
DART: Dropouts meet Multiple Additive Regression Trees

K. V. Rashmi

Department of Electrical Engineering and Computer Science
UC Berkeley

Ran Gilad-Bachrach

Machine Learning Department
Microsoft Research

Abstract

MART (Friedman, 2001, 2002), an ensemble model of boosted regression trees, is known to deliver high prediction accuracy for diverse tasks, and it is widely used in practice. However, it suffers an issue which we call *over-specialization*, wherein trees added at later iterations tend to impact the prediction of only a few instances, and make negligible contribution towards the remaining instances. This negatively affects the performance of the model on unseen data, and also makes the model over-sensitive to the contributions of the few, initially added trees. We show that the commonly used tool to address this issue, that of *shrinkage*, alleviates the problem only to a certain extent and the fundamental issue of over-specialization still remains.

In this work, we explore a different approach to address the problem that of employing *dropouts*, a tool that has been recently proposed in the context of learning deep neural networks (Hinton et al., 2012). We propose a novel way of employing dropouts in MART, resulting in the *DART* algorithm. We evaluate DART on ranking, regression and classification tasks, using large scale, publicly available datasets, and show that DART outperforms MART in each of the tasks, with a significant margin. We also show that DART overcomes the issue of over-specialization to a considerable extent.

1 Introduction

Ensemble based algorithms have been shown to achieve high accuracy for a number of machine learning tasks (Caruana and Niculescu-Mizil, 2006). For ensembles to achieve better accuracy than the individual predictors that they are made of, these predictors need to be accurate but uncorrelated (Breiman, 2001). This helps to increase the accuracy of the model by reducing the sensitivity to specific features or instances that might exist in the individual predictors (Breiman, 2001; Hinton et al., 2012). While some classes of ensemble algorithms such as random forests (Breiman, 2001) learn each predictor in the ensemble independently, boosted ensemble algorithms such as AdaBoost (Freund and Schapire, 1995) and MART (Friedman, 2001, 2002)¹ iteratively add each predictor.

Boosting algorithms add predictors that focus on improving the current model, and this is achieved by modifying the learning problem between iterations. While this guarantees that the added predictor is different than the ones in the ensemble, the new predictors typically focus on a small subset of the problem and hence do not have a strong predictive power when measured on the original problem. This increases the risk of adding models that over-fit specific instances. This is a well-known problem in the context of boosting (Freund, 2001) as well as in MART, which is an ensemble of boosted regression trees. Here, trees added at later iterations tend to impact the prediction of only a few instances, and they make negligible contribution towards the prediction of all the remaining instances. This, in turn, can negatively impact the performance of the algorithm on unseen data by increasing the capacity of the model without making significant improvement in its training error. This also makes the model over-sensitive to the contributions of the few,

Appearing in Proceedings of the 18th International Conference on Artificial Intelligence and Statistics (AISTATS) 2015, San Diego, CA, USA. JMLR: W&CP volume 38. Copyright 2015 by the authors.

¹This algorithm is known by many names, including Gradient TreeBoost, boosted trees, and Multiple Additive Regression Trees (MART). We use the latter to refer to this algorithm.

initially added trees. We call this issue of subsequent trees affecting the prediction of only a small fraction of the training instances *over-specialization*. We discuss this issue in greater detail in Section 2 with an example from a regression task on a real-world dataset.

The most common approach employed to combat the problem of over-specialization in MART is *shrinkage* (Friedman, 2001, 2002). Here, the contribution of each new tree is reduced by a constant value called the *shrinkage factor*. As we will see in Section 2, shrinkage does help in reducing the impact of the first trees, nevertheless, however, as the size of the ensemble increases, the problem of over-specialization reappears.

In this work, we explore a different approach to address the issue of over-specialization in MART. We propose employing *dropouts*, a tool that has been recently proposed in the context of learning deep neural networks (Hinton et al., 2012). In neural networks, dropouts are used to mute a random fraction of the neural connections during the learning process. Therefore, nodes at higher layers of the network cannot rely on a few connections to deliver the information needed for the prediction. This method has contributed significantly to the success of deep neural networks for many tasks including, for example, object classification in images (Krizhevsky et al., 2012).

