
Self-Optimizing Random Forests

Anonymous¹

¹Anonymous Institution

Abstract Random Forests (RFs) have become a standard pick of data scientists for classification and regression tasks over the last two decades. Several modifications of RFs have been proposed to eliminate the limitation of standard RFs to axis-aligned splits through the notion of Oblique Random Forests (ORFs), which allow that splits are defined over linear combinations of attributes rather than a single attribute. However, like there is no single best learner for all supervised learning problems, there is also not a single best RF type. In this paper, we present self-optimizing random forests (SORFs). SORFs incrementally create a forest by growing different tree types and identifying the best tree type based on extrapolations of the forest performance curves obtained from out-of-bag performance estimate. Our exhaustive empirical evaluation shows that SORFs consistently achieve the maximum performance across three considered types of RFs while requiring only half the time for training all these forests on average. At the same time, SORFs outperform standard RFs statistically significantly and by at least 0.01 in accuracy on 25% of the considered datasets and even substantially beyond 0.35 in two cases.

1 Introduction

Random Forests (RFs) (Breiman, 2001) have become a standard pick of data scientists for classification and regression tasks over the last two decades. RFs are Bagging ensembles (Breiman, 1996) of decision trees (Quinlan, 1997) in which the decision trees only have access to a random subset of features in each inner node (Ho, 1998). The main strength of RFs is that they exhibit oftentimes strong performance without the need of hyperparameter tuning.

Several modifications of RFs have been proposed to eliminate the limitation of standard RFs to axis-aligned splits through the notion of Oblique Random Forests (ORFs). ORFs are RFs in which the split is not defined over a single original feature but a linear combination thereof, which amounts to an arbitrarily oriented hyperplane in the input space. ORFs were already discussed in Breiman's original work (Breiman, 2001), but optimizing the decision direction was, to our knowledge, first considered by Lemmond et al. who used a Linear Discriminant Analysis (LDA) to identify the split direction (Lemmond et al., 2008). Follow-up approaches considered a solution by regression (Menze et al., 2011), the application of a (potentially kernelized) Principal Component Analysis (PCA), (Zhang and Suganthan, 2014; Zhang et al., 2014; Wang et al., 2020) or even a combination of PCA and LDA (Zhang and Suganthan, 2014). ORFs have a close connection to Rotation Forests (Rodríguez et al., 2006), which compute the projection not in the inner leafs but previously once for the whole tree. Thereby, Rotation Forests can be seen as one type of ORFs as well.

As often observed in algorithm selection or configuration, there is not a single best rotation strategy. In many situations, it is best to not use any rotation at all while sometimes and LDA-based rotation or PCA-based rotation is best. Based on this observation, Zhang et al. proposed ensemble forests (Zhang and Suganthan, 2014), which optimize over different possible rotations in each inner node (including no rotation). However and perhaps surprisingly, we could verify that this type of forest does not consistently achieve the best performance obtained by any of the three forest types while exhibiting a substantially increased training time.

This paper makes two contributions. The minor contribution is to bring Oblique Random Forests to the attention of the AutoML community, which has not taken notice of this interesting extension

of RFs so far. The main contribution is then the introduction of Self-Optimizing Random Forests (SORFs). The idea is to incrementally create a forest by growing different tree *types*. Extrapolating the performance curves obtained from the out-of-bag error of sub-forests of the different tree types, SORFs early detect the type of tree that is most suitable for the dataset on which they are being trained. As soon as the performance curve models have stabilized, the best tree type is picked, and only trees from this type are grown. Here, we focus on standard random trees, LDA-rotated trees, and PCA-rotated trees. Besides, SORFs use this performance curve to automatically decide whether to add additional trees to the forest. So SORFs do *not* require the number of trees as a hyperparameter, which makes them even more straightforward to use than standard RFs.

Our experimental evaluation reveals that SORFs consistently achieve close-optimal performance compared to the best of the three RF types while requiring substantially less time to be trained. On 96 classification datasets, SORFs present an average regret of only 0.001 in accuracy compared to the single best forest type. On 25% of the datasets, significant improvements of at least 0.01 in accuracy can be achieved over standard RFs, in two cases even above 0.3. At the same time, SORFs exhibit substantially lower runtime compared to evaluating all of the three forest types separately, which would be the time required by an AutoML tool that treat them as a black box.

Therefore, we consider that a broader impact of SORFs to AutoML or even data science in general could be that SORFs replace standard RFs and even portfolios containing Oblique RFs. Thanks to the obsolete forest size hyperparameter this will even *reduce* the search space size of an AutoML tool that replaces standard RFs by SORFs.

2 Background on Oblique Random Forests

Random Forests (RFs) are ensembles of decision trees with two special properties (Breiman, 2001). First, RFs are *bagging ensembles* (Breiman, 1996), which means that each decision tree is being trained not on the original data but on a bootstrap sample thereof. Second, in each inner node of the tree, the training mechanism considers a random subspace (Ho, 1998), which means that it has access only to a random subset of the features. Both mechanisms diversify the ensemble.

Oblique RFs (ORFs) are a generalization of RFs that consist of decision trees in which the split points are not necessarily axis-aligned. Formally, a split point is then not formulated as $x_j \leq c$, where x_j is the value of the j -th feature of some instance x , but it is formulated over a *projection* of the considered attributes to some line β , i.e., $x^T \beta \leq c$. ORFs have been introduced by Breiman himself (Breiman, 2001) with random projections but received attention rather sporadically over time (Lemmond et al., 2008; Menze et al., 2011; Zhang and Suganthan, 2014; Zhang et al., 2014; Wang et al., 2020). A special type of ORF is a *rotation forest* (Rodríguez et al., 2006), in which the projections are defined *per tree* and cannot vary among inner nodes.

The general learning algorithm of an oblique random forest consisting of L trees is as follows:

1. For each $i = 1, \dots, L$
2. Create a dataset D_i by sampling from the original dataset D with replacement (bootstrapping).
3. At each node with purity $< \pi$ and with more than η instances associated:
 - (a) randomly select r out of the original d features.
 - (b) transform the data available in the node from the r selected features according to some projection to the new directions β_1, \dots, β_r
 - (c) calculate the best split in the transformed space, and create a left and right child with instances partitioned according to the split point.
 - (d) repeat this procedure with the left and right child respectively.

The only difference to the standard RF is step 3b, which is only used by ORFs. The transformation used in that part must be stored in the node because it needs to be applied to new instances arriving at prediction time. A new instance x is classified like in a standard RF except that the (original) instance is projected in each inner node according to the projection stored in it. The ORF approaches mainly differ in how the projections in step 3b are chosen. We briefly recap four of them.

The simplest approach is to create random projections. Such an approach was mentioned by Breiman himself (Breiman, 2001). In this case, the coefficients of β are uniformly drawn from $[-1, 1]$. Notably, random projections do not even need to be rotations of the original data.

Another technique to create the split is to map the class labels to numbers and to address the direction question as a regression problem. Menze et al. (Menze et al., 2011) introduce an approach in which the coordinate of a point on the line onto which it is projected is directly interpreted as a class prediction. Treating classes as numbers, the distance between the true label and this prediction is computed in order to identify a smooth loss. The best projection line is then found with a ridge regression approach.

Another broadly applied technique is to transform the data via a Principal Component Analysis (PCA) or Linear Discriminant Analysis (LDA). The PCA approach with a full rotation without projection is for example taken in (Zhang and Suganthan, 2014). If desired, only a subset of the first k principal components can be considered for projection. An extreme case is taken in (Wang et al., 2020), which only consider the case of $k = 1$ and hence project the data onto a line. An approach based on LDA projection is taken for example in (Lemmond et al., 2008; Zhang and Suganthan, 2014). In this paper, we refer to RFs of these types, i.e., RFs that project the data in each inner node, as PCA RFs and LDA RFs respectively.

While oblique RFs have the potential to improve prediction performance of standard RFs, they sacrifice interpretability to do so. Axis-parallel splits as created by standard trees can be explained easily. The same does not hold for PCA or LDA trees, in which splits are based on linear combinations of attributes. It is still possible to describe rules in terms of these linear combinations, but they make the rules longer and less intuitive.

To make projection-based RFs more stable and consistently superior to standard RFs, Zhang et al. (Zhang and Suganthan, 2014) have introduced projection ensemble RFs (PERFs). This has motivated the application of PERFs, which consider possibly *several* projections *in addition* to the standard random space used by the standard RF. The term “ensemble” here stems from the fact that an ensemble of projections (including no projection) is considered in each inner node. In (Zhang and Suganthan, 2014), the authors consider the identity (no) projection, a PCA, and an LDA.

However, PERFs are not a good final answer to the question of rotations. While experimental evaluations have shown that such projections, specifically LDA, oftentimes improve over standard RFs, there are also some counter-examples in which results deteriorate. Besides, the training time of PERFs is enormous corresponding to the one of training all forest types.

