0.006
bike_sharing_casual	0.270 ±0.006	0.288±0.006	0.281±0.006	<u>0.278</u> ±0.008	0.282±0.006
bike_sharing_total	0.208 ±0.006	0.229±0.007	0.215±0.006	<u>0.212</u> ±0.005	0.214±0.004
carbon_nanotubes_u	0.007 ±0.001	0.012±0.000	0.009±0.001	0.009±0.001	0.015±0.002
carbon_nanotubes_v	0.007 ±0.001	0.012±0.000	0.009±0.000	0.009±0.001	0.014±0.001
carbon_nanotubes_w	0.049 ±0.014	<u>0.051</u> ±0.012	<u>0.051</u> ±0.012	<u>0.050</u> ±0.012	<u>0.051</u> ±0.012
chess_krvk	0.091 ±0.005	0.108±0.006	0.342±0.005	0.267±0.023	0.411±0.028
cycle_power_plant	0.207±0.005	0.215±0.004	<u>0.184</u> ±0.005	0.183 ±0.005	<u>0.186</u> ±0.007
electrical_grid_stability_simulated	0.143 ±0.004	0.185±0.003	0.193±0.003	0.208±0.003	0.225±0.005
facebook_comment_volume	0.611±0.054	0.618±0.049	0.577 ±0.054	<u>0.594</u> ±0.055	<u>0.585</u> ±0.061
facebook_live_sellers_thailand_shares	0.534±0.077	0.486±0.059	0.463±0.071	<u>0.452</u> ±0.048	0.438 ±0.049
five_cities_beijing_pm25	0.278 ±0.008	0.423±0.007	0.381±0.006	0.359±0.007	0.368±0.008
five_cities_chengdu_pm25	0.262 ±0.006	0.375±0.010	0.348±0.005	0.302±0.009	0.321±0.015
five_cities_guangzhou_pm25	0.393 ±0.014	0.521±0.011	0.499±0.011	0.459±0.015	0.475±0.015
five_cities_shanghai_pm25	0.305 ±0.011	0.414±0.012	0.438±0.009	0.396±0.012	0.397±0.017
five_cities_shenyang_pm25	0.332 ±0.023	0.513±0.031	0.455±0.016	0.430±0.020	0.445±0.018
gas_sensor_drift_class	<u>0.079</u> ±0.010	0.078 ±0.007	0.120±0.005	0.120±0.007	0.120±0.007
gas_sensor_drift_conc	0.146 ±0.015	<u>0.150</u> ±0.010	0.163±0.013	0.167±0.014	0.162±0.015
indoor_loc_alt	0.100 ±0.004	0.185±0.005	0.181±0.004	0.137±0.003	0.160±0.006
indoor_loc_lat	0.079 ±0.004	0.108±0.003	0.122±0.004	0.090±0.003	0.105±0.007
indoor_loc_long	0.059 ±0.004	0.075±0.003	0.097±0.003	0.068±0.003	0.083±0.006
insurance_benchmark	0.975±0.031	0.977±0.031	<u>0.973</u> ±0.030	0.970 ±0.031	<u>0.971</u> ±0.032
metro_interstate_traffic_volume_long	0.460±0.004	0.469±0.005	0.398±0.004	0.392 ±0.008	<u>0.392</u> ±0.008
metro_interstate_traffic_volume_short	0.458±0.005	0.467±0.004	0.394±0.004	<u>0.388</u> ±0.009	0.386 ±0.008
naval_propulsion_comp	0.005 ±0.001	0.037±0.002	0.062±0.002	0.058±0.001	0.062±0.004
naval_propulsion_turb	0.014 ±0.001	0.054±0.006	0.094±0.002	0.091±0.007	0.095±0.006
nursery	0.080 ±0.004	0.088±0.004	0.125±0.002	0.114±0.005	0.121±0.006
online_news_popularity	0.987±0.121	0.981±0.116	0.980±0.117	0.977 ±0.118	0.979±0.117
parking_birmingham	0.292±0.004	0.301±0.004	0.284±0.004	0.285±0.004	0.279 ±0.004
parkinson_motor	0.098 ±0.015	0.386±0.022	0.219±0.011	0.187±0.011	0.213±0.025
parkinson_total	0.114 ±0.013	0.395±0.041	0.219±0.010	0.177±0.011	0.207±0.022
protein_tertiary_structure	0.566 ±0.004	0.573±0.005	0.607±0.003	<u>0.567</u> ±0.004	0.589±0.009
skill_craft	0.620 ±0.009	0.665±0.011	0.623±0.008	<u>0.628</u> ±0.014	0.631±0.010
sml2010_dining	0.029 ±0.001	0.068±0.004	0.076±0.002	0.085±0.005	0.089±0.006
sml2010_room	0.029 ±0.002	0.068±0.005	0.075±0.002	0.083±0.004	0.087±0.006