The technique of dropouts has been used successfully in other learning models (Maaten et al., 2013; Wang and Manning, 2013), for example, in logistic regression (Wager et al., 2013). In these cases, dropouts are used to mute a random fraction of the input features during the training phase. In the context of ensemble of trees, this approach makes them similar to the approach employed by random forests for diversification (Breiman, 2001), wherein each tree in the ensemble is learned (independently) using a different random fraction of the features.

In this paper, we propose a novel way of employing dropouts for ensemble of trees: muting complete trees as opposed to muting features.² We employ this approach in MART and call the resulting algorithm *DART*. We evaluate DART on three different tasks: ranking, regression and classification, using large scale, publicly available datasets. Our results show that DART outperforms MART and random forest in each of the tasks, with significant margins (see Section 4). We note that both MART and random forest are known to be highly successful models for many learning tasks (Caruana and Niculescu-Mizil, 2006), for example, the winners of the ‘Yahoo! learning to rank’ challenge employed a MART model (Chapelle

and Chang, 2011). Therefore, it is both surprising and encouraging that we can squeeze out even higher accuracy out of MART. One of the reasons for the improved performance of DART is that it addresses the issue of over-specialization and results in more balanced contribution from all the trees in the ensemble (see Section 2).

2 Overcoming the Over-specialization in MART

As we briefly discussed in Section 1, boosting, in particular the MART algorithm, suffers from the issue of over-specialization: trees added at later iterations tend to impact the prediction of only a few instances, and make negligible contribution towards the prediction of all the remaining instances. In this section, we will demonstrate this issue and the impact of using dropouts as employed in DART through an example from a regression task on the CTSlice data (see Section 4.2 for a description of the dataset and the task). We note that similar observations were made on the other datasets used in the evaluation (Section 4) as well.

Figure 1 presents the average contribution of the trees in the ensemble, where the average contribution of a tree T is defined to be $|E_x [T(x)]|$ with the expectation taken with respect to the training data. We can see that the MART algorithm (without using shrinkage) starts with a single tree that makes significant contribution and the rest of the trees add negligible contributions. We observed that even if we replace the term $|E_x [T(x)]|$ with $E_x [|T(x)|]$, the first tree has orders of magnitude larger contribution than the rest of the trees in the ensemble. This behavior is inherent in the algorithm: if one would add a constant value to all the labels in the training data, only the first tree will get modified (with this constant value added to all its leaves) and the rest of the trees will remain with a small contribution to the model. Therefore, in a sense, the first tree learns the bias of the problem while the rest of the trees in the ensemble learn the deviation from this bias. This makes the ensemble very sensitive to the decisions made by the first tree. This can be seen in Figure 2 as well, which depicts a few trees in the ensemble trained by different methods for the above mentioned task. We can see that the MART algorithm (without using shrinkage) adds trees that make negligible contribution to the overall prediction for most of the data points as indicated by the large yellow leaves in the first column.

As discussed briefly in Section 1, shrinkage (Friedman, 2001, 2002) is the most common approach employed to combat the issue of over-specialization. Since shrink-

²Muting trees and muting features can be done at the same time and indeed we do this in our experiments.




Figure 1: The average contribution of the trees in the ensemble for different learning algorithms (the graph presents the absolute value of the average). The shrinkage factor used is 0.1.

age reduces the impact of each tree by a constant value, the first tree cannot compensate for the entire bias of the problem. We can see the impact of this strategy in Figure 1 as well as in Figure 2. We observe that the contribution of later trees do drop, but at a much slower rate than in the case where shrinkage is not used. For example, while the contribution of the 100th tree in MART without shrinkage is about 15 orders of magnitude smaller than the contribution of the first tree, this factor in MART with shrinkage drops to “only” 4 orders of magnitude. In figure 2 we see that the large yellow leaves, representing the fact that a tree “abstains” on many of the instances, appear later in the ensemble. As we can see, the differences in the contributions from the trees in the ensemble are more gradual when shrinkage is used, nevertheless they are still notable.