3 Self-Optimizing Random Forests

The core idea of Self-Optimizing Random Forests (SORFs) is to maintain different tree *types* and to detect during training which of the tree types works best. At a high level, SORFs grow different sub-forests, one for each tree type, and observe which of these forests evolves best. Projecting the performance curves of each of these sub-forest, the best tree type is chosen as soon as the performance projection is stable for each of the types. The chosen forest type is then grown until its performance curve plateaus or a predefined maximum number of trees has been reached.

The final forest will only contain the trees of the chosen type. That is, the trees grown for the tree types not finally chosen are simply discarded. While it would be possible in principle to create a mixed forest with different tree types, we abstain from this option in this paper to create a simple baseline for then perhaps more sophisticated follow-up works. The same holds for other tree types like kernel PCA trees etc.

The main difference to PERFs is that a PERFs make the rotation decision within each inner node of each tree while SORFs commit to a rotation type once for all trees. That is, a single tree in a PERF can contain several different rotation technique (one at each inner node). In contrast, a single tree in a SORF can only contain one (or no) rotation technique, which is used in *all* of the nodes. While several tree type exist during the training phase of a SORF, the final SORF model will contain only trees of a single type, so that all split points in the whole forest are of the same rotation type. SORFs are hence less flexible than PERFs in terms of the space of forests they can build. Again, we conjecture that the performance of SORFs can be improved when *mixing* different tree types, but we want to keep the analysis in this paper simple and defer analyses of this type to future work.

3.1 Creating Performance Curves Using the Out-of-Bag Error

If the trees of an RF are trained (semi-)sequentially, one can create a performance curve during training. This performance curve is similar to what Cortes et al. (Cortes et al., 1994) refer to as a *capacity* curve since it shows the performance of the ensemble, i.e., the accuracy or the cross-entropy-loss, as a function of the ensemble *size* and hence the capacity of the ensemble. We deliberately avoid the term of a learning curve here, because this is typically considered as a function of the number of training samples or time or iterations spent by an optimizer.

There are mainly two possibilities to create this type of performance curve in RFs. The first one is to separate a fraction of data points for this purpose. This is the same strategy most AutoML tools would apply when considering each forest type as one algorithm and optimizing over these (without of course building an explicit performance curve). The second approach is to exploit the fact that random forests have the capacity to estimate the generalization performance through the notion of the out-of-bag (OOB) estimate. The OOB estimate is computed simply by predicting the labels of all the data points in the training data but, for every instance x_i in the dataset only using those trees of the ensemble for predictions that did not have x_i in their bootstrap sample.

None of the two strategies is obviously superior over the other. Based on the claims that the OOB error is an unbiased estimate for the test error (Breiman, 2001; Zhang et al., 2010; Goldstein et al., 2011) and quickly becomes stable, one might consider the OOB error the better choice since it does not require to sacrifice data points. However, it has been known for a time that the OOB estimate can (even for forests with hundreds of trees and with thousands of instances) substantially deviate from the test performance (Bylander, 2002; Matthew et al., 2011; Janitza and Hornung, 2018). While literature mostly reports that OOB is too optimistic (Janitza and Hornung, 2018), we also could observe many cases in which the test performance was significantly *better* than the OOB estimate. So on large datasets it could be preferable to use a validation fold instead of an OOB estimate to create the performance curve.

In this paper, we stick to the OOB estimate despite its possible bias. Separating validation data can be unacceptable on small datasets, and providing a principled decision criterion for one of the two techniques based on the dataset at hand is not straight forward. The threat of using the OOB is that projected OOB estimate curves have a different best tree type than the true test performance curves. In that case, using the OOB leads to a sub-optimal decision. While this is a valid concern, our experiments show that using the OOB estimate yields convincing results and is, in spite of its shortcomings, at least order-preserving. We clearly stress that these results do *not* say that OOB is consistently a good estimator for the test performance but only that the order of tree types with respect to performances is apparently not severely affected, so that SORFs still achieve optimal performance. A study how this compares to explicitly building a validation fold is left for future work.

3.2 Performance Curve Extrapolation

In order to speed up the learning process, the performance curve of the RFs as a function of the number of trees is extrapolated. This is a regression problem in which the parameters of a

parametric curve model are optimized in order to estimate the value of the performance curve at any arbitrary point (here number of trees). Once these extrapolations are stable, the sub-optimal tree types are discarded in order to avoid that the full number of trees is built. Sec. 3.3 describes how this model is used for decision making in more detail.

In this paper, we follow an extrapolation technique based on a maximum likelihood estimate for the four-parametric Morgan-Mercer-Flodin (MMF) model class, which has been used to model performance curves before in terms of sample size (Gu et al., 2001) or iterations of a neural network (Domhan et al., 2015). This class models a performance curve by

$$f(x) = \frac{ab + cx^d}{b + x^d} \quad (1)$$

where x in our case is the number of trees in the forest. Given a set of observations, the parameters $a, b, c, d \geq 0$ can be easily estimated with standard interpolation techniques such as the Levenberg-Marquardt method (Bard, 1974). Once the parameters are estimated, one can predict the performance for any specific forest size. If no limit is given, the best possible performance of the forest is estimated by $\lim_{x \rightarrow \infty} f(x) = c$.

We chose the MMF curve class since it showed best results in preliminary experiments and other works (Gu et al., 2001). Several other curve models have been used in the past, most notably the inverse power law (IPL), which is by margin the most widely used model for extrapolation (Mohr and van Rijn, 2022). Due to its relevance in the literature, Appendix B.2 describes the detailed results that would have been obtained if the IPL model class would have been used instead of the MMF class. While there are still some gains in the runtime, the advantage is much less pronounced.

3.3 The Training Algorithm

The training algorithm is simple and sketched in Alg. 1. It consists of two stages: In the first stage, a forest is trained for each tree type in sequence until a stopping criterion is met. The stopping criteria are (i) a stagnation of the performance curve, (ii) stability in the forecast created by the performance curve extrapolation, or (iii) that a possibly predefined maximum number of trees L has been trained. Once the first stage is completed, the performance forecast is used to pick the tree type that is expected to perform best. The algorithm then adds trees to the forest of this tree type and updates the forecast until stagnation of the performance curve is detected or the upper limit of the number of trees is reached. Stagnation detection and forecast stability are checked by a learning curve model c based on a curve assumption discussed in Sec. 3.2, and we now describe more formally how those decisions are being made.

The stagnation stopping criterion `c.is_stale()` is based on two parameters. The parameter ϵ_{stag} controls by how much the accuracy of the forest must improve at least in order to keep training. However, the performance curves obtained from OOB estimates can exhibit a significant zig-zag behavior, because, in contrast to cross-validation, there is only one observation at each forest size. To reduce this effect, the check is not applied to the curve itself but a *smoothened* curve in which the performance at “size” k averages the values of the forest sizes $[k - w_{stag}, k]$. Observe that, since the increments in size are always 1, ϵ_{stag} is a threshold for the minimum *slope* of the (smoothened) performance curve required to continue learning.

Stability in the forecast of performance curves is considered to occur in a similar way based on windowing. In each iteration, the algorithm memorizes the expected performance of the considered tree type as a result of the performance curve extrapolation. These forecasts then generate a forecast history over time. `c.stable_forecast()` returns true iff the *variance* in the $w_{forecast}$ most recent forecasts is less than $\epsilon_{forecast}$. In other words, if the prediction about the best possible performance of a forest type does not change substantially anymore, we consider the prediction model to be sufficiently accurate and rely on its prediction. Unless being configured with a very tight (small) $\epsilon_{forecast}$, it will usually be the case that this criterion will apply prior to stagnation.

Algorithm 1: SelfOptimizingRandomForest

Input : Maximum number of trees L
Stagnation parameters $(\epsilon_{stag}, w_{stag})$
Forecast stability parameters $(\epsilon_{forecast}, w_{forecast})$
Output: A trained forest of size up to L with the best tree type

```
1  $C \leftarrow [], M \leftarrow [];$ 
2 foreach tree type  $t$  do
3    $m \leftarrow \text{new RandomForest}(t);$ 
4    $c \leftarrow \text{new CurveModel}(\epsilon_{stag}, w_{stag}, w_{forecast}, \epsilon_{forecast});$ 
5   while  $m.size() < L \wedge \neg c.is\_stale() \wedge \neg c.stable\_forecast()$  do
6      $m.train\_and\_add\_tree();$ 
7      $c.tell(m.size(), m.get\_oob\_estimate());$ 
8    $M.append(m);$ 
9    $C.append(c);$ 
10  $i^* \leftarrow \arg \max_i \{C[i].predict\_performance(L)\};$ 
11  $m \leftarrow M[i^*];$ 
12  $c \leftarrow C[i^*];$ 
13 while  $m.size() < L \wedge \neg c.is\_stale()$  do
14    $m.train\_and\_add\_tree();$ 
15    $c.tell(m.size(), m.get\_oob\_estimate());$ 
16 return  $m$ 
```

4 Experimental Evaluation

Our evaluation seeks to answer the following research questions:

1. How often are standard RFs outperformed by LDA RFs or PCA RFs?
2. How consistent are SORFs in their ability to identify the best forest type?
3. What is the regret of (limited) SORFs compared to the optimal RF type?
4. How much faster are (limited) SORFs compared to growing three forests?