superconductivity	0.289±0.008	0.297±0.007	0.287±0.007	0.278 ±0.008	<u>0.281</u> ±0.006
travel_review_ratings	0.498±0.014	0.538±0.019	0.475±0.011	<u>0.463</u> ±0.012	0.460 ±0.011
wall_follow_robot_2	0.044 ±0.024	0.090±0.015	<u>0.060</u> ±0.020	<u>0.066</u> ±0.025	0.210±0.007
wall_follow_robot_24	0.167±0.021	0.309±0.021	<u>0.095</u> ±0.016	<u>0.097</u> ±0.019	0.080 ±0.016
wall_follow_robot_4	<u>0.047</u> ±0.027	0.134±0.014	<u>0.055</u> ±0.019	<u>0.065</u> ±0.025	0.043 ±0.016
wine_quality_all	0.755±0.018	0.779±0.015	0.731±0.016	0.708 ±0.019	<u>0.711</u> ±0.016
wine_quality_white	0.729±0.014	0.767±0.013	0.714±0.010	0.697 ±0.011	<u>0.702</u> ±0.011

Table D.5: Classification error of *untuned* methods on datasets in $\mathcal{B}_{\text{class}}^{\text{test}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-TD	TabR-S-D	MLP-D	CatBoost-TD	LGBM-TD	XGB-TD	RF-D
ada	<u>0.148</u> ±0.012	0.151±0.009	0.149±0.011	0.139 ±0.008	0.141±0.011	<u>0.140</u> ±0.011	<u>0.144</u> ±0.007
airlines	0.335±0.001	<u>0.332</u> ±0.001	0.337±0.001	0.332 ±0.001	<u>0.333</u> ±0.001	0.337±0.001	0.382±0.001
amazon-commerce-reviews	<u>0.209</u> ±0.021	0.402±0.028	0.372±0.019	0.197 ±0.017	0.285±0.020	0.290±0.020	0.407±0.022
Bioreponse	0.238±0.010	0.228±0.013	0.232±0.006	<u>0.205</u> ±0.010	0.204 ±0.008	<u>0.211</u> ±0.007	<u>0.205</u> ±0.009
car	0.008 ±0.006	<u>0.011</u> ±0.006	<u>0.011</u> ±0.006	0.019±0.009	0.021±0.006	0.019±0.004	0.071±0.017
christine	0.293±0.012	0.289±0.007	0.284±0.015	<u>0.269</u> ±0.012	0.266 ±0.012	<u>0.273</u> ±0.015	0.281±0.013
churn	0.044 ±0.003	0.055±0.006	0.065±0.007	0.050±0.004	<u>0.048</u> ±0.005	0.048±0.004	0.064±0.005
cmc	0.465±0.019	<u>0.449</u> ±0.025	0.441 ±0.015	0.460±0.017	<u>0.457</u> ±0.017	0.467±0.018	0.472±0.014
cnac-9	0.068±0.010	0.066±0.011	0.053 ±0.011	0.076±0.010	0.308±0.017	0.091±0.012	0.087±0.012
connect-4	0.130 ±0.003	0.135±0.003	0.149±0.002	0.143±0.003	0.136±0.003	0.142±0.003	0.181±0.003
covertype	0.029 ±0.001	<u>0.029</u> ±0.001	0.069±0.002	0.105±0.000	0.058±0.001	0.072±0.000	0.055±0.001
credit-g	<u>0.257</u> ±0.017	<u>0.252</u> ±0.025	0.269±0.017	0.250 ±0.018	<u>0.252</u> ±0.019	<u>0.255</u> ±0.013	<u>0.256</u> ±0.025
Diabetes130US	0.402±0.003	0.401±0.003	0.400±0.003	0.383 ±0.002	0.398±0.002	0.455±0.005	0.398±0.002
dilbert	0.010 ±0.002	0.020±0.002	0.024±0.003	0.013±0.002	0.013±0.002	<u>0.012</u> ±0.002	0.039±0.004
dionis	0.089 ±0.002	0.099±0.002	0.114±0.001	0.199±0.008	0.128±0.023	0.435±0.003	0.123±0.002
dna	<u>0.044</u> ±0.004	0.063±0.007	0.056±0.006	0.046±0.003	0.040 ±0.004	<u>0.041</u> ±0.003	0.050±0.005
fabert	0.312±0.009	0.354±0.009	0.370±0.009	0.285 ±0.006	0.386±0.008	0.299±0.006	0.317±0.008
Fashion-MNIST	0.097±0.001	0.106±0.002	0.109±0.002	0.099±0.001	0.091 ±0.001	<u>0.092</u> ±0.002	0.122±0.002
gina	<u>0.053</u> ±0.005	0.080±0.006	0.090±0.005	0.047 ±0.005	0.053±0.005	0.061±0.005	0.077±0.008
guillermo	0.175±0.004	0.271±0.010	0.243±0.006	0.165 ±0.004	0.171±0.003	0.179±0.004	0.197±0.005
