



Published in final edited form as:

Int Conf Big Data Smart Comput. 2021 January ; 2021: . doi:10.1109/bigcomp51126.2021.00023.

MP-Boost: Minipatch Boosting via Adaptive Feature and Observation Sampling

Mohammad Taha Toghani,

Department of Electrical and Computer Engineering, Rice University, Houston, TX 77005 USA.

Genevera I. Allen

Departments of Electrical and Computer Engineering, Statistics, and Computer Science, Rice University, Houston, TX 77005 USA, and Jan and Dan Duncan Neurological Research Institute, Baylor College of Medicine, Houston, TX 77030 USA

Abstract

Boosting methods are among the best general-purpose and off-the-shelf machine learning approaches, gaining widespread popularity. In this paper, we seek to develop a boosting method that yields comparable accuracy to popular AdaBoost and gradient boosting methods, yet is faster computationally and whose solution is more interpretable. We achieve this by developing MP-Boost, an algorithm loosely based on AdaBoost that learns by adaptively selecting small subsets of instances and features, or what we term *minipatches* (MP), at each iteration. By sequentially learning on tiny subsets of the data, our approach is computationally faster than other classic boosting algorithms. Also as it progresses, MP-Boost adaptively learns a probability distribution on the features and instances that upweight the most important features and challenging instances, hence adaptively selecting the most relevant minipatches for learning. These learned probability distributions also aid in interpretation of our method. We empirically demonstrate the interpretability, comparative accuracy, and computational time of our approach on a variety of binary classification tasks.

Keywords

Minipatch Learning; AdaBoost; Adaptive Observation Selection; Adaptive Feature Selection; Internal Validation

1. Introduction

Boosting algorithms adaptively learn a series of weak learners that overall yield often state-of-the-art predictive accuracy, but are computationally slow. Huge datasets with many instances and features (e.g., data recorded from sensors, texts, and images) highlight this "slow learning" behavior. For example, AdaBoost [1] and gradient boosting [2] suffer from a slow training speed since they try to achieve a proper performance neglecting the size of the data. Further, most boosting methods are "black-box", meaning that their

interpretation [3] lacks transparency and simplicity. Identification of features and examples with the most impact helps to obtain a more interpretable model [4], [5]. In this paper, our goal is to develop an AdaBoost-based algorithm that learns faster computationally and also yields interpretable solutions.

We propose to achieve this by adaptively sampling tiny subsets of both instances and features simultaneously, something we refer to as a *minipatch learning* (Fig. 1). Subsampling is widely used in machine learning and has been shown to have both computational and predictive advantages. For instance, bagging [6] uses the bootstrap technique [7] to reduce the dependency of weak learners to the training data. Random forest [8], as a specific type of bagging, reduces the number of efficient features in each weak learner as well. The computational advantages of the ensemble of uniform random minipatches have been investigated in [9].

There is preliminary evidence that uniform minipatch selection yields implicit regularization [10], [11]; however, dropout is not precisely what we refer to as minipatch learning. In fact, the combination of minibatch selection (i.e., stochastic optimization) and dropout in the first layer lies in the category of minipatch learning algorithms. Minipatch sampling not only can accelerate ensemble algorithms, but it can also implicitly regularize them, i.e., reduce the prediction's sensitivity to individual instances or features.

Subsampling can speed up iterative algorithms; however, uniform sampling can appropriately be modified with an adaptive procedure as the importance and difficulty of the observations and features vary. To give an example, in classification problems, observations closer to the decision boundary play the critical role in the ultimate model [12]. Moreover, we can use adaptive sampling to interpret the final results. For instance, SIRUS algorithm [5], [13] suggests how to leverage the frequency of splits in random forest trees to generate an interpretable model with fewer splits. Our MP-Boost algorithm will incorporate the advantages of adaptivity in order to learn distributions on the observations and features.

More closely related to our work, several have proposed to employ sampling schemes on the top of AdaBoost. Works along this line usually target either subsampling the features or observations adaptively. Exploiting bandit algorithms, adaptive feature selection methods [14], [15] have been suggested to reduce the number of effective features accessed by each weak learner. Similarly, Tasting [16] and Laminating [17], [18] propose score-based feature selection algorithms in each iteration of AdaBoost. They subsample features in each iteration adaptively but either utilize all instances to train weak learners or subsample them uniformly.

On the other hand, algorithms like MadaBoost [19] and FilterBoost [20] suggest adaptive oracle subsampling of the observations. However, they do not subsample features and the size of selected samples increases per iteration, thus get slower during their progress. In contrast, the algorithm that we will propose exploits the effect of minipatch selection; in this way, it will take the importance of both the observations and features into account.

A series of algorithms have been proposed to reduce the computational complexity of gradient boosting. For example, stochastic gradient boosting (SGB) [21] suggests

subsampling the observations randomly. XGBoost [22] by introducing a novel structure for decision trees, a fast training algorithm, as well as several other modifications such as features subsampling techniques through a pre-sorting and histogram-based algorithm, extensively optimizes the computational complexity of gradient boosting.

In this regard, LightGBM [23] and CatBoost [24] subsample observations adaptively proportionate to their gradient values and use an algorithm called exclusive feature bundling (EFB) to reduce the number of effective features by categorizing them. In contrast, MP-Boost will be designed to learn a probability distribution on features gradually during its progress instead of categorizing them initially. The minimal variance sampling (MVS) [25] algorithm is proposed to select the observations according to their gradient values provably with a minimal variance; however, it lacks subsampling over the features. Note that MP-Boost will be designed based on AdaBoost; thus, its adaptive observation selection will entirely be different from algorithms that are designed based on gradient boosting.