All insights and conclusions are limited to the datasets considered here.

4.1 Experiment Setup

To answer the above questions, we computed the voting scheme of RFs of the three tree types for forest sizes between 1 and 100 each on 96 datasets, all of which are available at openml.org (Vanschoren et al., 2013) for reproducibility. The 96 datasets are a mixture of datasets from the AutoML benchmark (Gijssbers et al., 2019) and previous works on Oblique RFs (Zhang and Suganthan, 2014). For technical reasons, we excluded datasets of the AutoML benchmark with a sparse data representation. The exact collection of datasets is described in Appendix A. All code for experiments as well as the result resource in this paper are public.¹

For every dataset, 20 random train/test splits of 80%/20% were created. From these voting schemes, which are shared with the community, it is possible to recover the behavior of the three simple RFs (standard, LDA, PCA) as well as the behavior of SORFs. This procedure not only saved computational resources but also allows other researchers to reproduce the results without explicitly training the forests again. Even more, it is possible to conduct a post-hoc analysis of

¹<https://anonymous.4open.science/r/sorf-2B66/>

other extrapolation techniques not tried for this paper. All reported performances are summarized test-fold performances.

The RFs were configured as follows. To enable better comparison, the maximum number of trees was set to 100 even though this hyperparameter is optional for SORFs. The purity parameter π was set to 0.9, and the minimum number of instances in a leaf η was set to 5. PCA RFs were configured to consider all of the dimensions (PCA simply used for rotation but not projection). The SORF hyperparameters were set to $w_{stag} = 10$, $\epsilon_{stag} = 10^{-4}$, $w_{forecast} = 10$, and $\epsilon_{forecast} = 10^{-3}$, which we would arguably consider reasonable default values; they have not been tuned. While a sensitivity study of these parameters is beyond our scope, we do not expect big surprises or high variance in such a study but that the performance is rather robust.

The computations for this experimentation required 58.5 CPU days in a compute center with Linux machines, each of them equipped with 2.6Ghz Intel Xeon E5-2670 processors and 16GB memory. No timeouts were imposed when training the forests and computing the voting schemes.

4.2 Results

4.2.1 How often are standard RFs outperformed by LDA RFs or PCA RFs? We are not the first to provide an answer to this question. Zhang et al. (Zhang and Suganthan, 2014) have reported similar results. However, since we could not reproduce their results on ensemble forests, we aimed at verifying that ORFs indeed can substantially improve the performance over standard RFs.

An answer to this question can be found in Fig. 1. The left figure compares the performances of a standard RF with the LDA and PCA RF respectively. Every point represents one dataset: blue points compare the standard RF against an LDA RF, and orange points compare it against an PCA RF trained on the same dataset. The scores are test set accuracies averaged over the 20 random train/test splits of size 80%/20%. The exact values can be found in Appendix B.1.

The plot suggests that none of the tree types is dominant. On many datasets, the standard RF without any rotation outperforms both the LDA RF and the PCA RF (points below the line). So none of the two alternatives dominates the standard RF. On the other side, for both RF types LDA RF and PCA RF there are datasets on which they perform substantially better than a standard RF. The most extreme example is the HILL-VALLEY dataset (1566) on which the two rotation based forest types achieve virtually a 100% test accuracy while the standard RF only achieves a 60% accuracy.

While the presentation of results in the scatter plot suggests that LDA RF and PCA RF only occasionally are the best choice, the other two plots in Fig. 1 show that indeed the standard RF is the best choice in only approximately a little less than 50% of the cases. In fact, LDA RF is the best choice in 29% of the cases a PCA RF in 23% of the cases. Here “cases” refers to splits of datasets, because on some splits of a dataset one algorithm might work better than on a different split. With these results, there is a clear motivation to identify the best forest type for a given dataset.

4.2.2 How consistent are SORFs in their ability to identify the best forest type? This question is answered by the the middle and right plot of Fig. 1, now also considering the actual decisions of the SORF (y-axes). Summing up the probabilities on the diagonal of the confusion matrix, we observe that SORFs identify the correct tree type in only 75% of the cases. The middle plot reflects this by rather impure bars in which the bar for every best choice contains the correct choice in the majority but certainly not all of the cases.

The explanation for this supposedly disappointing result is that, on many datasets, there is only very little difference between the algorithm performances and the best one is only very marginally better than the other or at least the second best one. In such cases, it is possible that the forecast of the performance curve is not precise enough to identify the best algorithm.

Interestingly, the more intuitive explanation for these results, namely the mediocre quality of the out-of-bag (OOB) estimate, is not a relevant factor here. Clearly, since the OOB does not accurately resemble the test error on many datasets, one would intuitively expect that this misguidance leads




Figure 1: Left: Comparison of the accuracies of standard RFs vs LDA RFs (blue) and PCA RFs (orange). Middle and Right: Best possible RF type choice and the type choices made by a SORF.

to wrong decisions. However and perhaps surprisingly, considering the same confusion plots not with the best choice on the test data but the best choice according to the OOB estimate, the results look worse: the SORF then identifies the best tree type (according to OOB) in only 64% of the cases.

Based on these observations, the next research question about the regret becomes even more urgent. Apparently SORFs are not able to consistently identify the best tree type. So the question is how much they are outperformed by the indeed best single RF type in the situation in which they don't pick the best type.

4.2.3 What is the regret of (limited) SORFs compared to the optimal RF type?

To analyze this question, Fig. 2 shows the performance difference between SORFs and each of the three RF types. The left plot compares the accuracy of SORFs (x-axis) against the other three tree type RFs. Blue points show the comparison to LDA RFs, orange points to PCA RFs, and green points to standard RFs. Every point is the average accuracy of the two compared algorithms on a single dataset across the 20 train/test splits. Bullets show statistically significant differences according to a Wilcoxon signed rank test (p-value 0.05), and circles show datasets on which the difference is not significant in this sense. To give a different perspective on the same data, the right plot summarizes the *improvements* in accuracy the SORF achieves over each of the three individual forest types. For each of them, there is one boxplot summarizing the improvements over *all* the datasets, and one only about those on which the mean value could be identified to be statistically significantly different (same test as for left plot). Besides, there is one boxplot showing the improvement over the "oracle", which is effectively the regret of the SORF since the oracle picks the best RF type and hence cannot be improved upon. Three vertical dotted lines serve as visual aids for a deterioration of 0.01 (red line), and improvements of 0.01 and 0.03 (black lines). The vertical dashed line is simply the 0 improvement (no improvement but also no regret). The exact values can be found in Appendix B.1.

These results now show SORFs in a much more positive light than the plots in Fig. 1. There are almost *no* datasets in which SORFs significantly underperform the best RF type for that dataset (points below the line in the left plot). In fact, on only one dataset a deterioration of more than 0.01 in accuracy can be observed, and this deterioration is not statistically significant. Among the statistically significantly deviating results, the regret is always under 0.007, and in over 70% of the cases there is no regret at all since SORF obtains the optimal performance among all tree types.

From a different viewpoint, the regret of the standard RF is above 0.01 in 25% of the cases. This can be seen in the bottom boxplot of the right panel, which shows that the 75%-quantile is exactly at 0.01. In other words, we should not expect SORFs to significantly outperform standard RFs all too often, but it *can* happen, and in such situations the improvement is often not only statistically significant but also quite substantial in terms of absolute improvement.




Figure 2: Left: Comparison of accuracies of a SORF compared to standard RFs (green), LDA RFs (blue) and PCA RFs (orange). Right: Summary of improvements of the SORF over the RF types on the different datasets. The boxplots labeled with “sign.” only summarize the datasets on which the difference of the mean performance is statistically significant. The oracle entry refers to the comparison to the best RF type on a dataset (hence always negative).

All this being said, we conclude that SORFs have negligible regret when being compared to an oracle that picks the best forest type. The average regret is 0.001, and the maximum statistically significant regret is below 0.01. In other words and putting this into the light of the discussion in Sec. 4.2.2, SORFs not always identify the best forest type, but in the case where they don’t, the chosen alternative is almost equally good. The regret here is the price one has to pay getting the runtime advantages of SORFs in contrast to training all forest types and choosing the best one. We next assess precisely this runtime advantage.