helena	0.617±0.002	0.602 ±0.002	0.623±0.002	0.631±0.003	0.638±0.003	0.718±0.003	0.647±0.002
Higgs	0.250 ±0.001	0.255±0.001	0.253±0.001	0.257±0.001	0.259±0.001	0.260±0.001	0.271±0.001
Internet-Advertisements	0.024±0.003	0.026±0.004	0.026±0.005	<u>0.024</u> ±0.005	0.025±0.003	0.025±0.005	0.020 ±0.003
jannis	0.273±0.002	0.271 ±0.002	0.291±0.002	0.282±0.002	0.282±0.002	0.285±0.002	0.302±0.002
jasmine	0.207±0.014	0.206±0.012	0.207±0.012	0.187 ±0.011	<u>0.190</u> ±0.010	<u>0.195</u> ±0.012	<u>0.189</u> ±0.008
jungle_chess_2pcs_raw_endgame_complete	0.004 ±0.001	0.098±0.010	0.107±0.003	0.133±0.002	0.134±0.003	0.136±0.002	0.204±0.002
kc1	0.140 ±0.007	0.143±0.009	0.145±0.011	0.147±0.010	0.143±0.007	<u>0.144</u> ±0.010	0.141±0.007
KDDCup99	0.000±0.000	0.000±0.000	0.000±0.000	0.000 ±0.000	0.002±0.000	0.000±0.000	0.000±0.000
kick	0.099±0.001	0.100±0.001	0.098±0.001	0.096 ±0.001	0.097±0.001	0.138±0.008	0.098±0.001
madeline	0.258±0.013	0.425±0.022	0.413±0.012	0.136 ±0.008	0.198±0.011	0.195±0.018	0.262±0.008
mfeat-factors	0.016 ±0.004	0.024±0.005	0.029±0.005	0.021±0.004	0.028±0.005	0.030±0.005	0.031±0.006
MiniBooNE	0.050 ±0.001	<u>0.050</u> ±0.001	0.052±0.001	0.053±0.001	0.053±0.001	0.055±0.002	0.065±0.001
numera128.6	0.479±0.004	0.421 ±0.002	0.480±0.002	0.480±0.003	0.481±0.003	0.483±0.004	0.489±0.003
okcupid-stem	0.253±0.004	0.247±0.003	0.249±0.004	0.243 ±0.003	0.246±0.003	0.410±0.016	0.262±0.003
pc4	<u>0.095</u> ±0.012	<u>0.101</u> ±0.019	0.094 ±0.008	<u>0.099</u> ±0.012	<u>0.099</u> ±0.013	<u>0.100</u> ±0.013	<u>0.104</u> ±0.012
philippine	0.284±0.011	0.305±0.012	0.301±0.008	0.249 ±0.012	<u>0.251</u> ±0.011	<u>0.253</u> ±0.009	0.254±0.010
phoneme	<u>0.097</u> ±0.007	<u>0.101</u> ±0.007	0.120±0.013	0.097 ±0.007	0.100±0.005	<u>0.102</u> ±0.007	0.098±0.006
porto-seguro	<u>0.038</u> ±0.000	0.038 ±0.000	0.038 ±0.000	<u>0.038</u> ±0.000	0.038 ±0.000	<u>0.038</u> ±0.000	0.038 ±0.000
qsar-biodeg	<u>0.126</u> ±0.015	<u>0.133</u> ±0.013	0.121 ±0.016	0.139±0.014	0.137±0.016	<u>0.131</u> ±0.012	0.136±0.016
riccardo	0.002 ±0.000	0.004±0.001	0.006±0.001	0.003±0.001	0.003±0.001	0.003±0.001	0.048±0.003
robert	0.488±0.008	0.574±0.004	0.579±0.007	0.487±0.006	0.464 ±0.006	0.471±0.008	0.570±0.008
Satellite	<u>0.006</u> ±0.002	<u>0.007</u> ±0.002	<u>0.006</u> ±0.001	<u>0.006</u> ±0.002	<u>0.005</u> ±0.002	0.005 ±0.002	<u>0.006</u> ±0.002
segment	<u>0.077</u> ±0.009	<u>0.071</u> ±0.007	0.082±0.008	0.069 ±0.008	<u>0.071</u> ±0.006	<u>0.070</u> ±0.006	0.072±0.006
shuttle	<u>0.000</u> ±0.000	0.001±0.000	0.001±0.000	<u>0.000</u> ±0.000	0.002±0.001	0.000 ±0.000	0.000±0.000
steel-plates-fault	0.241±0.019	0.232±0.011	0.250±0.014	<u>0.223</u> ±0.014	<u>0.223</u> ±0.011	0.220 ±0.011	0.241±0.013
sylvine	<u>0.054</u> ±0.005	0.060±0.007	0.075±0.006	0.052 ±0.005	<u>0.057</u> ±0.006	0.058±0.005	0.067±0.005
volkert	0.282±0.003	0.223 ±0.003	0.271±0.003	0.299±0.002	0.291±0.002	0.296±0.002	0.341±0.002