The remainder of the paper is organized as follows. In Section 2, stating the problem formulation, we present MP-Boost. In Section 3, we investigate the efficacy of the adaptive subsampling, interpretability of the proposed algorithm, and compare the generalization accuracy of MP-Boost with that of AdaBoost, gradient boosting, and random forest. We end with discussing and concluding remarks in Section 4.

2. MP-Boost

Our goal is to develop a boosting algorithm utilizing adaptive sampling of features and observations that enhances both scalability and interpretability.

2.1. MP-Boost Algorithm

We begin by focusing on binary classification tasks. Let our data be $X \subset \mathbb{R}^{N \times M}$ for N observations or instances and M features; for each instance, we observe a label, (\mathbf{x}_i, y_i) with $y_i \in \{-1, +1\}$. We seek to learn a classifier $y_i = \text{sgn}(F(\mathbf{x}_i))$.

To achieve this, we propose an adaptive sampling based version of boosting inspired by AdaBoost. Our method relies on learning a weak learner from a tiny subset of observations and features at each iteration. We call this tiny subset a *minipatch*, termed based on the use of "patches" in image processing, and minibatches as small subsamples of observations commonly used in machine learning. Our approach is to take an ensemble of minipatches, or minipatch learning as shown in Figure 1, where each minipatch is sampled adaptively. Formally, we define a minipatch by $(\mathbf{X}_{\mathcal{R}}, \mathcal{E}, y_{\mathcal{R}})$, where \mathcal{R} is a subset of observations with size n , and \mathcal{E} is a subset of features with size m . By learning weak learners on tiny subsets or minipatches, our boosting algorithm will have major computational advantages for large N and/or large M datasets.

We define \mathcal{H} to be the *class of weak learners* [26], where each $h \in \mathcal{H}$ is a function $h: \mathbb{R}^m \rightarrow \{-1, +1\}$. Our algorithm is generic to the type of weak learners, which can be either simple or expressive. However, we select decision trees [27] as the default weak

learner in MP-Boost. We consider both depth- k trees as well as saturated trees that are split until each terminal leaf consists of samples from the same class.

The core of our algorithm uses adaptive sampling of observations to achieve the adaptive slow learning properties [28] of the AdaBoost algorithm. Similar to [25] for gradient boosting, MP-Boost subsamples observations according to an adaptive probability distribution. Let p be the probability distribution on observations (i.e., $\sum_{i=1}^N p_i = 1$) and initially set p to be uniform ($\mathbf{U}_{[N]}$). We define $\text{Sample}(N, n, p)$ as sampling a subset of $[N]$ of size n according to the probability distribution p without replacement.

Let $F: \mathbb{R}^M \rightarrow \mathbb{R}$ be the *ensemble function*. Our algorithm selects a minipatch, trains a proper weak learner on it, and computes the summation of weak learners, $F(\mathbf{x}_i)^{(t)} = \sum_{k=1}^t h^{(k)}(\mathbf{x}_i)_{\mathcal{C}(k)}$. Misclassified samples are more difficult to be learned, so we need to increase their probabilities to be sampled more frequently. Let $\mathcal{L}: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}^+$ be a function that measures the similarity between the ensemble outputs and labels, i.e., positive yF yields smaller $\mathcal{L}(y, F)$ and vice versa. MP-Boost assigns a probability proportional to $\mathcal{L}(y_i, F(\mathbf{x}_i))$ to the i^{th} observation. Table 1 shows the choices for function \mathcal{L} inspired by the weighting function in AdaBoost and LogitBoost [29]. Note that "Soft functions" are sensitive to the ensemble output value, while "Hard functions" merely care about its sign.

Full-batch boosting algorithms reweight all of the observations and train a new weak learner on their weighted average in each iteration [1], [2], [29]. In contrast, stochastic algorithms use each sample's frequency to take the effect of its weight into account [19], [20], [25]. We update the probability of the observations similar to FilterBoost [20]. However, FilterBoost increases n during its progress and uses negative sampling, and is thus slower than ours.

Another major goal is to increase the interpretability of boosting approaches. To accomplish this, we also propose to adaptively select features that are effective for learning. Similar to p , let q be the probability distribution on features. Our algorithm requires a criterion to compute the importance of the selected m features based on the structure of \mathcal{H} . There exist several choices for computing features importance based on \mathcal{H} . Some of these inspection techniques are model agnostic, hence proper for MP-Boost to incorporate weak learners from different classes. For example, the permutation importance method [8], [30] shuffling each feature infers its importance according to the difference in the prediction score.

Nevertheless, specific metrics like *impurity reduction score* [31] are defined for decision trees. We utilize this quantity to define our probability distribution over the features. Let \mathcal{I}^h denote the normalized feature importance vector for a weak learner h , wherein each entry determines the relative importance of the corresponding feature compared to other features in the minipatch. In each iteration, q is updated through computing the weighted average of q and \mathcal{I} according to a momentum μ . The hyperparameter $\mu \in (0, 1)$ determines the ratio of exploration vs. exploitation. MP-Boost only modifies the probability of features inside the minipatch, in each iteration.