Now, let us see the effect of using dropouts as employed in DART. The last column in Figure 2 depicts trees learned by the DART algorithm. First, compared to MART and MART with shrinkage, we see that trees specialized at a significantly slower rate as indicated by the much slower emergence of large yellow leaves. This can be seen in Figure 1 as well, where we see that the expected contribution of the trees added in later iterations do not drop much.³ Therefore, the sensitivity to the contribution of the individual trees is drastically reduced. At the same time, unlike random forest, DART continues to learn trees to compensate for the deficiencies of the existing trees in the ensemble. It, however, does so in a controlled manner to strike a balance between diversity and over-specialization. We will see in Section 3 that both MART and random

forest can be viewed as extreme cases of the DART algorithm.

3 Description of the DART Algorithm

We start our presentation with the MART algorithm as the foundation on which DART builds. MART can be viewed as a gradient descent algorithm (Friedman, 2001): at every iteration, MART computes the derivative of the loss function for the current predictions and adds a regression tree that fits the inverse of these derivatives to the ensemble. More formally, the input to the algorithm includes a set of points and their labels, (x, y) , where the points x are in some space \mathcal{X} and the labels y are in a label space. The algorithm also takes as input a loss generating function which is tuned to the task at hand (for example, regression, classification, ranking, etc.). Using the loss generating function and the labels, the algorithm defines the loss for every point x , $L_x : \mathcal{Y} \mapsto \mathbb{R}$ where \mathcal{Y} is the prediction space, typically the reals. For example, if the task is regression then the loss may be defined as $L_x(\hat{y}) = (\hat{y} - y)^2$ where y is the true label of x .

At every iteration, let the current model be denoted by $M : \mathcal{X} \mapsto \mathcal{Y}$ and $M(x)$ denotes the prediction of the current model for point x . Let $L'_x(M(x))$ be the derivative of the loss function at $M(x)$. MART creates an intermediate dataset in which a new label, $-L'_x(M(x))$, is associated with every point x in the training data. A tree is trained to predict this inverse derivative and added to the ensemble as a step in the inverse direction of the derivative (in order to minimize the loss).

The choice of the loss makes the MART algorithm applicable to a variety of learning tasks. As discussed earlier, the squared loss is used for regression

³Linear regression on this data suggests that there might be a slow decline in the average contribution of the trees at a rate of 0.0003.