yeast	<u>0.403</u> ±0.021	<u>0.404</u> ±0.019	0.411±0.019	0.411±0.019	<u>0.401</u> ±0.017	<u>0.409</u> ±0.024	0.391 ±0.017

Table D.6: Classification error of *tuned* methods on datasets in $\mathcal{B}_{\text{class}}^{\text{test}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-HPO	MLP-HPO	CatBoost-HPO	LGBM-HPO	XGB-HPO
ada	0.147±0.008	0.152±0.010	0.138 ±0.012	0.141±0.013	0.140±0.012
airlines	0.334±0.001	0.335±0.001	0.331±0.001	0.329 ±0.001	0.329±0.001
amazon-commerce-reviews	0.207 ±0.022	0.337±0.028	<u>0.216</u> ±0.015	0.264±0.022	0.300±0.021
Bioresponse	0.219±0.010	0.233±0.012	<u>0.209</u> ±0.010	<u>0.206</u> ±0.012	0.204 ±0.011
car	0.004 ±0.003	0.012±0.004	0.017±0.008	0.017±0.010	0.022±0.007
christine	<u>0.280</u> ±0.014	0.300±0.013	<u>0.270</u> ±0.009	0.268 ±0.012	<u>0.270</u> ±0.013
churn	0.042 ±0.003	0.058±0.004	0.048±0.004	0.048±0.005	0.048±0.005
cmc	<u>0.472</u> ±0.022	<u>0.467</u> ±0.016	0.471±0.021	<u>0.470</u> ±0.016	0.454 ±0.018
cnae-9	0.079±0.020	0.054 ±0.011	0.066±0.009	0.079±0.013	0.095±0.015
connect-4	0.132 ±0.002	0.141±0.003	0.139±0.003	0.136±0.002	0.145±0.001
covertypes	0.028 ±0.001	0.040±0.001	0.062±0.001	0.033±0.001	0.040±0.003
credit-g	0.262±0.023	0.266±0.016	0.234 ±0.024	0.268±0.027	<u>0.248</u> ±0.017
Diabetes130US	0.395±0.002	0.399±0.003	0.384 ±0.003	0.390±0.002	0.387±0.002
dilbert	0.007 ±0.001	0.027±0.003	0.014±0.002	0.014±0.002	0.022±0.004
dionis	0.088 ±0.001	0.107±0.004	0.104±0.002	0.109±0.003	0.122±0.003
dna	<u>0.043</u> ±0.005	0.059±0.007	<u>0.043</u> ±0.004	0.040 ±0.003	<u>0.040</u> ±0.003
fabert	0.309±0.006	0.365±0.007	0.286 ±0.006	0.298±0.007	0.303±0.007
Fashion-MNIST	<u>0.093</u> ±0.003	0.105±0.002	0.097±0.002	0.091 ±0.001	0.094±0.002
gina	0.046 ±0.006	0.085±0.009	0.053±0.005	<u>0.050</u> ±0.005	0.058±0.005
guillermo	0.165 ±0.002	0.238±0.005	0.170±0.002	<u>0.167</u> ±0.002	0.169±0.003
helena	0.614 ±0.003	0.621±0.005	0.622±0.002	0.624±0.003	0.626±0.002
Higgs	0.247 ±0.002	0.251±0.001	0.258±0.001	0.255±0.001	0.257±0.001
Internet-Advertisements	<u>0.024</u> ±0.002	0.023 ±0.005	<u>0.025</u> ±0.005	<u>0.025</u> ±0.004	<u>0.026</u> ±0.004
jannis	0.269 ±0.002	0.296±0.004	0.281±0.002	0.278±0.002	0.279±0.002
jasmine	0.213±0.012	0.215±0.016	0.202±0.011	<u>0.196</u> ±0.016	0.188 ±0.006
jungle_chess_2pcs_raw_endgame_complete	0.003 ±0.001	0.031±0.006	0.133±0.002	0.133±0.003	0.134±0.002
kc1	<u>0.143</u> ±0.010	<u>0.146</u> ±0.013	0.142 ±0.008	<u>0.143</u> ±0.010	<u>0.144</u> ±0.005
KDDCup99	0.000±0.000	0.000±0.000	0.000 ±0.000	0.000±0.000	0.000±0.000
kick	0.098±0.001	0.098±0.001	0.096 ±0.001	<u>0.096</u> ±0.001	<u>0.097</u> ±0.001
madeline	0.166±0.012	0.411±0.016	0.150 ±0.013	<u>0.153</u> ±0.010	<u>0.162</u> ±0.014
mfeat-factors	0.015 ±0.003	0.031±0.008	0.022±0.003	0.029±0.005	0.034±0.006
MiniBooNE	0.049 ±0.001	0.052±0.001	0.053±0.001	0.052±0.001	0.053±0.001
numera128.6	0.479 ±0.004	<u>0.479</u> ±0.002	<u>0.480</u> ±0.003	<u>0.479</u> ±0.004	0.481±0.002
okcupid-stem	0.250±0.004	0.261±0.005	0.242 ±0.003	<u>0.245</u> ±0.003	0.254±0.004