Algorithm 1

MP-Boost

MP-Boost ($\mathbf{X}, \mathbf{y}, n, m, \mu$)

Initialization ($t = 0$):

$p^{(1)} = \mathbf{U}_{[N]}$ // observation probabilities

$q^{(1)} = \mathbf{U}_{[M]}$ // feature probabilities

$F^{(1)}(\mathbf{x}_i) = 0, \forall i \in [N]$ // ensemble output

$G^{(1)}(\mathbf{x}_i) = 0, \forall i \in [N]$ // out-of-patch output

while Stopping – Criterion($\text{oop}^{(t)}$) not met **do** $t \leftarrow t + 1$

1) **Sample a minipatch:**

a) $\mathcal{R}^{(t)} = \text{Sample}(N, n, p^{(t)})$ // select n instances

b) $\mathcal{C}^{(t)} = \text{Sample}(M, m, q^{(t)})$ // select m features

c) $(\mathbf{X}^{(t)}, \mathbf{y}^{(t)}) = (\mathbf{X}_{\mathcal{R}^{(t)}}, \mathbf{y}_{\mathcal{R}^{(t)}})$ // minipatch

2) **Train a weak learner on the minipatch:**

a) $h^{(t)} \in \mathcal{H}$: weak learner trained on $\mathbf{X}^{(t)}, \mathbf{y}^{(t)}$

3) **Update outputs:**

a) $F^{(t)}(\mathbf{x}_i) = F^{(t-1)}(\mathbf{x}_i) + h^{(t)}(\mathbf{x}_i)_{\mathcal{C}^{(t)}}, \forall i \in [N]$

4) **Update probability distributions:**

a) $p_i^{(t+1)} = \frac{\mathcal{L}(y_i, F^{(t)}(\mathbf{x}_i))}{\sum_{k=1}^N \mathcal{L}(y_k, F^{(t)}(\mathbf{x}_k))}, \forall i \in [N]$

b) $q_j^{(t+1)} = (1 - \mu)q_j^{(t)} + \mu r \mathcal{I}_j^{h^{(t)}}, j \in \mathcal{C}^{(t)}$

 where, $r = \sum_{j \in \mathcal{C}^{(t)}} q_j^{(t)}$

5) **Out-of-Patch Accuracy:**

a) $G^{(t)}(\mathbf{x}_i) = G^{(t-1)}(\mathbf{x}_i) + h^{(t)}(\mathbf{x}_i)_{\mathcal{C}^{(t)}}, \forall i \notin \mathcal{R}^{(t)}$

b) $\text{oop}^{(t)} = \frac{1}{N} \sum_{i=1}^N \mathbb{1}\{\text{sgn}(G^{(t)}(\mathbf{x}_i)) = y_i\}$

end while

Return $\text{sgn}(F^{(T)}), p^{(T)}, q^{(T)}$

Finally, many boosting algorithms are designed to run for a fixed number of iterations [2] or use a validation criterion [19], [20], [32] in order to determine when to stop. Internal validation approaches often have better performance and are computationally much faster. For instance, consider the out-of-bag criterion [33] in bagging and random forest that uses internal validation properties without incurring any additional computational cost. Similarly to bagging, our MP-Boost has access to out-of-patch instances, which we can use for internal validation. Therefore, for each sample $i \in [N]$, we accumulate the output of weak learners that do not have it in their minipatch. Thus, we define *out-of-patch output* to be a function $G: \mathbb{R}^M \rightarrow \mathbb{R}$ as follows:

$$G^{(t)}(\mathbf{x}_i) = \sum_{k=1}^t h^{(k)}(\mathbf{x}_i) \mathbb{1}_{\{i \notin \mathcal{R}^{(k)}\}} \quad (1)$$

for an arbitrary \mathbf{x}_i . Accordingly, the out-of-patch accuracy, oop, can easily be quantified.

Out-of-patch accuracy is a conservative estimate of the test accuracy. Hence it can assist MP-Boost to track the progress of the generalization (test) performance internally and decide at which iteration to stop. In a nutshell, observing the oop value, the algorithm finds where it is saturated. Algorithm 2 (in Appendix B) is a heuristic algorithm that takes $\text{oop}^{(t)}$, compares it with its previous values, and finally decides when it becomes saturated. In fact, the stopping algorithm follows a general rule; if the current value of oop increases with some margin, then the algorithm needs more time to improve; otherwise, the generalization performance is saturated.

We put all of this together in a summary of our MP-Boost algorithm in Algorithm 1. Notice here that selecting minipatches (step 1) reduces the computational complexity imposed per iteration, thus improves the scalability. Other variations of AdaBoost usually subsample either features or observations while ours exploits both. Therefore, in addition to the predictive model F , our algorithm learns probability distributions p and q that express the importance of observations and features, respectively. Since learning p, q is a part of the iterative procedure, it does not incur an extra computational cost. In addition, MP-Boost exploits an internal validation that yields an automatic stopping criterion when the algorithm ceases to learn. Hence, steps (4) and (5) of Algorithm 1 highlight the main differences of MP-Boost with other sampling-based boosting algorithms.

2.2. Hyperparameter Tuning

The minipatch size is a crucial hyperparameter of our algorithm. Large n or m slows down weak learners' training but is more likely to yield better performance. Note that m must be large enough to provide a wide range of features for comparison and update features probability properly. Similarly, n must be large enough such that each minipatch represents a meaningful subset of the observations. On the other hand, small n results in a oop value akin to the generalization accuracy. Selecting around ten percent of the observations and features seems to be a proper choice for our algorithm, as evidenced in our empirical studies. Additionally, our studies reveal that the results are fairly robust to small changes in n and m . μ is the other important hyperparameter in our algorithm where a moderate value for it (e.g., $\mu = 0.5$) strikes a balance between exploration and exploitation. While our default hyperparameter settings seem to perform well and are robust in a variety of settings (see Section 3), one could always select these in a datadriven manner using our oop criterion as well.

2.3. Extensions

Our proposed algorithm is initially developed for the binary classification problem. Here, we discuss its extension to *regression* and *multiclass classification* problems. First, for multiclass classification, algorithms like AdaBoost.M2, AdaBoost.MH, and AdaBoost.OC

[34], [35] are proposed as multiclass extensions of AdaBoost. We can incorporate similar modifications to extend MP-Boost as well. Further, [36], [37] have suggested regression extensions like AdaBoost.R2 and AdaBoost.RT to the vanilla AdaBoost. Obviously, for regression, we will have to change our observation probability function and the out-of-patch accuracy, thus we can employ techniques in [33] to address the regression problem. All these can be further extensions to our approach.