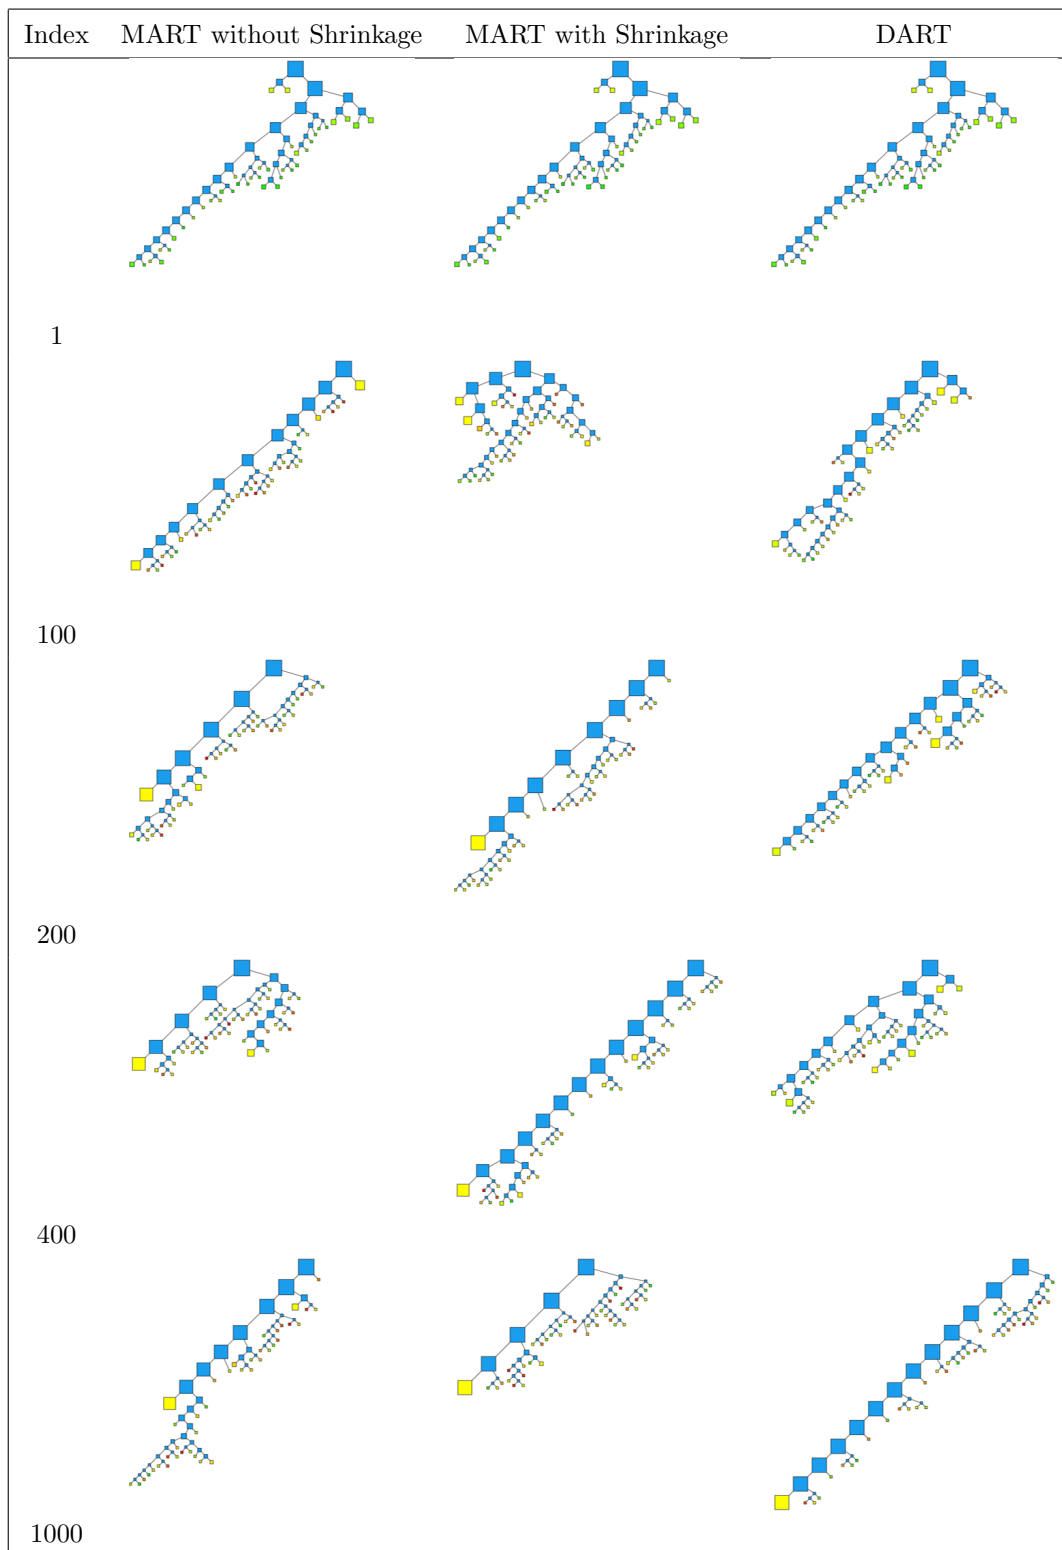


Figure 2: Examples of trees in the ensemble for the regression task on CT slice dataset (Section 4.2). Each column represents a different learning algorithm (MART (without shrinkage), MART+shrinkage, and DART). Each row represents a different index in the ensemble: 1st, 100th, 200th, 400th and 1000th tree in the ensemble. In each tree, the size of nodes is proportional to the percentage of the instances that reach this node. The color gradient of leaves represent the range of values where green stands for the positive extreme, yellow for zero, and red for negative extreme.

tasks. The logistic loss function is used for classification tasks. Here, the loss function is defined to be $L_x(\hat{y}) = (1 + \exp(\lambda y \hat{y}))^{-1}$ where λ is a parameter. For ranking tasks, the loss function would depend on the relative ordering of the points in the predicted ranking. In our evaluation (Section 4) for ranking tasks, we use the definition of the LambdaMart method (Burges, 2010). The main idea here is to directly define the gradient of the loss function:

$$L'_x(M(x)) := \sum_{x'} \frac{s(x, x') \lambda}{1 + \exp(\lambda(M(x) - M(x')))}$$

where λ is a parameter and $s(x, x')$ is the NDCG loss that results from reversing the order of the points x and x' , and the summation is over all the points which relate to the same query. See Burges et al. (2007) for more details.