4.2.4 How much faster are (limited) SORFs compared to growing three forests?. To analyze the training time behavior, Fig. 3 relates the numbers of trained trees and the actual training times of SORFs and the individual RF types in several ways. The panel with the two left plots shows the relative training time compared to the full forest of the respective types (left plot) and the number of trees grown for each of the types on average based on the chosen forest type (middle left plot). There is one row for each of the tree types. Every built SORF trained on any dataset contributes to the row belonging to the tree type it eventually selected (on that dataset (and split)). Since each SORF builds also other tree types during training, we want to know how much trees are built of finally discarded tree types and how much time is hence lost in that. So the figure shows the time spent to train such trees (leftmost plot) or the number of trees of the different types trained (middle left plot) in the different columns of the matrix; the diagonal shows the time/number of trees trained of the type that was eventually chosen by the SORF. While the left panel only shows mean values averaged over the datasets, the panel with the two boxplots on the right summarizes the relative reduction in training time (middle right plot) or numbers of trained tree (right plot) per tree type without taking the chosen tree type into account. Again, it is important to note that the SORFs were limited to a forest size of 100 for the finally chosen tree type. Without this restriction, some forests would have built more trees requiring more training time while possibly leading to better results. This analysis is beyond our scope.

Looking first at the left plots, we can observe a substantial runtime reduction on average. The fourth column of the first plot shows that SORFs exhibit a training time that is roughly only 40% of the time required to fully train all three forest types with 100 trees and then pick the best one. It also shows that the eventual choice of the tree types barely affects this number even though PCA RFs have a significantly higher runtime than the other two forest types on average (not shown in the figure). Consistently, the middle left plot shows that the number of trained trees is only 40% of the trees that would be trained otherwise, i.e. approximately 120 instead of 300. What is particularly interesting is that, on average, even for the chosen tree type only roughly 50 instead of 100 trees are trained. This suggests the usefulness of the automated stopping criterion applied by




Figure 3: Left: Mean training time compression in SORFs for different tree types depending on the finally chosen tree type. Middle left: Mean number of trees trained per tree depending on the finally chosen tree type. Middle and outer right: Mean reduction per dataset in training time/trained trees per tree type and in total summarized over all datasets.

SORFs, which could of course also be used independently by the RF training algorithms. Note that the actual numbers of trained trees of the selected type ranges between 10 and 100, and the fact that it is on average close to 50 is only a coincidence.

The right panel offers a more detailed view on the range of the concrete reductions. They summarize across the datasets, for each tree type but also in total, how much the training time or number of trees of that tree type is reduced on average on the respective dataset. For both boxplots it shows two visual aids at 50% and 25% through the dotted lines. Remarkably, we can see that the compression of both time and trained trees is better (below) than 50% in more than 75% of the cases. In roughly 25% of the analyzed datasets, these quantities are even reduced by 75% or more.

Putting everything together, we can conclude, with the usual limitation of a limited number of datasets used for analysis, that SORFs achieve the same performance as a portfolio consisting of three RF types while requiring substantially less time for training. Thanks to the projection of the performance curves and early stopping based on the detection of stale training progress, the training time can be significantly reduced while sacrificing only a marginal portion of accuracy. The trivial approach of considering full RFs for all of the types seems only justified if the user is concerned about accuracy differences on an order that is below 0.01 in accuracy.

5 Conclusion

In this paper, we have proposed Self-Optimizing Random Forests (SORFs). SORFs build sub-forests with different tree types up to a size in which the final performance of each type can be safely predicted based on a performance curve model. The experimental evaluation over 96 datasets shows that SORFs with three tree types (standard, with LDA rotations, and with PCA rotations) exhibit optimal performance compared to a portfolio of RFs of these tree types while requiring only 40% of the training time on average. While SORFs can be seen as a RF-centered AutoML approach itself, the results suggest to replace standard RFs in AutoML tools by SORFs since they cover a broader class of learners than standard RFs with potentially significant performance improvements while exhibiting a much lower training time compared to adding each of these RF types separately to the portfolio of the AutoML tool.

There are plenty of options for immediate future work. One straight forward question is on the potential improvements of SORFs if no tree limit is imposed. A second question is on the evaluation mechanism and whether the usage of a validation set instead of using the OOB estimate can improve the performance or runtime of the SORF. Finally, analyzing the prediction performance of *mixed* forests that consist of trees of *different* types would be interesting, because this could reduce the training time even more while adding more diversity to the ensemble.

References

- Bard, Y. (1974). *Nonlinear parameter estimation*. Academic Press. 401
- Breiman, L. (1996). Bagging predictors. *Mach. Learn.*, 24(2):123–140. 402
- Breiman, L. (2001). Random forests. *Mach. Learn.*, 45(1):5–32. 403
- Bylander, T. (2002). Estimating generalization error on two-class datasets using out-of-bag estimates. *Mach. Learn.*, 48(1-3):287–297. 404
- Cortes, C., Jackel, L. D., and Chiang, W. (1994). Limits in learning machine accuracy imposed by data quality. In Tesauro, G., Touretzky, D. S., and Leen, T. K., editors, *Advances in Neural Information Processing Systems 7, [NIPS Conference, Denver, Colorado, USA, 1994]*, pages 239–246. MIT Press. 405
- Domhan, T., Springenberg, J. T., and Hutter, F. (2015). Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 3460–3468. AAAI Press. 406
- Gijsbers, P., LeDell, E., Thomas, J., Poirier, S., Bischl, B., and Vanschoren, J. (2019). An open source automl benchmark. *CoRR*, abs/1907.00909. 407
- Goldstein, B. A., Polley, E. C., and Briggs, F. B. (2011). Random forests for genetic association studies. *Statistical applications in genetics and molecular biology*, 10(1). 408
- Gu, B., Hu, F., and Liu, H. (2001). Modelling classification performance for large data sets. In Wang, X. S., Yu, G., and Lu, H., editors, *Advances in Web-Age Information Management, Second International Conference, WAIM 2001, Xi’an, China, July 9-11, 2001, Proceedings*, volume 2118 of *Lecture Notes in Computer Science*, pages 317–328. Springer. 409
- Ho, T. K. (1998). The random subspace method for constructing decision forests. *IEEE Trans. Pattern Anal. Mach. Intell.*, 20(8):832–844. 410
- Janitza, S. and Hornung, R. (2018). On the overestimation of random forest’s out-of-bag error. *PloS one*, 13(8):e0201904. 411
- Lemmond, T. D., Hatch, A. O., Chen, B. Y., Knapp, D., Hiller, L., Mugge, M., and Hanley, W. G. (2008). Discriminant random forests. In Stahlbock, R., Crone, S. F., and Lessmann, S., editors, *Proceedings of The 2008 International Conference on Data Mining, DMIN 2008, July 14-17, 2008, Las Vegas, USA, 2 Volumes*, pages 55–61. CSREA Press. 412
- Matthew, W. et al. (2011). Bias of the random forest out-of-bag (oob) error for certain input parameters. *Open Journal of Statistics*, 2011. 413
- Menze, B. H., Kelm, B. M., Splitthoff, D. N., Köthe, U., and Hamprecht, F. A. (2011). On oblique random forests. In Gunopulos, D., Hofmann, T., Malerba, D., and Vazirgiannis, M., editors, *Machine Learning and Knowledge Discovery in Databases - European Conference, ECML PKDD 2011, Athens, Greece, September 5-9, 2011, Proceedings, Part II*, volume 6912 of *Lecture Notes in Computer Science*, pages 453–469. Springer. 414
- Mohr, F. and van Rijn, J. N. (2022). Learning curves for decision making in supervised machine learning - A survey. *CoRR*, abs/2201.12150. 415

Quinlan, J. R. (1997). Decision trees and instance-based classifiers. In Tucker, A. B., editor, *The Computer Science and Engineering Handbook*, pages 521–535. CRC Press. 440
441

Rodríguez, J. J., Kuncheva, L. I., and Alonso, C. J. (2006). Rotation forest: A new classifier ensemble method. *IEEE Trans. Pattern Anal. Mach. Intell.*, 28(10):1619–1630. 442
443

Vanschoren, J., van Rijn, J. N., Bischl, B., and Torgo, L. (2013). Openml: networked science in machine learning. *SIGKDD Explor.*, 15(2):49–60. 444
445

Wang, F., Wang, Q., Nie, F., Yu, W., Wang, R., and Li, Z. (2020). A forest of trees with principal direction specified oblique split on random subspace. *Neurocomputing*, 379:413–425. 446
447

Zhang, G., Zhang, C., and Zhang, J. (2010). Out-of-bag estimation of the optimal hyperparameter in subbag ensemble method. *Commun. Stat. Simul. Comput.*, 39(10):1877–1892. 448
449

Zhang, L., Ren, Y., and Suganthan, P. N. (2014). Towards generating random forests via extremely randomized trees. In *2014 International Joint Conference on Neural Networks, IJCNN 2014, Beijing, China, July 6-11, 2014*, pages 2645–2652. IEEE. 450
451
452

Zhang, L. and Suganthan, P. N. (2014). Random forests with ensemble of feature spaces. *Pattern Recognit.*, 47(10):3429–3437. 453
454