3. Experiments

We begin by using an illustrative case study to show how our method works and how it aids interpretability. Next, we compare our algorithm to other popular boosting and tree-based methods, focusing on accuracy and scalability.

3.1. Illustrative Case Study

We use a series of experiments to demonstrate how our algorithm works and show how to interpret the results. Specifically, we ask: how does adaptive sampling boost the performance; how does the out-of-patch accuracy relate to test accuracy and yield a datadriven stopping criterion; how do we interpret the results via the final value of p and q ?

To answer these questions, we focus our investigations on an explicable binary classification task: detecting digit 3 versus 8 in MNIST [38]. This dataset includes handwritten digits as images of size 28×28 . The training data is huge ($N > 10000$) and high-dimensional ($M = 784$). We use cross-validation to tune all hyperparameters, yielding $n = 500$, $m = 30$, and $\mu = 0.5$ as well as the Soft-Logistic function 1 as L .

To measure the effect of adaptive observation and/or feature selection, we train MP-Boost on MNIST(3,8). Then, we turn off the adaptive updates for p and replace it with the uniform distribution and resulting random observation sampling. We do the same for q analogously for the features. Finally, we turn off adaptive sampling for both features and observations, and repeat each experiment 5 times. Figure 2a shows the superiority of joint adaptive sampling of both observations and features in terms of performance on the test data. Further, we compare the training, out-of-patch, and test accuracy curves for MP-Boost in Figure 2b. The dashed line indicates the stopping time of our algorithm based on the oop curve and our stopping heuristic Algorithm 2. To observe the behavior of the three curves, we let MP-Boost continue after the stopping criterion is satisfied. As shown in Figure 2b, and unlike the training curve, the trend in the out-of-patch curve is similar to that of the test curve. These results demonstrate the power that adaptive sampling of both observations and features brings to boost performance as well as the ability to use the oop curve to assess algorithm progress.

Next, we illustrate how to use p and q to interpret the observations and features. First, observations that are difficult to classify are upweighted in p . Hence, we can use p to identify the most challenging samples, yielding a similar type of interpretation commonly employed in support vector machines [12]. To visualize the final value of p , we project observations on a two-dimensional space using PCA with sizes of each observation proportional to p . We show this visualization before and after the training in Figure 3.

Further, we highlight a few of the observations with large values of p , illustrating that these samples are indeed difficult to distinguish between the two classes.

Finally, we show how to use q to find the most important features, and also illustrate how MP-Boost learns these features probability in Figure 4. Here, the color of each pixel (feature) is proportional to its probability, with darker pixels indicating the feature has been upweighted. We expect a sparse representation for q which matches with our result. Moreover, this example clearly shows how to interpret the most relevant features; in Figure 4f, two regions are darker compared to other pixels corresponding with the complementary area for digit 3 versus 8.

3.2. Comparative Empirical Results

Here, we compare the speed and performance of our algorithm with AdaBoost, gradient boosting, and random forest on multiple binary classification tasks. To this end, we select large real datasets (usually $N > 1000$ and $M > 100$) from UCI machine learning repository [39], MNIST [38], and CIFAR-10 [40]. Moreover, we use a sparse synthetic dataset of two high-dimensional cones. The size of datasets is provided in Appendix A.

To be fair, we choose the oracle hyperparameters for every method. To this end, we pick decision trees with different maximum depths ($\text{depth} \in \{1, 2, 3, 4, 5, 6, 7\}$) or depth-saturated trees as weak learners. Note that we use scikit-learn [41] modules to implement our algorithm and all competitors so that all time-based comparisons are fair. For our method, we select from the following choices of hyperparameters: $n \in \{50, 100, 200, 500\}$, $m \in \{5, 10, 15, 20, \sqrt{M}\}$ and $\mu \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$.

For each dataset, we select the best performance of MP-Boost, versus that of AdaBoost, gradient boosting, and random forest constrained to the runtime of MP-Boost. Table 2 shows the best performance of each algorithm within the fixed runtime (MP-Boost training time). Results indicate that MP-Boost achieves a better performance faster than the other three algorithms. We also provide a more comprehensive comparison in Table 3, where we can see that without any runtime constraint, MP-Boost is much faster with a comparable accuracy across a wide variety of datasets.

4. Discussion

In this work, we proposed a new boosting approach using adaptive minipatch learning. We showed that our approach is computationally faster, as well as more interpretable, compared to standard boosting algorithms. Moreover, we showed that our algorithm outperforms AdaBoost and gradient boosting in a fixed runtime and has a comparable performance without any time constraint.

Our approach would be particularly useful for large data settings where both observations and features are large. Further, our MP-Boost algorithm can be particularly fitting to datasets with sparse features or noisy features because we learn the important features. Further, our adaptive observation selection can be used in active learning problems.

We suggested extensions of MP-Boost for multiclass classification and regression problems in this paper. In future work, we plan on comparing them empirically to existing algorithms. Other aspects of future work would be to use this minipatch learning scheme in not only AdaBoost-like methods but also in gradient boosting-based methods. Additionally, one could explore other updating schemes for both the observation and feature probabilities. Theoretical analysis of MP-Boost would help to interpret the performance of our algorithm and its properties. Efficient implementations can improve the speed of current MP-Boost for different schemes. Likewise, efficient memory allocation is another essential work ahead of this project, where reducing the required hardware enables industry-scale problems to exploit properties in MP-Boost.

Acknowledgments

The authors acknowledge support from NSF DMS-1554821, NSF NeuroNex-1707400, and NIH 1R01GM140468.

Appendix A.