As discussed in Section 1 and Section 2, the gradient-descent style boosting that MART employs may lead to over-specialization, and a common approach employed to address this issue is to use shrinkage. Under this method, MART operates as described above when learning the new tree in every iteration. However, before adding this newly learned tree to the ensemble, its leaf values are reduced in magnitude by multiplying them with a constant value in $(0, 1)$. Shrinkage helps in alleviating the problem of over-specialization to a certain extent as we observed in Section 2.

We now move on to describing the DART algorithm, which is presented as Algorithm 1. DART diverges from MART at two places. First, when computing the gradient that the next tree will fit, only a random subset of the existing ensemble is considered. Let us say that the current model M after n iterations is such that $M = \sum_{i=1}^n T_i$, where T_i is the tree learned in the i 'th iteration. DART first selects a random subset $I \subset \{1, \dots, n\}$ and creates a model $\hat{M} = \sum_{i \in I} T_i$. Given this model, it learns a regression tree T to predict the inverse derivative of the loss function with respect to this modified model by creating the intermediate dataset $\left\{ \left(x, -L'_x(\hat{M}(x)) \right) \right\}$.

The second place at which DART diverges from MART is when adding the new tree to the ensemble where DART performs a normalization step. The rationale behind the normalization step is that the new trained tree T is trying to close the gap between \hat{M} and the optimal predictor, however, the dropped trees are also trying to close the same gap. Therefore, introducing both the new tree and the dropped trees will result in the model overshooting the target. Furthermore, assuming that the number of trees dropped from the ensemble to create I that result in the model \hat{M} is k , the new tree T has roughly k times larger magnitude than

each of the individual trees in the set of dropped trees. Therefore, DART scales the new tree T by a factor of $1/k$ such that it will have the same order of magnitude as the dropped trees. Following this, the new tree and the dropped trees are scaled by a factor of $k/(k+1)$ and the new tree is added to the ensemble. Scaling by the factor of $k/(k+1)$ ensures that the combined effect of the dropped trees together with the new tree remains the same as the effect of the dropped trees alone before the introduction of the new tree.

As seen in Figure 1 and Figure 2, DART reduces the problem of over-specialization. Therefore, it can be viewed as regularization where the number of trees dropped controls the amount of regularization. On one extreme, if no tree is dropped, DART is no different than MART. On the other extreme, if all the trees are dropped, the DART is no different than random forest. Therefore, the size of the dropped set allows DART to vary between the ‘‘aggressive’’ MART mode to a ‘‘conservative’’ random-forest mode.

There are many ways to select the trees to be dropped. In the experiments reported here, we have employed what we call the *Binomial-plus-one* technique. In this technique, each of the existing trees in the ensemble is dropped with a probability p_{drop} . However, if no tree was selected to be dropped using the above binomial selection, a single tree is selected uniformly at random to be dropped. Therefore, at least one tree will be dropped at each iteration.

If p_{drop} is set to a very small value, the random selection boils down to simply dropping a single tree in each round. We have experimented with this mode as well, and we denote this mode by defining p_{drop} to be ε in the evaluation results presented in Section 4.

4 Evaluation

We evaluated DART for three different tasks: ranking, regression and classification. For each of the tasks, we used large scale, publicly available datasets. In our evaluation, we compare DART to MART with different shrinkage factors. Furthermore, since random forests (RF) can be considered as an extreme case of DART, we compare to this algorithm as well whenever applicable.