Reproducibility Checklist 455

1. For all authors... 456
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper’s contributions and scope? 457
Yes 458
459
 - (b) Did you describe the limitations of your work? 460
Yes, we explicitly state that the conclusions are limited to the considered datasets. 461
 - (c) Did you discuss any potential negative societal impacts of your work? 462
No, because we think that this does not apply here. None of the points in the ethic guidelines applies to our paper. 463
464
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? 465
Yes. 466
 - (e) Did you include the code, data, and instructions needed to reproduce the main experimental results, including all requirements (e.g., requirements.txt with explicit version), an instructive README with installation, and execution commands (either in the supplemental material or as a URL)? 467
Yes, included as URL. 468
469
470
471
 - (f) Did you include the raw results of running the given instructions on the given code and data? 472
Yes, the raw results are part of the repository coming along with this paper. 473
474
 - (g) Did you include scripts and commands that can be used to generate the figures and tables in your paper based on the raw results of the code, data, and instructions given? 475
Yes, the repository contains the python notebook that was used to generate all the figures and result tables. 476
477
478

- (h) Did you ensure sufficient code quality such that your code can be safely executed and the code is properly documented? 479
We hope so. 480
481
- (i) Did you specify all the training details (e.g., data splits, pre-processing, search spaces, fixed hyperparameter settings, and how they were chosen)? 482
Yes, in parts this is described in the paper. The exact technical details (how the train/test splits were created) at the implementation level are part of the code repository. 483
484
485
- (j) Did you ensure that you compared different methods (including your own) exactly on the same benchmarks, including the same datasets, search space, code for training and hyperparameters for that code? 486
Yes 487
488
489
- (k) Did you run ablation studies to assess the impact of different components of your approach? 490
We believe that this does not apply to our paper since SORFs do not have several components. 491
One could, maybe, evaluate SORFs for different configurations, but this was beyond our scope and probably not yield relevant insights. 492
493
- (l) Did you use the same evaluation protocol for the methods being compared? 494
Of course. 495
- (m) Did you compare performance over time? 496
Does not apply. But yes we also considered the aspect of time. 497
- (n) Did you perform multiple runs of your experiments and report random seeds? 498
Yes. The random seeds are documented in the code base and not relevant in the paper itself. 499
500
- (o) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? No, but we applied a significance test. 501
502
- (p) Did you use tabular or surrogate benchmarks for in-depth evaluations? 503
Effectively, we created our own benchmark for this problem, which can be re-used by other researchers. 504
505
- (q) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? Yes. 506
507
- (r) Did you report how you tuned hyperparameters, and what time and resources this required (if they were not automatically tuned by your AutoML method, e.g. in a NAS approach; and also hyperparameters of your own method)? Does not apply since we did not tune parameters. 508
509
510
511
2. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets... 512
- (a) If your work uses existing assets, did you cite the creators? 513
Not all of them. 514
- (b) Did you mention the license of the assets? 515
No. 516
- (c) Did you include any new assets either in the supplemental material or as a URL? Yes. The implementation of SORFs. 517
518
- (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? 519
We only use data from openml.org where consent is part of the upload policy. 520
521

(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? 522
Does not apply. 523
524

Appendix

525

A Datasets

526

id	name	inst.	features	classes	id	name	inst.	features	classes
3	kr-vs-kp	3196	36	2	1487	ozone-level-8hr	2534	72	2
6	letter	20000	16	26	1488	parkinsons	195	22	2
11	balance-scale	625	4	3	1489	phoneme	5404	5	2
12	mfeat-factors	2000	216	10	1494	qsar-biodeg	1055	41	2
13	breast-cancer	286	9	2	1499	seeds	210	7	3
23	cmc	1473	9	3	1515	micro-mass	571	1300	20
30	page-blocks	5473	10	5	1566	hill-valley	1212	100	2
31	credit-g	1000	20	2	1591	connect-4	67557	126	0
54	vehicle	846	18	4	4134	Bioresponse	3751	1776	2
55	hepatitis	155	19	2	4137	Dorothea	1150	100000	2
60	waveform-5000	5000	40	3	4534	PhishingWebsites	11055	30	2
181	yeast	1484	8	10	4538	GesturePhaseSegmenta	9873	32	5
201	pol	15000	48	0	40498	wine-quality-white	4898	11	7
299	libras_move	360	90	0	40664	car-evaluation	1728	21	4
336	SPECT	267	22	2	40670	dna	3186	180	3
346	aids	50	4	2	40677	led24	3200	24	10
380	SyskillWebert-Bands	61	2	3	40685	shuttle	58000	9	7
446	prnn_crabs	200	7	2	40687	solar-flare	1066	12	6
1042	gina_prior	3468	784	2	40701	churn	5000	20	2
1049	pc4	1458	37	2	40713	dis	3772	29	2
1067	kc1	2109	21	2	40900	Satellite	5100	36	2
1083	mouseType	214	45101	7	40910	Speech	3686	400	2
1084	BurkittLymphoma	220	22283	3	40971	collins	1000	23	30
1085	anthracyclineTaxaneC	159	61359	2	40975	car	1728	6	4
1086	ovarianTumour	283	54621	3	40978	Internet-Advertiseme	3279	1558	2
1087	hepatitisC	283	54621	3	40981	Australian	690	14	2
1088	variousCancers_final	383	54675	9	40982	steel-plates-fault	1941	27	7
1128	OVA_Breast	1545	10936	2	40983	wilt	4839	5	2
1130	OVA_Lung	1545	10936	2	40984	segment	2310	19	7
1134	OVA_Kidney	1545	10936	2	40994	climate-model-simula	540	20	2
1138	OVA_Uterus	1545	10936	2	41027	jungle_chess_2pcs_ra	44819	6	3
1139	OVA_Omentum	1545	10936	2	41142	christine	5418	1636	2
1142	OVA_Endometrium	1545	10936	2	41143	jasmine	2984	144	2
1146	OVA_Prostate	1545	10936	2	41144	madeline	3140	259	2
1161	OVA_Colon	1545	10936	2	41145	philippine	5832	308	2
1233	eating	945	6373	7	41146	sylvine	5124	20	2
1441	KungChi3	123	39	2	41156	ada	4147	48	2
1448	KnuggetChase3	194	39	2	41157	arcene	100	10000	2
1450	MindCave2	125	39	2	41158	gina	3153	970	2
1457	amazon-commerce-revi	1500	10000	50	41159	guillermo	20000	4296	2
1461	bank-marketing	45211	16	2	41163	dilbert	10000	2000	5
1464	blood-transfusion-se	748	4	2	41164	fabert	8237	800	7
1465	breast-tissue	106	9	6	41165	robert	10000	7200	10
1468	cnae-9	1080	856	9	41166	volkert	58310	180	10
1475	first-order-theorem-	6118	51	6	41168	jannis	83733	54	4
1477	gas-drift-different-	13910	129	6	41972	Indian_pines	9144	220	8
1479	hill-valley	1212	100	2	42809	kits	1000	27648	2
1485	madelon	2600	500	2	42810	PCam	4000	27648	2

527

B Results in Detail

528

Test accuracies achieved by the different approaches. Bold entries are best (mean values, averaged over 20 seeds) and underlined are not statistically significantly different from the best according to a Wilcoxon signed rank test with a p-value of 0.05.