Dataset

According to our specific problem, we select datasets listed in Table 4 from UCI machine learning repository [39], MNIST [38], and CIFAR [40]. Note that we consider multiclass datasets like MNIST and CIFAR-10 and select two classes that are harder to be distinguished like digits 3&8 from MNIST or "Truck & Car" classes from CIFAR-10. Additionally, we partition all classes into two for different multiclass datasets, like "Odd & Even" digits in MNIST. About 20% of each dataset is selected as the test data. "Cones" is a sparse synthetic dataset of two originally 10dimensional cones. Adding 490 noise features to them, we transform data points to \mathbb{R}^{500} .

Appendix B.

Stopping Criterion Algorithm

The trend in oop is an approximation of the generalization accuracy; however, there are numerical oscillations in it. To make the stopping algorithm (Algorithm 2) robust to such numerical vibrations, instead of the largest previous value of oop, it keeps the k highest previous values of oop in a list, \mathcal{A} , and compares the $\text{oop}^{(t)}$ with their minimum. If the oop value does not improve after k sequential iterations, then MP-Boost halts.

Algorithm 2

Stopping-Criterion

$\gamma = 1 + \frac{\log(n)}{N}$ // ratio of the growth must be $\gamma > 1$

$k = \log(N)$ // number of top oop

$v = 0$ // number of successive iterations without progress $\mathcal{A} =$ a list containing k zeros // keeping the top k largest values of oop up to each iteration

Stopping-Criterion ($\text{oop}^{(t)}$)

1) **Detect the best iteration:**

if $\text{oop}^{(t)} > \max(\mathcal{A})$, set $T = t$

2) **Check the stopping criteria:**

if $v > k$ or $t > T_{\max}$, halt Algorithm 1 // T_{\max} : maximum number of iterations

3) **Check the progress of oop:**

$$v = \begin{cases} v + 1 & \text{oop}^{(t)} < \gamma \cdot \min(\mathcal{A}) \\ 0 & \text{otherwise} \end{cases}$$

4) **Update \mathcal{A} :**

if $\text{oop}^{(t)} > \min(\mathcal{A})$, then replace the minimum entry with $\text{oop}^{(t)}$

References

- [1]. Freund Y and Schapire RE, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of computer and system sciences*, vol. 55, no. 1, pp. 119–139, 1997.
- [2]. Friedman JH, "Greedy function approximation: a gradient boosting machine," *Annals of statistics*, pp. 1189–1232, 2001.
- [3]. Lipton ZC, "The mythos of model interpretability," *Queue*, vol. 16, no. 3, pp. 31–57, 2018.
- [4]. Biau G and Scornet E, "A random forest guided tour," *Test*, vol. 25, no. 2, pp. 197–227, 2016.
- [5]. Bénard C, Biau G, Da Veiga S, and Scornet E, "Interpretable random forests via rule extraction," *arXiv preprint arXiv:2004.14841*, 2020.
- [6]. Breiman L, "Bagging predictors," *Machine learning*, vol. 24, no. 2, pp. 123–140, 1996.
- [7]. Efron B and Tibshirani R, "An introduction to the bootstrap," 1993.
- [8]. Breiman L, "Random forests," *Machine learning*, vol. 45, no. 1, pp. 5–32, 2001.
- [9]. Louppe G and Geurts P, "Ensembles on random patches," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 346–361, Springer, 2012.
- [10]. Srivastava N, Hinton G, Krizhevsky A, Sutskever I, and Salakhutdinov R, "Dropout: a simple way to prevent neural networks from overfitting," *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [11]. LeJeune D, Javadi H, and Baraniuk RG, "The implicit regularization of ordinary least squares ensembles," *arXiv preprint arXiv:1910.04743*, 2019.
- [12]. Cortes C and Vapnik V, "Support-vector networks," *Machine learning*, vol. 20, no. 3, pp. 273–297, 1995.
- [13]. Bénard C, Biau G, Da Veiga S, and Scornet E, "Sirus: making random forests interpretable," *arXiv preprint arXiv:1908.06852*, 2019.
- [14]. Busa-Fekete R and Kegl B, "Bandit-aided boosting," in *OPT 2009: 2nd NIPS Workshop on Optimization for Machine Learning*, 2009.
- [15]. Busa-Fekete R and Kegl B, "Fast boosting using adversarial bandits," in *27th International Conference on Machine Learning (ICML 2010)*, pp. 143–150, 2010.
- [16]. Dubout C and Fleuret F, "Tasting families of features for image classification," in *2011 International Conference on Computer Vision*, pp. 929–936, IEEE, 2011.
- [17]. Dubout C and Fleuret F, "Boosting with maximum adaptive sampling," in *Advances in Neural Information Processing Systems*, pp. 1332–1340, 2011.
- [18]. Dubout C and Fleuret F, "Adaptive sampling for large scale boosting," *The Journal of Machine Learning Research*, vol. 15, no. 1, pp. 1431–1453, 2014.
- [19]. Domingo C, Watanabe O, et al., "Madaboost: A modification of adaboost," in *COLT*, pp. 180–189, Citeseer, 2000.