4.1 Ranking

MART is commonly used for ranking tasks. For example, in the Yahoo! learning to rank challenge, the winners employed boosted trees (Chapelle and Chang, 2011) based on the LambdaMart method (Burges, 2010). We introduced dropouts as explained in Section 3 into LambdaMart and tested it on the MSLR-

Parameter	MART	DART
Shrinkage	0.05, 0.1, 0.2, 0.4	-
Dropout rate	-	ε , 0.015, 0.03, 0.045
Number of trees	100	
Leaves per tree	40	
Loss function parameter	0.2,0.4,0.6,0.8,1,1.2	
Fraction of features scanned per leaf	0.5, 0.75, 1.0	

Table 1: Parameter values scanned for the ranking task.

algorithm	Shrinkage	Dropout	Loss function parameter	Feature fraction	NDCG@3
MART	0.4	0	1.2	0.75	46.31
DART	1	0.03	1.2	0.5	46.70

Table 2: NDCG scores for MART and DART on the ranking task. For NDCG scores, higher is better.

Algorithm 1 The DART algorithm

Let N be the total number of trees to be added to the ensemble
 $S_1 \leftarrow \{x, -L'_x(0)\}$
 T_1 be a tree trained on the dataset S_1
 $M \leftarrow \{T_1\}$
for $t = 2, \dots, N$ **do**
 $D \leftarrow$ the subset of M such that $T \in M$ is in D
with probability p_{drop}
 if $D = \emptyset$ **then** $D \leftarrow$ a random element from M
 end if
 $\hat{M} \leftarrow M \setminus D$
 $S_t \leftarrow \{x, -L'_x(\hat{M}(x))\}$
 T_t be a tree trained on the dataset S_t
 $M \leftarrow M \cup \left\{ \frac{T_t}{|D|+1} \right\}$
 for $T \in D$ **do**
 Multiply T in M by a factor of $\frac{|D|}{|D|+1}$
 end for
end for
Output M

WEB10K dataset.⁴ This dataset contains $\sim 1.2M$ query-URL pairs for 10K different queries and the task is to rank the URLs for each query according to their relevance using the 136 available features.

The dataset is partitioned into five parts such that 60% of the data is used for training, 20% is used for validation, and 20% for testing. We scanned the values of various parameters for both algorithms by training on the training data and comparing their performance on the validation data. We selected the best performing models based on their scores on the validation set, and applied them to the test set to obtain the reported results. The different parameters scanned are summarized in Table 1. For each of the parameter combinations experimented, we computed the NDCG score at position 3 and used this as the metric for selecting the parameter values. NDCG (Burges et al., 2005) is a common metric used to evaluate web-ranking tasks. Moreover, the loss functions used were designed to optimize this metric (Burges, 2010).

Table 2 presents the main results for the ranking task. DART gains ~ 0.4 NDCG points over MART. Moreover, when checking the NDCG scores at positions 1 and 2 we see significant gains as well (0.2 points gain in position 1 and 0.38 points gain in position 2). To put this observed improvement in perspective, in the Yahoo! learning to rank challenge, the gap, in terms of NDCG, between the winners and the team who ranked 5th was 0.35 points (Chapelle and Chang, 2011).

⁴<http://research.microsoft.com/en-us/projects/mslr/default.aspx>

Parameter	MART	DART	Random Forest
Shrinkage	0.05 , 0.1, 0.2, 0.3, 0.5	-	-
Dropout rate	-	ϵ , 0.01 , 0.025, 0.05, 0.1, 0.2	-
Fraction of instances used per tree	1.0	1.0	0.25, 0.5, 0.75, 1.0
Number of trees	25, 50,100,250,500, 1000		
Leaves per tree	50 ,100,250,500,1000	50 ,100,250,500,1000	50,100,250,500, 1000
Fraction of features scanned per leaf	0.05,0.1, 0.2 ,0.4, 0.8,1.0	0.05,0.1,0.2, 0.4 , 0.8,1.0	0.01,0.025,0.05,0.1,0.2, 0.4 , 0.5,0.8,1.0

Table 3: Parameter values scanned for the regression task. The parameter values that yielded the lowest loss under each algorithm are highlighted.

Ensemble size	25	50	100	250	500	1000
MART	35.13	31.79	30.92	30.07	29.76	29.28
DART	32.50	30.50	29.66	28.14	28.11	27.98
Random Forest	32.76	33.21	32.88	32.36	32.66	32.33

Table 4: L2 error of optimal parameter combinations for DART, MART and random forest on the regression task for various ensemble sizes. DART outperforms MART and random forest for all the ensemble sizes tested (the best result for every ensemble size is boldfaced).

4.2 Regression

To test the merits of using dropouts for regression tasks we have used the CT slices dataset (Graf et al., 2011) available at the UCI repository (Bache and Lichman, 2013). This dataset contains 53500 histograms created from CT scans of 74 individuals. The task is to infer the location on the axial axis where the image was taken from. Each image is represented as a 386 dimensional feature vector. We scanned values for various parameters involved and these are summarized in Table 3. We have used 10 fold cross validation to compare the algorithms. The folds were selected such that either all the images of an individual are in the train set or all of them are in the test set.

The evaluation results for the regression task are presented in Table 4. For every ensemble size, the best DART model, outperformed both the best MART and the best RF models. We observed that DART outperforms MART and RF even when DART is restricted to drop only a single tree in every iteration (that is the dropout rate is ϵ).

Furthermore, we observed that RF requires large trees to achieve low losses. For example, when the tree sizes are limited to 50 and 100 leaves, the best RF model achieved a loss of 44.48 and 36.29 respectively. On the

other hand, MART and DART achieve their lowest loss values with trees comprising only 50 leaves.

4.3 Classification

The performance of DART on classification tasks was evaluated using the face detection (fd) dataset from the Pascal Large Scale Learning Challenge.⁵ This dataset contains 30x30 gray scale images and the goal is to infer whether there is a face in the image or not. We used the first 300K examples for training, the next 200K examples for validation and the next 200K examples for testing. The parameters scanned for this task are summarized in Table 5.

We used the validation set to select the best performing parameters for the MART, DART and random forest models and evaluated them on the test set. Table 6 presents the results for the classification task. Both MART and DART achieve the highest accuracy with ensembles of 250 trees. Although the difference in accuracies is small, it is statistically significant ($P < 0.0001$), since the two models disagree on 1106 predictions on the test set and the MART model gets only 481 of them right while the DART model

⁵<http://largescale.ml.tu-berlin.de/instructions/>

Parameter	MART	DART	Random Forest
Shrinkage	0.2, 0.3 , 0.4, 0.5	-	-
Dropout rate	-	ϵ , 0.015, 0.03, 0.045	-
Fraction of instances per tree	1.0	1.0	0.25 , 0.5, 0.75, 1.0
Number of trees	50, 100, 250 , 500, 1000	50, 100, 250 , 500, 1000	50, 100, 250, 500 , 1000
Leaves per tree	40	40	50, 100, 250, 500, 1000
Loss function parameter	0.2, 0.3, 0.4, 0.5	0.2, 0.3, 0.4, 0.5	-
Fraction of features per leaf	0.5, 0.75, 1.0	0.5, 0.75, 1.0	0.5, 0.75, 1.0

Table 5: Parameter values scanned for the classification task. The parameter values that yielded the highest accuracy under each algorithm are highlighted.

Ensemble size	50	100	250	500	1000
MART	0.9687	0.9699	0.9707	0.9704	0.9695
DART	0.9676	0.9692	0.9714*	0.9693	0.9699
Random Forest	0.9627	0.9629	0.9629	0.9630	0.9628

Table 6: Accuracies on the test set for DART, MART and random forest on the face-detection classification task for various ensemble sizes. The results are comparable between DART and MART: while MART “wins” on 3 out of the 5 different ensembles sizes, however, the best model is a DART model.

gets 625 of them correct. The main difference between the models is in their recall where MART has a recall rate of 0.665 while DART has a recall rate of 0.672. This is a significant difference for this dataset due to its highly skewed nature: only $\sim 8.6\%$ of the instances are labeled positive. Random forest exhibits lower accuracy for this task.

In our experiments, random forest did not compare well against MART or DART. Since MART and random forest are the two extremes of the DART algorithm, it serves us to show that the optimal point between these two extremes is not trivial.

5 Conclusions

Dropouts (Hinton et al., 2012) have been shown to improve the accuracies of Neural Network models significantly. On the other hand, Multiple Additive Regression Trees (MART) (Friedman, 2001; Elith et al., 2008) have been found to be the most accurate models for many tasks (Caruana and Niculescu-Mizil, 2006), most notably the web ranking task (Chapelle and Chang, 2011). Motivated by the observation that MART adds trees with significantly diminishing contributions, we hypothesize that dropouts can provide efficient regularization for MART and propose the DART algorithm. Our experiments show that this is indeed the

case: trees in the ensemble created by DART contribute more evenly towards the final prediction, as shown in Figure 1. In addition, this results in considerable gains in accuracies for ranking, regression and classification tasks.

This study opens the door to several future directions. For example, using the same technique proposed in this work, it is possible to introduce dropouts in other models such as AdaBoost (Freund and Schapire, 1995). The simplicity of these models may allow us to improve our understanding of dropouts. Another direction is to further tune the DART algorithm by experimenting different ways of selecting the dropped set and the normalization techniques. Furthermore, the even contribution of the trees in DART may allow using it for learning tasks with drifting targets. This can be achieved, for example, by periodically dropping a subset of the existing trees and learning new trees, with new data, to replace them.

Acknowledgments

This research was conducted while the first author was an intern at the machine learning department at Microsoft Research.

References

- Kevin Bache and Moshe Lichman. UCI machine learning repository, 2013. URL <http://archive.ics.uci.edu/ml>.
- Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- Chris Burges, Tal Shaked, Erin Renshaw, Ari Lazier, Matt Deeds, Nicole Hamilton, and Greg Hullender. Learning to rank using gradient descent. In *Proceedings of the 22nd international conference on Machine learning*, pages 89–96. ACM, 2005.
- Christopher JC Burges. From ranknet to lambdarank to lambdamart: An overview. *Learning*, 11:23–581, 2010.
- Christopher J.C. Burges, Robert Ragno, and Quoc Viet Le. Learning to rank with nonsmooth cost functions. *NIPS07*, 19:193, 2007.
- Rich Caruana and Alexandru Niculescu-Mizil. An empirical comparison of supervised learning algorithms. In *Proceedings of the 23rd international conference on Machine learning*, pages 161–168. ACM, 2006.
- Olivier Chapelle and Yi Chang. Yahoo! learning to rank challenge overview. In *Yahoo! Learning to Rank Challenge*, pages 1–24, 2011.
- Jane Elith, John R Leathwick, and Trevor Hastie. A working guide to boosted regression trees. *Journal of Animal Ecology*, 77(4):802–813, 2008. ISSN 1365-2656. doi: 10.1111/j.1365-2656.2008.01390.x. URL <http://dx.doi.org/10.1111/j.1365-2656.2008.01390.x>.
- Yoav Freund. An adaptive version of the boost by majority algorithm. *Machine learning*, 43(3):293–318, 2001.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *Computational learning theory*, pages 23–37. Springer, 1995.
- Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of Statistics*, pages 1189–1232, 2001.
- Jerome H Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- Franz Graf, Hans-Peter Kriegel, Matthias Schubert, Sebastian Pölsterl, and Alexander Cavallaro. 2d image registration in ct images using radial image descriptors. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2011*, pages 607–614. Springer, 2011.
- Geoffrey E Hinton, Nitish Srivastava, Alex Krizhevsky, Ilya Sutskever, and Ruslan R Salakhutdinov. Improving neural networks by preventing co-adaptation of feature detectors. *arXiv preprint arXiv:1207.0580*, 2012.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
- Laurens Maaten, Minmin Chen, Stephen Tyree, and Kilian Q. Weinberger. Learning with marginalized corrupted features. In Sanjoy Dasgupta and David Mcallester, editors, *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, volume 28, pages 410–418. JMLR Workshop and Conference Proceedings, 2013. URL <http://jmlr.csail.mit.edu/proceedings/papers/v28/vandermaaten13.pdf>.
- Stefan Wager, Sida Wang, and Percy Liang. Dropout training as adaptive regularization. In *Advances in Neural Information Processing Systems*, pages 351–359, 2013.
- Sida Wang and Christopher Manning. Fast dropout training. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 118–126, 2013.