529

530

531

B.1 Results Using MMF For Extrapolation

532

id	RF	LDA RF	PCA RF	SORF	id	RF	LDA RF	PCA RF	SORF
3	97.2±0.2	97.52±0.3	97.03±0.3	97.35±0.4	1487	93.72±0.3	93.72±0.3	93.72±0.3	93.72±0.3
6	<u>94.08±0.1</u>	<u>94.95±0.1</u>	95.03±0.1	<u>95.02±0.1</u>	1488	<u>87.87±1.6</u>	<u>88.21±1.0</u>	<u>88.23±1.3</u>	88.67±1.1
11	<u>85.44±0.8</u>	89.35±0.9	<u>89.03±0.9</u>	<u>89.29±1.1</u>	1489	<u>88.17±0.3</u>	<u>88.55±0.3</u>	<u>88.58±0.3</u>	88.59±0.3
12	95.28±0.4	97.1±0.3	96.73±0.4	97.09±0.4	1494	85.86±0.8	87.06±0.6	86.63±0.7	86.89±0.7
13	<u>73.12±1.5</u>	<u>72.63±1.7</u>	<u>72.81±1.9</u>	<u>73.06±1.7</u>	1499	<u>90.53±1.1</u>	<u>91.67±1.3</u>	<u>91.49±1.4</u>	91.86±1.3
23	52.81±1.2	52.78±1.2	52.37±1.4	53.04±1.3	1515	83.06±1.5	67.94±2.1	76.34±1.4	82.75±2.0
30	<u>91.9±1.5</u>	<u>91.89±1.4</u>	92.49±1.8	<u>92.44±1.8</u>	1566	<u>59.35±1.2</u>	100.0±0.0	<u>99.99±0.0</u>	100.0±0.0
31	<u>74.51±1.0</u>	<u>74.19±1.3</u>	<u>74.41±1.0</u>	74.67±1.1	1591	81.28±0.1	81.24±0.1	78.69±0.1	81.29±0.1
54	<u>74.4±1.1</u>	<u>78.78±1.0</u>	<u>78.73±0.9</u>	78.93±0.9	4134	79.03±0.6	77.33±0.4	76.19±0.5	78.66±1.0
55	84.5±2.1	85.07±1.9	84.79±2.2	85.78±1.7	4137	<u>90.24±0.4</u>	90.24±0.4	<u>90.24±0.4</u>	90.24±0.4
60	84.36±0.4	85.92±0.4	85.65±0.4	85.75±0.4	4534	94.73±0.3	95.36±0.2	94.89±0.2	95.24±0.3
181	<u>61.14±0.9</u>	<u>61.11±1.1</u>	61.5±0.8	<u>61.35±1.0</u>	4538	63.42±0.5	61.76±0.3	61.97±0.4	63.25±0.6
201	84.18±0.2	83.96±0.2	83.2±0.2	84.17±0.2	40498	<u>65.86±0.6</u>	<u>66.67±0.6</u>	66.99±0.5	<u>66.86±0.6</u>
299	<u>74.31±2.2</u>	<u>82.99±1.7</u>	<u>83.6±1.9</u>	84.0±1.9	40664	<u>92.92±0.6</u>	<u>94.27±0.6</u>	94.67±0.6	<u>94.37±0.8</u>
336	<u>82.42±1.7</u>	<u>81.96±1.5</u>	82.48±1.2	<u>82.36±1.5</u>	40670	94.56±0.3	93.1±0.4	93.58±0.5	94.55±0.3
346	<u>42.62±6.6</u>	<u>34.7±6.2</u>	35.35±8.1	39.71±7.1	40677	71.61±0.7	71.82±0.6	71.59±0.7	71.86±0.7
380	62.89±4.9	62.89±4.9	<u>62.36±5.3</u>	62.89±4.9	40685	99.52±0.0	99.36±0.1	99.51±0.0	99.52±0.0
446	<u>88.52±1.8</u>	<u>98.92±0.7</u>	<u>99.03±0.8</u>	99.13±0.8	40687	75.27±1.0	74.82±0.9	74.69±1.1	75.02±1.0
1042	92.18±0.3	90.22±0.3	90.73±0.3	92.18±0.3	40701	89.44±0.2	88.9±0.3	89.14±0.3	89.4±0.3
1049	89.76±0.5	90.07±0.5	89.66±0.4	90.07±0.4	40713	98.42±0.1	98.42±0.1	98.42±0.1	98.42±0.1
1067	<u>85.85±0.6</u>	<u>86.12±0.5</u>	86.38±0.5	<u>86.27±0.5</u>	40900	98.58±0.1	98.58±0.1	98.58±0.1	98.58±0.1
1083	65.99±2.7	60.29±3.0	59.99±2.7	65.87±2.8	40910	98.32±0.2	98.32±0.2	98.32±0.2	98.32±0.2
1084	81.44±1.5	88.8±1.8	79.64±1.8	88.8±1.8	40971	23.61±1.7	23.91±1.2	24.11±1.4	24.64±1.0
1085	<u>54.91±4.4</u>	<u>52.79±3.5</u>	55.06±3.3	<u>55.0±3.9</u>	40975	80.15±1.0	81.88±1.0	84.14±0.9	<u>84.09±1.0</u>
1086	<u>88.21±1.0</u>	<u>88.19±1.3</u>	<u>88.31±1.2</u>	88.64±1.2	40978	90.54±0.3	92.33±0.3	91.93±0.3	92.27±0.3
1087	<u>88.21±1.0</u>	<u>88.19±1.3</u>	<u>88.31±1.2</u>	88.64±1.2	40981	<u>86.49±1.2</u>	<u>86.39±1.2</u>	<u>86.7±0.8</u>	86.8±1.0
1088	79.33±1.5	76.01±1.7	74.55±1.6	79.04±1.7	40982	76.23±0.6	74.86±0.6	75.63±0.6	76.05±0.7
1128	95.24±0.5	97.25±0.4	89.21±1.0	97.25±0.4	40983	94.66±0.2	94.66±0.2	94.66±0.2	94.66±0.2
1130	91.96±0.4	91.96±0.4	91.96±0.4	91.96±0.4	40984	<u>92.24±0.4</u>	93.8±0.5	93.49±0.4	93.76±0.5
1134	96.64±0.4	98.62±0.2	95.83±0.4	98.62±0.2	40994	91.36±0.8	91.36±0.8	91.36±0.8	91.36±0.8
1138	92.09±0.4	92.09±0.4	92.09±0.4	92.09±0.4	41027	81.42±0.2	81.22±0.1	81.32±0.2	81.36±0.2
1139	95.08±0.3	95.08±0.3	95.08±0.3	95.08±0.3	41142	71.69±0.4	72.72±0.4	71.07±0.5	72.56±0.5
1142	96.0±0.2	96.0±0.2	96.0±0.2	96.0±0.2	41143	80.86±0.5	79.86±0.8	80.02±0.5	80.66±0.6
1146	95.55±0.2	95.55±0.2	95.55±0.2	95.55±0.2	41144	74.83±0.9	61.5±0.8	63.26±0.9	74.83±0.9
1161	95.98±0.4	96.78±0.3	90.33±1.3	96.78±0.3	41145	74.15±0.5	72.37±0.4	72.48±0.5	74.09±0.5
1233	48.02±1.4	57.79±1.1	36.73±2.0	57.79±1.1	41146	<u>92.27±0.2</u>	<u>92.53±0.2</u>	92.76±0.2	<u>92.69±0.2</u>
1441	<u>86.07±2.3</u>	<u>86.24±2.1</u>	86.73±1.9	<u>86.69±1.9</u>	41156	84.84±0.5	84.49±0.5	83.7±0.8	84.79±0.5
1448	<u>79.35±1.7</u>	<u>79.91±2.1</u>	<u>80.55±1.8</u>	80.77±1.7	41157	71.96±3.5	71.24±3.7	66.25±6.1	72.43±3.5
1450	<u>69.01±2.8</u>	<u>69.56±3.2</u>	70.22±3.8	<u>69.68±3.2</u>	41158	91.4±0.4	87.65±0.5	86.56±0.5	91.4±0.4
1457	30.87±1.6	21.19±1.1	17.82±1.0	30.29±2.6	41159	79.89±0.2	62.63±0.4	60.59±0.4	79.89±0.2
1461	89.34±0.1	89.24±0.1	89.17±0.1	89.31±0.1	41163	94.79±0.2	94.39±0.2	93.45±0.3	94.45±0.6
1464	77.83±1.3	78.25±1.2	77.86±1.2	78.08±1.4	41164	67.57±0.5	60.38±0.4	60.18±0.5	67.57±0.5
1465	40.71±4.9	<u>40.14±6.0</u>	<u>39.68±4.7</u>	<u>40.25±4.8</u>	41165	39.29±0.6	36.74±0.5	27.52±0.4	39.29±0.6
1468	89.33±0.7	91.17±0.6	87.27±0.9	90.92±1.0	41166	64.62±0.1	62.44±0.2	62.33±0.2	64.62±0.1
1475	60.92±0.4	<u>60.69±0.4</u>	<u>60.77±0.4</u>	<u>60.91±0.4</u>	41168	69.23±0.1	68.24±0.2	68.8±0.2	69.1±0.2
1477	<u>98.41±0.1</u>	<u>98.97±0.1</u>	99.19±0.0	<u>99.18±0.0</u>	41972	<u>88.9±0.2</u>	<u>91.88±0.3</u>	92.48±0.2	<u>92.45±0.2</u>
1479	<u>55.84±1.4</u>	<u>92.21±0.5</u>	92.45±0.6	<u>92.43±0.6</u>	42809	73.38±0.9	64.89±1.8	58.43±1.9	73.38±0.9
1485	68.3±1.1	57.5±0.9	57.84±1.1	68.3±1.1	42810	74.59±0.5	68.81±0.5	63.83±0.7	74.59±0.5

533

B.2 Results Using IPL For Extrapolation

534

The IPL is defined as $f(x) = a - bx^{-c}$ (Cortes et al., 1994) and the most widely used class model for learning curves (Mohr and van Rijn, 2022). Using it for extrapolation, we obtain the following results, which are much less exciting than those obtained using MMF.

535

536




Figure 4: Computations with the IPL. Semantics identical to Fig. 3.