- [20]. Bradley JK and Schapire RE, "Filterboost: Regression and classification on large datasets," in Advances in neural information processing systems, pp. 185–192, 2008.
- [21]. Friedman JH, "Stochastic gradient boosting," Computational statistics & data analysis, vol. 38, no. 4, pp. 367–378, 2002.
- [22]. Chen T and Guestrin C, "Xgboost: A scalable tree boosting system," in Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining, pp. 785–794, ACM, 2016.
- [23]. Ke G, Meng Q, Finley T, Wang T, Chen W, Ma W, Ye Q, and Liu T-Y, "Lightgbm: A highly efficient gradient boosting decision tree," in Advances in neural information processing systems, pp. 3146–3154, 2017.
- [24]. Dorogush AV, Ershov V, and Gulin A, "Catboost: gradient boosting with categorical features support," arXiv preprint arXiv:1810.11363, 2018.
- [25]. Ibragimov B and Gusev G, "Minimal variance sampling in stochastic gradient boosting," in Advances in Neural Information Processing Systems, pp. 15061–15071, 2019.
- [26]. Mohri M, Rostamizadeh A, and Talwalkar A, Foundations of machine learning. MIT press, 2018.
- [27]. Breiman L, Friedman J, Stone CJ, and Olshen RA, Classification and regression trees. CRC press, 1984.
- [28]. Wyner AJ, Olson M, Bleich J, and Mease D, "Explaining the success of adaboost and random forests as interpolating classifiers," The Journal of Machine Learning Research, vol. 18, no. 1, pp. 1558–1590, 2017.
- [29]. Friedman J, Hastie T, Tibshirani R, et al. , "Additive logistic regression: a statistical view of boosting (with discussion and a rejoinder by the authors)," The annals of statistics, vol. 28, no. 2, pp. 337–407, 2000.
- [30]. Altmann A, Tolos L, Sander i, O., and Lengauer T, "Permutation importance: a corrected feature importance measure," Bioinformatics, vol. 26, no. 10, pp. 1340–1347, 2010. [PubMed: 20385727]
- [31]. Nembrini S, Konig IR, and Wright MN, "The revival of the gini importance?," Bioinformatics, vol. 34, no. 21, pp. 3711–3718, 2018. [PubMed: 29757357]
- [32]. Meijer DWJ and Tax DMJ, "Regularizing adaboost with validation sets of increasing size," 2016 23rd International Conference on Pattern Recognition (ICPR), pp. 192–197, 2016.
- [33]. Breiman L, "Out-of-bag estimation," 1996.
- [34]. Schapire RE, "Using output codes to boost multiclass learning problems," in ICML, vol. 97, pp. 313–321, Citeseer, 1997.
- [35]. Chengsheng T, Huacheng L, and Bing X, "Adaboost typical algorithm and its application research," in MATEC Web of Conferences, vol. 139, p. 00222, EDP Sciences, 2017.
- [36]. Drucker H, "Improving regressors using boosting techniques," in ICML, vol. 97, pp. 107–115, 1997.
- [37]. Solomatine DP and Shrestha DL, "Adaboost. rt: a boosting algorithm for regression problems," in 2004 IEEE International Joint Conference on Neural Networks (IEEE Cat. No. 04CH37541), vol. 2, pp. 1163–1168, IEEE, 2004.
- [38]. LeCun Y, "The mnist database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, 1998.
- [39]. Dua D and Graff C, "UCI machine learning repository," 2017.
- [40]. Krizhevsky A, Hinton G, et al., "Learning multiple layers of features from tiny images," 2009.
- [41]. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg V, Vanderplas J, Passos A, Cournapeau D, Brucher M, Perrot M, and Duchesnay E, "Scikit-learn: Machine learning in Python," Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.

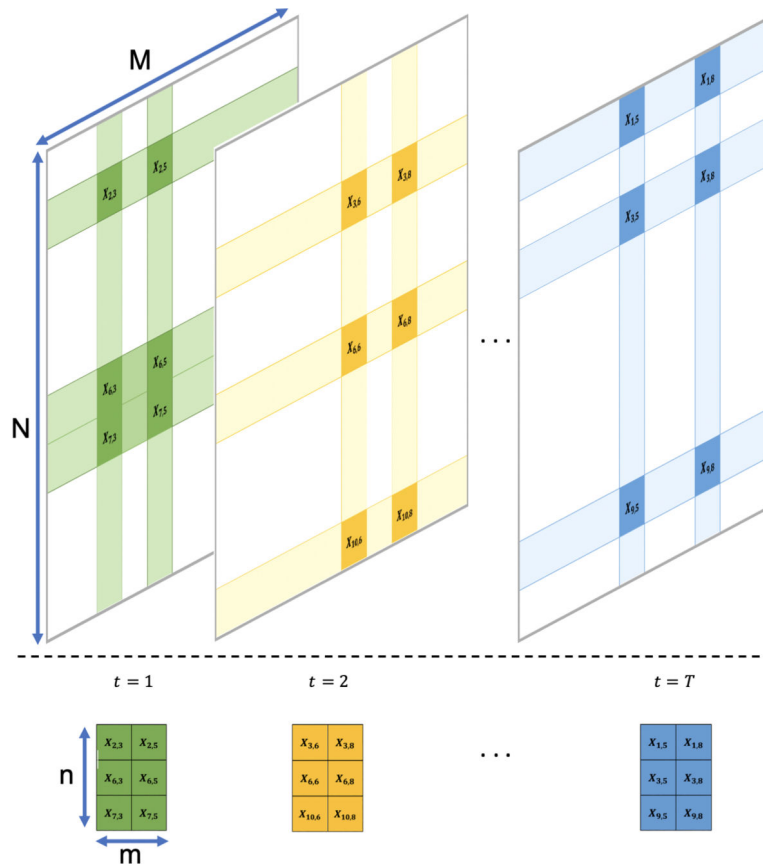


Figure 1: Minipatch Learning-
 Selecting random minipatches of the dataset during the training of an ensemble algorithm.