pc4	0.103±0.009	0.093 ±0.011	<u>0.096</u> ±0.012	<u>0.103</u> ±0.015	<u>0.104</u> ±0.014
philippine	0.273±0.015	0.299±0.012	0.250±0.008	0.241 ±0.010	<u>0.245</u> ±0.007
phoneme	0.098 ±0.007	0.109±0.007	0.102±0.004	0.102±0.008	0.112±0.006
porto-seguro	<u>0.038</u> ±0.000	0.038 ±0.000	<u>0.038</u> ±0.000	<u>0.038</u> ±0.000	<u>0.038</u> ±0.000
qsar-biodeg	<u>0.129</u> ±0.014	0.127 ±0.008	0.135±0.012	0.136±0.012	0.134±0.020
riccardo	0.002 ±0.001	0.005±0.001	0.003±0.001	<u>0.003</u> ±0.001	0.003±0.001
robert	0.478±0.007	0.559±0.008	0.486±0.008	0.467 ±0.005	0.475±0.008
Satellite	<u>0.006</u> ±0.001	<u>0.007</u> ±0.002	<u>0.007</u> ±0.002	<u>0.006</u> ±0.002	0.005 ±0.002
segment	0.079±0.007	0.085±0.011	0.070 ±0.007	<u>0.072</u> ±0.008	<u>0.073</u> ±0.006
shuttle	0.000 ±0.000	0.001±0.000	<u>0.000</u> ±0.000	<u>0.000</u> ±0.000	<u>0.000</u> ±0.000
steel-plates-fault	0.239±0.014	0.246±0.013	0.222 ±0.012	0.223±0.008	0.232±0.017
sylvine	0.053 ±0.005	0.081±0.005	0.055±0.005	0.053 ±0.004	0.058±0.003
volkert	0.272±0.004	0.256 ±0.003	0.301±0.003	0.285±0.004	0.290±0.003
yeast	<u>0.407</u> ±0.020	0.443±0.011	<u>0.393</u> ±0.020	<u>0.404</u> ±0.022	0.391 ±0.017

Table D.7: RMSE of *untuned* methods on datasets in $\mathcal{B}_{\text{reg}}^{\text{test}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the $\#\text{splits} \rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-TD	TabR-S-D	MLP-D	CatBoost-TD	LGBM-TD	XGB-TD	RF-D
airfoil_self_noise	0.183 \pm 0.013	0.267 \pm 0.026	0.313 \pm 0.020	0.217 \pm 0.010	0.236 \pm 0.017	0.242 \pm 0.016	0.304 \pm 0.016
Airlines_DepDelay_10M	0.986 \pm 0.011	<u>0.990</u> \pm 0.011	<u>0.990</u> \pm 0.011	<u>0.987</u> \pm 0.011	<u>0.992</u> \pm 0.011	1.007 \pm 0.011	1.020 \pm 0.010
Allstate_Claims_Severity	0.658 \pm 0.013	<u>0.671</u> \pm 0.013	<u>0.667</u> \pm 0.012	<u>0.662</u> \pm 0.015	<u>0.666</u> \pm 0.014	0.956 \pm 0.152	0.691 \pm 0.012
auction_verification	0.192 \pm 0.019	0.104 \pm 0.028	0.191 \pm 0.029	0.063 \pm 0.018	0.201 \pm 0.034	<u>0.063</u> \pm 0.014	0.119 \pm 0.018
black_friday	0.692 \pm 0.002	0.688 \pm 0.002	0.693 \pm 0.002	0.678 \pm 0.002	<u>0.678</u> \pm 0.002	0.680 \pm 0.002	0.743 \pm 0.002
brazilian_houses	0.515 \pm 0.327	<u>0.471</u> \pm 0.392	<u>0.387</u> \pm 0.254	0.277 \pm 0.129	<u>0.505</u> \pm 0.336	<u>0.471</u> \pm 0.318	0.294 \pm 0.177
Buzzinsocialmedia_Twitter	0.343 \pm 0.023	<u>0.248</u> \pm 0.017	0.338 \pm 0.025	0.324 \pm 0.024	0.298 \pm 0.023	0.363 \pm 0.028	0.235 \pm 0.016
california_housing	0.421 \pm 0.006	0.364 \pm 0.008	0.434 \pm 0.008	0.401 \pm 0.007	0.404 \pm 0.007	0.410 \pm 0.007	0.440 \pm 0.008
concrete_compressive_strength	<u>0.291</u> \pm 0.023	<u>0.299</u> \pm 0.025	0.316 \pm 0.024	<u>0.296</u> \pm 0.017	0.290 \pm 0.021	<u>0.292</u> \pm 0.019	0.323 \pm 0.015
cps88wages	0.814 \pm 0.050	<u>0.814</u> \pm 0.050	<u>0.815</u> \pm 0.050	<u>0.821</u> \pm 0.048	<u>0.829</u> \pm 0.049	<u>0.832</u> \pm 0.049	0.904 \pm 0.047