537

openmlid	RF	LDA RF	PCA RF	SORF	openmlid	RF	LDA RF	PCA RF	SORF
3	97.2±0.2	97.52±0.3	97.03±0.3	97.5±0.3	1487	93.72±0.3	93.72±0.3	93.72±0.3	93.72±0.3
6	94.08±0.1	94.95±0.1	95.03±0.1	95.05±0.1	1488	87.87±1.6	88.21±1.0	88.23±1.3	88.6±1.3
11	85.44±0.8	89.35±0.9	89.03±0.9	89.38±1.0	1489	88.17±0.3	88.55±0.3	88.58±0.3	88.61±0.3
12	95.28±0.4	97.1±0.3	96.73±0.4	97.09±0.4	1494	85.86±0.8	87.06±0.6	86.63±0.7	86.98±0.7
13	73.12±1.5	72.63±1.7	72.81±1.9	72.95±1.7	1499	90.53±1.1	91.67±1.3	91.49±1.4	91.98±1.1
23	52.81±1.2	52.78±1.2	52.37±1.4	52.85±1.1	1515	83.06±1.5	67.94±2.1	76.34±1.4	83.06±1.5
30	91.9±1.5	91.89±1.4	92.49±1.8	92.45±1.8	1566	59.35±1.2	100.0±0.0	99.99±0.0	100.0±0.0
31	74.51±1.0	74.19±1.3	74.41±1.0	74.55±1.0	1591	81.28±0.1	81.24±0.1	78.69±0.1	81.31±0.1
54	74.4±1.1	78.78±1.0	78.73±0.9	78.98±1.0	4134	79.03±0.6	77.33±0.4	76.19±0.5	78.86±0.9
55	84.5±2.1	85.07±1.9	84.79±2.2	85.69±1.7	4137	90.24±0.4	90.24±0.4	90.24±0.4	90.24±0.4
60	84.36±0.4	85.92±0.4	85.65±0.4	85.92±0.4	4534	94.73±0.3	95.36±0.2	94.89±0.2	95.29±0.2
181	61.14±0.9	61.11±1.1	61.5±0.8	61.47±1.0	4538	63.42±0.5	61.76±0.3	61.97±0.4	63.42±0.5
201	84.18±0.2	83.96±0.2	83.2±0.2	84.18±0.2	40498	65.86±0.6	66.67±0.6	66.99±0.5	67.05±0.5
299	74.31±2.2	82.99±1.7	83.6±1.9	83.84±1.6	40664	92.92±0.6	94.27±0.6	94.67±0.6	94.71±0.5
336	82.42±1.7	81.96±1.5	82.48±1.2	82.21±1.5	40670	94.56±0.3	93.1±0.4	93.58±0.5	94.58±0.3
346	42.62±6.6	34.7±6.2	35.35±8.1	39.27±7.1	40677	71.61±0.7	71.82±0.6	71.59±0.7	71.74±0.6
380	62.89±4.9	62.89±4.9	62.36±5.3	62.89±4.9	40685	99.52±0.0	99.36±0.1	99.51±0.0	99.52±0.0
446	88.52±1.8	98.92±0.7	99.03±0.8	99.07±0.7	40687	75.27±1.0	74.82±0.9	74.69±1.1	75.03±0.9
1042	92.18±0.3	90.22±0.3	90.73±0.3	92.12±0.4	40701	89.44±0.2	88.9±0.3	89.14±0.3	89.41±0.3
1049	89.76±0.5	90.07±0.5	89.66±0.4	90.15±0.4	40713	98.42±0.1	98.42±0.1	98.42±0.1	98.42±0.1
1067	85.85±0.6	86.12±0.5	86.38±0.5	86.3±0.5	40900	98.58±0.1	98.58±0.1	98.58±0.1	98.58±0.1
1083	65.99±2.7	60.29±3.0	59.99±2.7	65.67±2.9	40910	98.32±0.2	98.32±0.2	98.32±0.2	98.32±0.2
1084	81.44±1.5	88.8±1.8	79.64±1.8	88.53±2.2	40971	23.61±1.7	23.91±1.2	24.11±1.4	24.24±1.1
1085	54.91±4.4	52.79±3.5	55.06±3.3	54.9±3.8	40975	80.15±1.0	81.88±1.0	84.14±0.9	84.13±0.9
1086	88.21±1.0	88.19±1.3	88.31±1.2	88.74±1.1	40978	90.54±0.3	92.33±0.3	91.93±0.3	92.27±0.3
1087	88.21±1.0	88.19±1.3	88.31±1.2	88.74±1.1	40981	86.49±1.2	86.39±1.2	86.7±0.8	86.87±1.0
1088	79.33±1.5	76.01±1.7	74.55±1.6	78.77±1.6	40982	76.23±0.6	74.86±0.6	75.63±0.6	76.15±0.6
1128	95.24±0.5	97.25±0.4	89.21±1.0	97.25±0.4	40983	94.66±0.2	94.66±0.2	94.66±0.2	94.66±0.2
1130	91.96±0.4	91.96±0.4	91.96±0.4	91.96±0.4	40984	92.24±0.4	93.8±0.5	93.49±0.4	93.79±0.5
1134	96.64±0.4	98.62±0.2	95.83±0.4	98.62±0.2	40994	91.36±0.8	91.36±0.8	91.36±0.8	91.36±0.8
1138	92.09±0.4	92.09±0.4	92.09±0.4	92.09±0.4	41027	81.42±0.2	81.22±0.1	81.32±0.2	81.41±0.2
1139	95.08±0.3	95.08±0.3	95.08±0.3	95.08±0.3	41142	71.69±0.4	72.72±0.4	71.07±0.5	72.63±0.5
1142	96.0±0.2	96.0±0.2	96.0±0.2	96.0±0.2	41143	80.86±0.5	79.86±0.8	80.02±0.5	80.8±0.6
1146	95.55±0.2	95.55±0.2	95.55±0.2	95.55±0.2	41144	74.83±0.9	61.5±0.8	63.26±0.9	74.83±0.9
1161	95.98±0.4	96.78±0.3	90.33±1.3	96.78±0.3	41145	74.15±0.5	72.37±0.4	72.48±0.5	74.15±0.5
1233	48.02±1.4	57.79±1.1	36.73±2.0	57.79±1.1	41146	92.27±0.2	92.53±0.2	92.76±0.2	92.75±0.2
1441	86.07±2.3	86.24±2.1	86.73±1.9	86.71±2.0	41156	84.84±0.5	84.49±0.5	83.7±0.8	84.78±0.5
1448	79.35±1.7	79.91±2.1	80.55±1.8	80.5±1.5	41157	71.96±3.5	71.24±3.7	66.25±6.1	72.43±3.3
1450	69.01±2.8	69.56±3.2	70.22±3.8	70.21±3.0	41158	91.4±0.4	87.65±0.5	86.56±0.5	91.4±0.4
1457	30.87±1.6	21.19±1.1	17.82±1.0	30.87±1.6	41159	79.89±0.2	62.63±0.4	60.59±0.4	79.89±0.2
1461	89.34±0.1	89.24±0.1	89.17±0.1	89.31±0.1	41163	94.79±0.2	94.39±0.2	93.45±0.3	94.72±0.3
1464	77.83±1.3	78.25±1.2	77.86±1.2	78.33±1.3	41164	67.57±0.5	60.38±0.4	60.18±0.5	67.57±0.5
1465	40.71±4.9	40.14±6.0	39.68±4.7	40.75±5.1	41165	39.29±0.6	36.74±0.5	27.52±0.4	39.29±0.6
1468	89.33±0.7	91.17±0.6	87.27±0.9	91.1±0.6	41166	64.62±0.1	62.44±0.2	62.33±0.2	64.62±0.1
1475	60.92±0.4	60.69±0.4	60.77±0.4	60.97±0.4	41168	69.23±0.1	68.24±0.2	68.8±0.2	69.23±0.1
1477	98.41±0.1	98.97±0.1	99.19±0.0	99.19±0.0	41972	88.9±0.2	91.88±0.3	92.48±0.2	92.46±0.2
1479	55.84±1.4	92.21±0.5	92.45±0.6	92.57±0.6	42809	73.38±0.9	64.89±1.8	58.43±1.9	73.38±0.9
1485	68.3±1.1	57.5±0.9	57.84±1.1	68.3±1.1	42810	74.59±0.5	68.81±0.5	63.83±0.7	74.59±0.5

538

C Training Times

All units are in seconds.