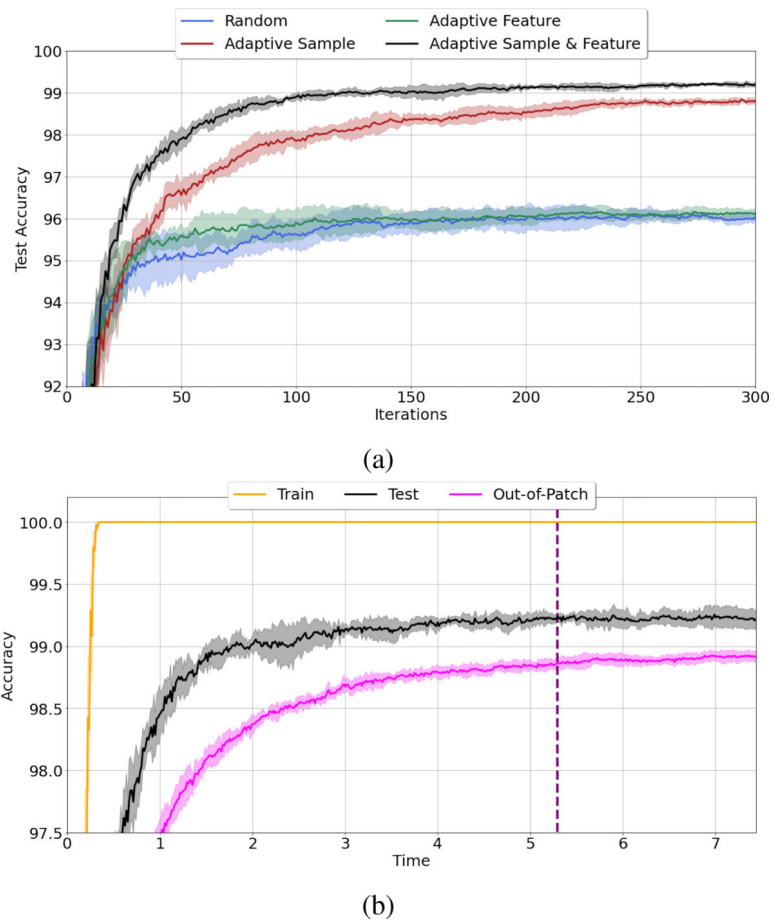


Figure 2:

(a) Effect of MP-Boost adaptive sampling on MNIST (3 vs. 8) test data. (b) Train, test, and out-of-patch accuracy of MNIST (3 vs. 8) using MP-Boost. The dashed line shows the stopping time based on Algorithm 2.

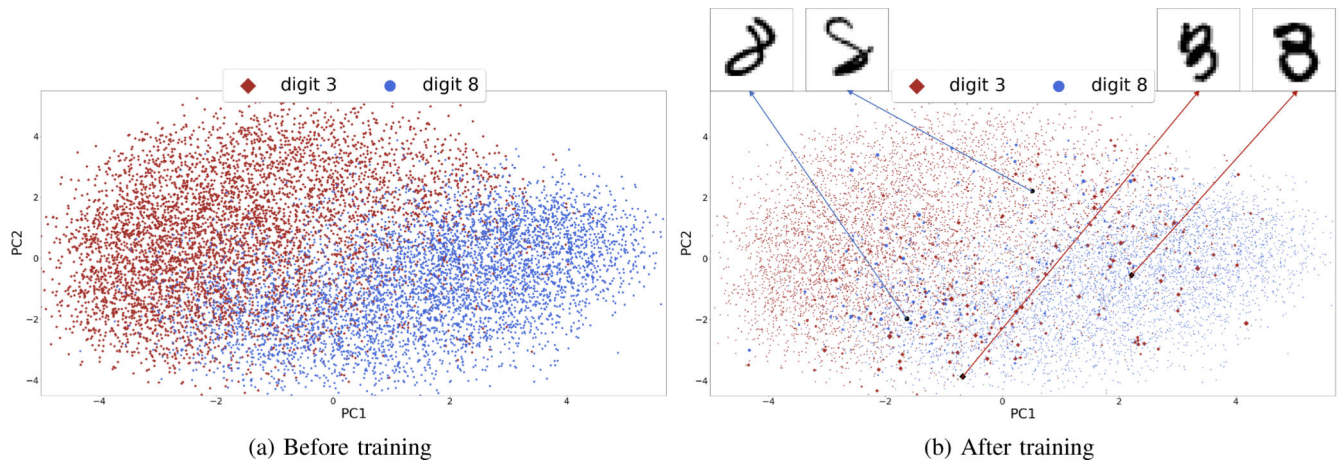


Figure 3: Probability distribution of MNIST (digit 3 vs. 8) samples before and after training MP-Boost-

Samples are Projected onto a 2-dimensional space using PCA. The size of each point (sample) indicates its relative probability. After training, samples close to the boundary have higher probabilities. The displayed samples are visually hard to be learned and have high probabilities.




Figure 4: Probability distribution of MNIST (digit 3 vs. 8) features during the training of MP-Boost-

The color of each pixel (feature) indicates the relative probability with respect to which the feature is selected. As MP-Boost progresses, the probability of ineffective features (pixels in the background) decreases while that of efficient features (pixels shaping the digits) increases. The final distribution resembles a blurred version of digits 3 and 8.

TABLE 1:

$\mathcal{L}(y_i, F(x_i))$ **choices for weighting samples**- \mathcal{L} is a decreasing function with respect to $y_i \times F(x_i)$. Hard functions depend on the sign of $y_i F(x_i)$ that makes MP-Boost less sensitive to outlier samples.

	Exponential	Logistic
Soft	$\exp(-y_i F(x_i))$	$\frac{1}{1 + \exp(y_i F(x_i))}$
Hard	$\exp(-y_i \text{sgn}(F(x_i)))$	$\frac{1}{1 + \exp(y_i \text{sgn}(F(x_i)))}$

Best test accuracy (%) of MP-Boost, AdaBoost, Random Forest, and Gradient Boosting on binary classification tasks during the runtime of MP-Boost.