cpu_activity	0.129 \pm 0.002	0.125 \pm 0.002	0.145 \pm 0.004	0.168 \pm 0.011	<u>0.127</u> \pm 0.010	<u>0.128</u> \pm 0.013	0.136 \pm 0.004
diamonds	0.136 \pm 0.002	0.136 \pm 0.003	0.144 \pm 0.002	<u>0.138</u> \pm 0.003	0.143 \pm 0.003	<u>0.138</u> \pm 0.003	0.148 \pm 0.005
elevators	0.280 \pm 0.003	0.723 \pm 0.010	0.745 \pm 0.012	0.322 \pm 0.005	0.333 \pm 0.004	0.345 \pm 0.004	0.420 \pm 0.009
fifa	0.452 \pm 0.018	0.497 \pm 0.020	0.496 \pm 0.022	0.477 \pm 0.015	<u>0.464</u> \pm 0.013	0.538 \pm 0.022	0.473 \pm 0.014
fps_benchmark	0.006 \pm 0.002	0.027 \pm 0.004	0.036 \pm 0.003	0.071 \pm 0.019	0.093 \pm 0.023	0.034 \pm 0.004	0.094 \pm 0.011
geographical_origin_of_music	0.877 \pm 0.036	0.891 \pm 0.024	0.893 \pm 0.024	<u>0.872</u> \pm 0.024	<u>0.870</u> \pm 0.027	0.866 \pm 0.027	<u>0.875</u> \pm 0.027
health_insurance	0.774 \pm 0.005	<u>0.775</u> \pm 0.005	<u>0.775</u> \pm 0.005	<u>0.775</u> \pm 0.005	0.779 \pm 0.004	0.783 \pm 0.005	0.826 \pm 0.004
house_16H	0.569 \pm 0.013	0.563 \pm 0.016	<u>0.569</u> \pm 0.020	0.580 \pm 0.015	<u>0.572</u> \pm 0.013	0.583 \pm 0.015	0.607 \pm 0.012
house_prices_nominal	<u>0.382</u> \pm 0.055	<u>0.385</u> \pm 0.052	<u>0.411</u> \pm 0.067	0.425 \pm 0.038	<u>0.379</u> \pm 0.047	<u>0.380</u> \pm 0.046	0.374 \pm 0.044
house_sales	0.317 \pm 0.012	0.338 \pm 0.015	0.331 \pm 0.009	<u>0.320</u> \pm 0.021	<u>0.322</u> \pm 0.020	<u>0.329</u> \pm 0.020	0.358 \pm 0.018
kin8nm	0.241 \pm 0.003	0.300 \pm 0.006	0.302 \pm 0.005	0.346 \pm 0.006	0.424 \pm 0.005	0.451 \pm 0.007	0.558 \pm 0.007
kings_county	<u>0.324</u> \pm 0.015	0.344 \pm 0.019	0.333 \pm 0.010	0.321 \pm 0.021	<u>0.326</u> \pm 0.020	0.341 \pm 0.018	0.358 \pm 0.019
Mercedes_Benz_Greener_Manufacturing	0.661 \pm 0.048	<u>0.668</u> \pm 0.050	<u>0.671</u> \pm 0.050	<u>0.661</u> \pm 0.046	<u>0.681</u> \pm 0.045	0.972 \pm 0.043	0.706 \pm 0.044
miami_housing	0.277 \pm 0.011	<u>0.274</u> \pm 0.012	0.294 \pm 0.011	<u>0.277</u> \pm 0.014	0.273 \pm 0.012	<u>0.276</u> \pm 0.012	0.305 \pm 0.014
MIP-2016-regression	0.840 \pm 0.041	0.806 \pm 0.035	0.829 \pm 0.022	<u>0.813</u> \pm 0.026	0.818 \pm 0.041	<u>0.813</u> \pm 0.034	0.841 \pm 0.032
nyc-taxi-green-dec-2016	0.698 \pm 0.037	<u>0.663</u> \pm 0.031	0.709 \pm 0.025	<u>0.654</u> \pm 0.029	<u>0.665</u> \pm 0.026	1.143 \pm 0.178	0.645 \pm 0.022
pol	0.067 \pm 0.003	0.144 \pm 0.005	0.133 \pm 0.007	0.101 \pm 0.004	0.102 \pm 0.005	0.109 \pm 0.007	0.120 \pm 0.007
pumadyn32nh	0.593 \pm 0.005	0.647 \pm 0.009	0.655 \pm 0.007	<u>0.596</u> \pm 0.007	0.608 \pm 0.005	0.611 \pm 0.006	0.605 \pm 0.006
QSAR-TID-10980	0.594 \pm 0.010	<u>0.597</u> \pm 0.009	0.618 \pm 0.008	<u>0.594</u> \pm 0.008	0.590 \pm 0.007	0.597 \pm 0.006	0.613 \pm 0.008
QSAR-TID-11	0.525 \pm 0.022	0.522 \pm 0.020	0.534 \pm 0.019	0.530 \pm 0.019	<u>0.524</u> \pm 0.020	<u>0.529</u> \pm 0.018	<u>0.541</u> \pm 0.020
quake	0.993 \pm 0.034	<u>0.991</u> \pm 0.038	0.988 \pm 0.033	0.987 \pm 0.034	<u>0.990</u> \pm 0.034	0.988 \pm 0.033	1.036 \pm 0.030
Santander_transaction_value	0.915 \pm 0.027	1.009 \pm 0.021	1.011 \pm 0.022	<u>0.879</u> \pm 0.029	0.877 \pm 0.030	<u>0.884</u> \pm 0.026	0.897 \pm 0.031
sarcos	0.117 \pm 0.002	0.107 \pm 0.002	0.133 \pm 0.003	0.125 \pm 0.001	0.129 \pm 0.002	0.134 \pm 0.002	0.172 \pm 0.003
SAT11-HAND-runtime-regression	0.473 \pm 0.039	0.462 \pm 0.031	<u>0.482</u> \pm 0.031	<u>0.489</u> \pm 0.033	0.555 \pm 0.028	<u>0.490</u> \pm 0.033	0.616 \pm 0.035
socmob	0.425 \pm 0.043	<u>0.378</u> \pm 0.062	0.485 \pm 0.099	0.377 \pm 0.036	0.437 \pm 0.071	0.405 \pm 0.079	0.491 \pm 0.069
solar_flare	0.899 \pm 0.113	0.882 \pm 0.105	0.897 \pm 0.115	0.919 \pm 0.130	<u>0.922</u> \pm 0.123	0.910 \pm 0.133	1.016 \pm 0.101
space_ga	0.562 \pm 0.041	<u>0.512</u> \pm 0.050	0.509 \pm 0.029	0.579 \pm 0.057	0.574 \pm 0.055	0.580 \pm 0.051	0.620 \pm 0.057
topo_2_l	0.907 \pm 0.012	<u>0.909</u> \pm 0.009	<u>0.908</u> \pm 0.012	<u>0.914</u> \pm 0.010	<u>0.915</u> \pm 0.013	<u>0.916</u> \pm 0.013	0.922 \pm 0.013
video_transcoding	0.056 \pm 0.004	0.081 \pm 0.007	0.095 \pm 0.005	<u>0.057</u> \pm 0.004	0.067 \pm 0.005	0.067 \pm 0.006	0.116 \pm 0.007
wave_energy	0.003 \pm 0.001	0.023 \pm 0.002	0.110 \pm 0.004	0.078 \pm 0.001	0.139 \pm 0.001	0.193 \pm 0.002	0.414 \pm 0.002
Yolanda	0.791 \pm 0.002	0.752 \pm 0.003	0.794 \pm 0.003	0.802 \pm 0.003	0.804 \pm 0.003	0.804 \pm 0.003	0.843 \pm 0.003
yprop_4_l	0.903 \pm 0.011	0.890 \pm 0.012	0.905 \pm 0.011	<u>0.899</u> \pm 0.012	<u>0.899</u> \pm 0.011	0.906 \pm 0.013	<u>0.902</u> \pm 0.016

Table D.8: RMSE of *tuned* methods on datasets in $\mathcal{B}_{\text{reg}}^{\text{test}}$, averaged over ten train-validation-test splits. When we write $a \pm b$, a is the mean error on the dataset and $[a - b, a + b]$ is an approximate 95% confidence interval for the mean in the #splits $\rightarrow \infty$ limit. The confidence interval is computed from the t -distribution using a normality assumption as in Appendix C.5. In each row, the lowest mean error is highlighted in bold, and errors whose confidence interval contains the lowest error are underlined.

Dataset	RealMLP-HPO	MLP-HPO	CatBoost-HPO	LGBM-HPO	XGB-HPO
airfoil_self_noise	0.177 ±0.011	0.234±0.013	0.230±0.010	0.244±0.016	0.251±0.012
Airlines_DepDelay_10M	0.985 ±0.011	<u>0.989</u> ±0.011	<u>0.986</u> ±0.011	<u>0.986</u> ±0.011	<u>0.988</u> ±0.011
Allstate_Claims_Severity	0.655 ±0.012	<u>0.666</u> ±0.012	<u>0.656</u> ±0.011	<u>0.661</u> ±0.014	<u>0.660</u> ±0.012
auction_verification	0.098±0.016	0.181±0.019	0.060 ±0.020	0.127±0.028	0.086±0.010
black_friday	0.686±0.003	0.693±0.002	<u>0.679</u> ±0.002	0.679 ±0.002	0.680±0.002
brazilian_houses	<u>0.327</u> ±0.119	<u>0.393</u> ±0.289	0.315 ±0.166	0.489±0.333	0.349±0.272
Buzzinsocialmedia_Twitter	0.267±0.021	0.271±0.014	0.322±0.026	0.291±0.026	0.223 ±0.018
california_housing	0.415±0.008	0.436±0.009	<u>0.403</u> ±0.007	0.399 ±0.007	<u>0.401</u> ±0.008
concrete_compressive_strength	<u>0.283</u> ±0.024	0.306±0.025	0.264 ±0.023	<u>0.273</u> ±0.018	<u>0.271</u> ±0.017
cps88wages	0.814±0.050	0.814±0.050	<u>0.815</u> ±0.050	0.815±0.050	0.814 ±0.050
cpu_activity	0.126±0.004	0.143±0.005	<u>0.123</u> ±0.003	<u>0.124</u> ±0.007	0.121 ±0.006
diamonds	0.134 ±0.002	0.142±0.004	<u>0.136</u> ±0.002	0.138±0.003	<u>0.135</u> ±0.002
elevators	0.276 ±0.003	0.765±0.014	0.306±0.004	0.317±0.003	0.322±0.005
fifa	0.449 ±0.020	0.509±0.017	<u>0.457</u> ±0.018	<u>0.458</u> ±0.024	0.478±0.019
fps_benchmark	0.005 ±0.002	0.010±0.001	0.038±0.019	0.018±0.002	0.033±0.007
geographical_origin_of_music	0.889±0.037	0.900±0.024	<u>0.862</u> ±0.019	<u>0.860</u> ±0.026	0.852 ±0.021
health_insurance	<u>0.774</u> ±0.005	<u>0.774</u> ±0.006	<u>0.773</u> ±0.005	<u>0.774</u> ±0.005	0.773 ±0.005
house_16H	0.563 ±0.017	<u>0.576</u> ±0.017	<u>0.572</u> ±0.014	<u>0.570</u> ±0.016	<u>0.574</u> ±0.018
house_prices_nominal	<u>0.405</u> ±0.068	<u>0.381</u> ±0.067	0.364 ±0.039	<u>0.385</u> ±0.040	<u>0.378</u> ±0.054
house_sales	<u>0.317</u> ±0.013	<u>0.337</u> ±0.012	0.308 ±0.019	<u>0.317</u> ±0.016	0.322±0.014
kin8nm	0.238 ±0.002	0.296±0.006	0.377±0.008	0.423±0.014	0.460±0.007
kings_county	0.325±0.017	0.334±0.010	0.307 ±0.016	<u>0.321</u> ±0.016	0.323±0.014
Mercedes_Benz_Greener_Manufacturing	<u>0.657</u> ±0.050	0.682±0.048	0.653 ±0.049	0.658±0.047	<u>0.653</u> ±0.049
miami_housing	<u>0.266</u> ±0.011	0.293±0.016	0.256 ±0.008	<u>0.268</u> ±0.013	0.270±0.014
MIP-2016-regression	0.848±0.041	<u>0.826</u> ±0.042	<u>0.792</u> ±0.031	0.792 ±0.035	<u>0.811</u> ±0.023
nyc-taxi-green-dec-2016	0.616 ±0.039	0.680±0.037	0.659±0.025	0.658±0.026	0.664±0.027
pol	0.059 ±0.004	0.134±0.004	0.118±0.003	0.106±0.005	0.108±0.006
pumadyn32nh	0.589 ±0.006	0.632±0.008	0.597±0.005	0.603±0.006	0.605±0.006
QSAR-TID-10980	0.599±0.011	0.610±0.009	0.597±0.008	0.583 ±0.008	0.591±0.008
QSAR-TID-11	<u>0.517</u> ±0.022	0.534±0.020	<u>0.528</u> ±0.020	0.511 ±0.020	<u>0.520</u> ±0.020
quake	<u>0.998</u> ±0.034	0.988 ±0.033	<u>0.992</u> ±0.032	<u>0.988</u> ±0.036	<u>0.991</u> ±0.031
Santander_transaction_value	0.892±0.031	0.947±0.029	0.886±0.027	0.856 ±0.032	<u>0.864</u> ±0.030
sarcos	0.102 ±0.002	0.112±0.002	0.136±0.002	0.128±0.002	0.132±0.002
SAT11-HAND-runtime-regression	0.442 ±0.052	0.481±0.030	0.513±0.031	0.499±0.029	0.528±0.041
socmob	<u>0.380</u> ±0.050	<u>0.432</u> ±0.093	0.367 ±0.073	<u>0.405</u> ±0.065	0.418±0.066
solar_flare	<u>0.925</u> ±0.099	<u>0.891</u> ±0.105	<u>0.902</u> ±0.111	0.889 ±0.105	<u>0.908</u> ±0.088
space_ga	0.502 ±0.040	<u>0.509</u> ±0.025	0.556±0.045	0.554±0.050	0.572±0.048
topo_2_1	<u>0.908</u> ±0.012	<u>0.907</u> ±0.011	<u>0.909</u> ±0.012	0.904 ±0.013	<u>0.908</u> ±0.014
video_transcoding	0.052 ±0.005	0.064±0.006	0.073±0.002	0.067±0.004	0.072±0.003
wave_energy	0.003 ±0.001	0.032±0.003	0.049±0.001	0.081±0.004	0.095±0.009
Yolanda	0.785 ±0.002	0.789±0.002	0.808±0.003	0.793±0.004	0.799±0.004
yprop_4_1	0.904±0.011	0.905±0.009	0.902±0.010	0.890 ±0.010	<u>0.894</u> ±0.012