539

540

id	RF	LDA RF	PCA RF	SORF	id	RF	LDA RF	PCA RF	SORF
3	42±3	25±2	116±8	80±31	1487	2±0	2±0	2±0	1±0
6	330±25	461±37	4244±311	3782±919	1488	1±0	1±0	3±0	2±0
11	4±0	5±1	11±1	8±1	1489	37±3	41±4	105±8	75±28
12	74±6	43±3	372±29	230±116	1494	13±1	7±1	36±3	20±4
13	5±0	5±0	16±1	7±3	1499	1±0	1±0	2±0	1±0
23	23±2	33±2	110±8	71±23	1515	91±6	63±4	158±11	234±55
30	6±1	8±2	31±9	11±9	1566	113±11	3±1	33±3	53±28
31	24±2	12±1	63±5	40±10	1591	2970±199	1980±111	8912±493	6987±1894
54	13±1	13±1	46±3	30±6	4134	849±49	301±20	1782±116	1739±559
55	1±0	1±0	3±0	1±0	4137	1287±281	1159±176	1269±225	1221±991
60	107±7	54±4	338±24	258±74	4534	110±7	57±4	302±24	133±27
181	18±1	34±3	97±7	64±12	4538	333±19	311±18	1339±83	1212±282
201	227±15	232±19	1603±102	758±201	40498	90±6	131±9	419±31	332±102
299	13±1	10±1	53±4	48±12	40664	15±1	18±2	58±5	48±15
336	2±0	2±0	7±1	3±1	40670	125±7	48±3	373±21	317±108
346	1±0	1±0	2±0	1±0	40677	58±4	89±8	305±22	280±71
380	3±0	2±0	5±1	3±2	40685	104±7	102±7	497±30	347±205
446	1±0	1±0	2±0	2±1	40687	11±1	13±1	63±5	39±14
1042	258±21	101±8	743±53	577±209	40701	29±3	17±1	94±8	36±8
1049	13±1	6±1	33±4	13±3	40713	1±0	1±0	1±0	1±0
1067	18±2	11±1	49±5	24±6	40900	2±0	2±0	2±0	2±0
1083	453±30	334±22	460±37	565±179	40910	11±1	11±1	11±1	9±1
1084	165±13	86±5	180±12	150±31	40971	29±2	41±3	226±14	158±41
1085	329±14	195±9	324±18	167±31	40975	14±1	24±3	68±5	46±9
1086	276±29	236±19	297±32	219±36	40978	232±18	95±6	250±21	142±19
1087	267±26	228±17	283±32	212±34	40981	8±1	6±0	23±2	16±5
1088	1009±59	683±33	1009±52	1343±231	40982	40±2	46±3	192±14	125±27
1128	467±34	245±14	697±48	409±50	40983	1±0	1±0	1±0	0±0
1130	123±10	123±8	127±12	128±93	40984	23±1	24±2	101±6	47±6
1134	327±19	246±18	468±34	290±39	40994	1±0	1±0	2±1	2±2
1138	125±12	125±8	129±11	53±9	41027	291±20	698±394	1645±710	892±456
1139	121±7	125±8	128±10	100±10	41142	1165±81	383±23	2339±131	2061±597
1142	119±7	122±7	124±7	98±8	41143	112±7	41±2	304±20	228±61
1146	123±9	125±9	126±10	51±4	41144	213±14	64±5	657±47	481±136
1161	359±24	252±18	553±48	320±55	41145	481±33	129±9	1146±75	774±204
1233	596±33	216±14	773±55	800±232	41146	44±3	30±2	148±10	77±21
1441	1±0	0±0	2±1	1±0	41156	59±4	37±3	191±13	115±37
1448	2±0	1±0	5±0	2±0	41157	35±3	14±1	35±2	21±7
1450	2±0	1±0	5±0	2±1	41158	276±16	102±6	782±42	532±101
1457	1314±73	1463±95	1587±101	2992±526	41159	12052±543	5404±391	28376±1663	18174±3579
1461	312±25	181±12	835±69	319±188	41163	2750±135	915±54	6448±420	8550±1119
1464	5±0	7±1	15±1	7±2	41164	6031±379	1532±102	10831±590	16058±3047
1465	2±0	3±0	6±1	3±1	41165	9604±545	5144±288	22115±1081	25829±5586
1468	217±10	119±6	648±35	706±139	41166	4341±275	3766±230	29830±1683	19960±3564
1475	226±17	212±16	1112±71	886±262	41168	4043±233	2314±155	14659±828	10002±3786
1477	434±31	166±13	1468±112	856±181	41972	448±25	175±10	1513±85	1023±235
1479	106±11	13±1	87±6	98±21	42809	1567±97	553±35	1138±74	1276±304
1485	270±15	76±6	725±52	392±116	42810	8424±462	3267±223	10019±580	12821±3749

541

D Built Trees

Average number of tree types built by SORF on the different datasets. The configured maximum per tree types was 100.

id	Standard Trees	LDA Trees	PCA Trees	Total	id	Standard Trees	LDA Trees	PCA Trees	Total
3	31±11	48±23	48±23	127±43	1487	15±0	13±0	13±0	41±0
6	27±5	43±28	83±26	153±12	1488	31±7	28±7	30±6	89±13
11	35±8	41±13	40±10	116±18	1489	34±8	47±20	40±18	121±33
12	32±7	60±22	48±28	140±47	1494	36±13	40±13	36±10	112±27
13	27±15	32±15	24±14	83±36	1499	26±7	30±14	34±11	90±26
23	47±20	44±21	41±17	133±42	1515	74±20	76±23	75±19	224±47
30	18±4	19±11	24±12	61±23	1566	32±19	43±26	46±30	122±62
31	40±7	39±10	41±14	120±16	1591	59±32	75±23	42±12	176±51
54	31±5	48±13	44±12	123±17	4134	63±24	42±18	60±23	165±41
55	31±10	32±12	28±11	91±20	4137	36±33	34±31	28±21	98±77
60	36±8	48±19	57±20	141±32	4534	27±6	36±11	27±9	90±10
181	43±10	48±10	41±10	132±17	4538	83±21	54±17	57±22	194±21
201	40±13	33±12	37±11	110±33	40498	40±12	61±19	52±23	152±36
299	58±15	62±19	66±20	186±42	40664	51±20	58±19	52±24	161±44
336	34±14	30±13	32±11	95±30	40670	72±21	43±19	55±22	171±43
346	19±5	17±4	18±4	53±9	40677	63±21	63±21	61±17	187±38
380	33±34	28±30	27±24	88±76	40685	61±35	35±27	50±36	145±75
446	36±16	42±20	41±20	119±47	40687	35±20	41±23	47±20	123±44
1042	64±18	34±8	51±21	149±41	40701	32±10	25±9	25±9	82±18
1049	23±8	25±8	24±8	72±15	40713	28±1	28±1	28±1	83±3
1067	30±12	31±15	32±11	93±25	40900	27±1	27±1	27±1	82±3
1083	54±18	43±13	37±14	135±36	40910	27±1	27±1	27±1	82±4
1084	34±8	46±13	30±9	110±22	40971	55±15	56±18	53±16	163±33
1085	20±6	18±5	20±5	59±10	40975	30±5	35±9	49±13	114±14
1086	28±6	30±8	25±5	83±12	40978	24±2	26±6	24±3	74±7
1087	28±6	30±8	25±5	83±12	40981	44±16	37±11	44±16	125±31
1088	56±18	58±14	38±11	152±23	40982	51±16	44±13	44±12	139±25
1128	26±4	37±7	28±5	92±10	40983	15±0	13±0	13±0	41±0
1130	34±22	34±22	33±18	101±62	40984	26±4	40±9	31±5	98±8
1134	26±4	32±4	27±4	85±9	40994	55±37	50±31	58±37	162±98
1138	15±1	13±1	14±2	42±4	41027	26±8	26±5	35±5	87±9
1139	27±2	27±2	27±2	81±6	41142	59±21	66±22	48±17	172±46
1142	27±2	27±2	27±2	81±6	41143	55±19	46±17	49±17	149±37
1146	15±0	13±0	13±0	41±0	41144	71±19	39±14	46±17	156±34
1161	27±5	34±7	24±3	86±11	41145	56±20	33±10	40±16	130±29
1233	57±23	73±19	39±17	169±44	41146	29±9	35±12	36±11	99±23
1441	20±5	19±5	21±5	60±8	41156	39±8	40±18	41±17	119±33
1448	17±4	24±8	23±6	63±12	41157	28±12	22±7	24±10	74±22
1450	30±20	22±13	25±12	77±33	41158	54±9	35±6	44±13	133±15
1457	95±17	55±22	61±27	210±35	41159	69±9	24±8	30±12	123±14
1461	27±18	21±5	24±18	73±36	41163	65±35	56±35	97±13	218±31
1464	29±13	28±13	25±9	82±24	41164	93±13	88±13	84±21	265±37
1465	30±13	28±11	30±12	87±23	41165	99±2	80±19	55±24	234±31
1468	43±21	79±21	80±19	202±33	41166	96±6	42±11	49±14	187±18
1475	48±19	60±22	59±20	167±41	41168	64±27	28±7	46±31	138±14
1477	30±3	33±13	46±11	109±22	41972	26±4	29±10	56±15	112±14
1479	24±7	69±19	73±20	166±31	42809	53±16	32±17	24±9	109±30
1485	57±17	27±12	30±13	114±28	42810	69±19	55±21	52±20	176±51

542

543

544

545

E Code

546

Please find the code for all the experiments and result resources of this paper at <https://anonymous.4open.science/r/sorf-2B66/>.

547

548