TABLE 2:

Dataset	MP-Boost	AdaBoost	Random Forest	Gradient Boosting
Cones	100.0 ± 0.0	85.56 ± 0.16	95.89 ± 0.97	94.01 ± 1.23
Hill-Valley	99.45 ± 0.19	96.28 ± 1.47	98.76 ± 0.34	96.69 ± 0.67
Christine	74.27 ± 0.32	69.28 ± 0.09	73.75 ± 0.04	70.27 ± 0.13
Jasmine	80.03 ± 0.76	79.53 ± 0.01	79.47 ± 0.55	79.47 ± 0.08
Philippine	71.01 ± 0.37	69.64 ± 0.01	70.07 ± 0.31	69.98 ± 0.01
SensIT Vehicle	86.23 ± 0.01	83.69 ± 0.01	85.98 ± 0.05	84.24 ± 0.02
Higgs Boson	83.42 ± 0.06	82.52 ± 0.01	83.57 ± 0.08	81.44 ± 0.01
MINIST (3,8)	99.31 ± 0.03	97.49 ± 0.08	98.84 ± 0.03	96.81 ± 0.03
MINIST (O,E)	98.15 ± 0.06	93.78 ± 0.01	97.74 ± 0.04	93.23 ± 0.01
CIFAR-10 (T,C)	73.67 ± 0.39	67.65 ± 0.02	73.08 ± 0.4	67.81 ± 0.15
GAS Drift	99.76 ± 0.04	96.54 ± 0.02	99.64 ± 0.08	96.54 ± 0.02
DNA	97.44 ± 0.07	96.86 ± 0.17	97.44 ± 0.2	96.81 ± 0.07
Volkert	74.14 ± 0.08	71.54 ± 0.01	77.95 ± 0.05	71.9 ± 0.04

Best test accuracy (%) and its corresponding training time (sec) of MP-Boost, AdaBoost, Random Forest, and Gradient Boosting on binary classification tasks.

TABLE 3:

Dataset	Criteria	MP-Boost	AdaBoost	Random Forest	Gradient Boosting
Cones	Test	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0	100.0 ± 0.0
	Time	2.81 ± 0.16	23.63 ± 0.14	14.2 ± 0.06	143.87 ± 1.23
Hill-Valley	Test	99.45 ± 0.19	99.45 ± 0.19	99.17 ± 0.06	98.76 ± 0.34
	Time	0.08 ± 0.01	0.37 ± 0.01	0.14 ± 0.01	2.92 ± 0.08
Christine	Test	74.27 ± 0.32	74.7 ± 0.01	75.32 ± 0.41	76.18 ± 0.03
	Time	2.22 ± 0.2	92.36 ± 0.53	6.06 ± 0.09	108.47 ± 0.51
Jasmine	Test	80.03 ± 0.76	81.43 ± 0.08	80.7 ± 0.7	80.87 ± 0.14
	Time	0.15 ± 0.04	2.4 ± 0.01	0.47 ± 0.00	2.2 ± 0.02
Philippine	Test	71.01 ± 0.37	72.73 ± 0.01	72.3 ± 0.16	73.1 ± 0.47
	Time	0.96 ± 0.08	11.45 ± 0.04	5.78 ± 0.02	74.56 ± 0.91
SensIT Vehicle	Test	86.23 ± 0.01	85.93 ± 0.01	86.88 ± 0.05	87.43 ± 0.01
	Time	25.7 ± 3.03	576.63 ± 19.69	303.6 ± 0.9	730.06 ± 0.75
Higgs Boson	Test	83.42 ± 0.06	82.95 ± 0.01	83.68 ± 0.04	82.74 ± 0.01
	Time	79.31 ± 3.55	270.28 ± 0.51	258.79 ± 3.31	397.28 ± 1.04
MNIST (3,8)	Test	99.31 ± 0.03	99.31 ± 0.02	98.98 ± 0.13	98.96 ± 0.01
	Time	5.25 ± 0.76	81.65 ± 0.53	6.19 ± 0.04	99.5 ± 0.58
MNIST (O,E)	Test	98.15 ± 0.06	98.17 ± 0.01	97.74 ± 0.04	97.95 ± 0.03
	Time	56.34 ± 0.34	711.22 ± 3.98	36.62 ± 0.1	807.53 ± 0.67
CIFAR-10 (T,C)	Test	73.67 ± 0.39	76.12 ± 0.01	74.24 ± 0.37	78.5 ± 0.12
	Time	21.44 ± 4.54	2282.46	53.44 ± 0.23	2224.67 ± 1.63
GAS Drift	Test	99.76 ± 0.04	99.78 ± 0.01	99.66 ± 0.03	99.68 ± 0.03
	Time	1.13 ± 0.07	42.17 ± 0.03	2.42 ± 0.03	64.86 ± 0.1

Dataset	Criteria	MP-Boost	AdaBoost	Random Forest	Gradient Boosting
DNA	Test	97.44 ± 0.07	97.49 ± 0.01	97.75 ± 0.07	97.65 ± 0.03
	Time	0.48 ± 0.06	1.84 ± 0.01	1.08 ± 0.01	0.97 ± 0.01
Volkert	Test	74.14 ± 0.08	75.87 ± 0.3	79.03 ± 0.16	77.66 ± 0.13
	Time	32.82 ± 0.97	541.49 ± 9.13	74.42 ± 0.12	553.35 ± 1.84

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

TABLE 4:

Size of the datasets in Tables 2 and 3

Dataset	N_{train}	M	N_{test}
Cones	20000	500	5000
Hill-Valley	1000	100	200
Christine	4300	1636	1100
Jasmine	2400	144	600
Philippine	4700	308	1100
SensIT Vehicle	78800	100	19700
Higgs Boson	200000	30	50000
MNIST (3 vs. 8)	12000	784	2000
MNIST (Odd vs. Even)	60000	784	10000
CIFAR-10 (Truck vs. Car)	10000	3072	2000
Gas Drift	11100	128	2800
DNA	2600	180	6000
Volkert	46600	180	11700
Fabert	6600	800	1600

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript