

Residual Likelihood Forests

Yan Zuo
yan.zuo@monash.edu

ARC Centre of Excellence for Robotic
Vision, Monash University, Australia

Tom Drummond
tom.drummond@monash.edu

Abstract

This paper presents a novel ensemble learning approach called Residual Likelihood Forests (RLF). Our weak learners produce conditional likelihoods that are sequentially optimized using global loss in the context of previous learners within a boosting-like framework (rather than probability distributions that are measured from observed data) and are combined multiplicatively (rather than additively). This increases the efficiency of our strong classifier, allowing for the design of classifiers which are more compact in terms of model capacity. We apply our method to several machine learning classification tasks, showing significant improvements in performance. When compared against several ensemble approaches including Random Forests and Gradient Boosted Trees, RLFs offer a significant improvement in performance whilst concurrently reducing the required model size.

1 Introduction

Ensemble and Boosting methods such as Random Forests [3] and AdaBoost [19] are often recognized as some of the best out-of-the-box classifiers, consistently achieving state-of-the-art performance across a wide range of computer vision tasks including applications in image classification [1], semantic segmentation [2], object recognition [2] and data clustering [16]. The success of these methods is attributed to their ability to learn models (strong learners) which possess low bias and variance through the combination of weakly correlated learners (weak learners). Forests reduce variance through averaging its weak learners over the ensemble. Boosting, on the other hand, looks towards reducing both bias and variance through sequentially optimizing under conditional constraints.

The commonality between both approaches is in the way each learner is constructed: both methods use a top-down induction algorithm (such as CART [2]) which greedily learns decision nodes in a recursive manner. This approach is known to be suboptimal in terms of objective maximization as there are no guarantees that a global loss is being minimized [4]. In practice, this type of optimization requires the non-linearity offered by several (very) deep trees, which results in redundancy in learned models with large overlaps of information between weak learners.

To address these limitations, the ensemble approaches of [11, 20] have utilized gradient information within a boosting framework. This allows weak learners to be fit via pseudo-residuals or to a set of adaptive weights and allows for the minimization of a global loss via gradient descent. Whilst these methods have demonstrated improved performance over their

Random Forest counterpart, they are more susceptible to overfitting, making them harder to tune. An alternative approach proposed by [17] tried to retain Random Forests as a baseline classifier, employing a global refinement technique combined with leaf pruning to gain improvements in performance and reduce model size. However, this method still required the relatively large overhead of constructing the original Random Forest before refinement can be performed (with an additional overhead associated with global refinement and pruning).

In this work, we propose a method which constructs a decision forest that benefits from the boosting approach of creating complementary weak learners through minimizing a global loss, yet retains the simplicity in model tuning of Random Forests. Our approach monotonically decreases the loss function of the ensemble using a novel, top-down induction algorithm, which unlike CART, accounts for the mutual information between weak learners already in the ensemble. Empirically, we extensively evaluate our approach across multiple datasets against several ensemble and boosting baselines. Our results show that our method improves classification performance at a higher parameter efficiency.

The main contributions of this paper are as follows:

- We propose a novel ensemble learning framework that includes a novel induction algorithm which accounts for the mutual information between weak learners it constructs.
- Theoretically, we show that our method allows optimization to occur in closed-form (in contrast to existing gradient boosting methods), which simplifies model tuning.
- Empirically, we show that our method offers dramatic reductions in required model size. Further, we show that our approach outperforms several competing boosting and ensemble baselines across 5 standard machine learning datasets.

2 Related Work

2.1 Random Forests

RFs became popular in the literature due to their flexibility and robustness when dealing with data-driven learning problems [8]. RFs were first developed by [15] to overcome the variance and stability issues found in binary decision trees by randomly choosing the subspace of features for tree training. This method was refined by [9], where the bootstrap aggregating technique of [9] was combined with the random feature selection of RFs. [13] developed Extremely Randomized Trees, where both the attribute and cut-point of the input feature were randomized. [2] applied RFs to the semantic segmentation problem, using them as a global image classifier to quickly and efficiently provide priors to be pipelined into the system developed in [21].

2.2 Boosting

The work of [18] first introduced the concept of boosting using an algorithm that selectively combined several weak hypotheses into a single strong hypothesis. [8] extended this algorithm to combine the outputs of all weak hypotheses. These methods however, did not take advantage of the adaptability of weak learners. This was addressed by [9] with Adaptive Boosting, where each subsequent weak learner introduced to the strong learner is biased to favor misclassified training samples using high-weight data points. [9] extended on this

concept of adaptable boosting by offering an alternative to using weights, formulating the method as a gradient descent problem with a specialized loss function.

3 Preliminaries

Conventional decision trees consist of internal decision nodes and terminal leaf nodes. The internal nodes of each decision tree are a set of \mathcal{N}_d decision nodes, $D = \{d_0, \dots, d_{\mathcal{N}_d-1}\}$, each of which holds a decision function $d_i(\mathbf{x}; \theta_i)$, where θ_i are the parameters for decision node i in the tree. For binary decision trees, the decision function for each decision node i is defined as $d_i(\mathbf{x}; \theta_i) : \mathcal{X} \rightarrow [0, 1]$, which routes a sample to a left or right subtree. Collectively, the decision nodes route data, \mathbf{x} , through the tree until a terminal leaf node is reached: $\ell = \delta(\mathbf{x}; \Theta)$. Here, δ signifies the routing function which directs data x to the terminal leaf node ℓ , given the tree's decision node parameters Θ .

For classification, the leaf nodes Q in a decision tree contain class label probability distributions, $\mathbf{q} = Q(\ell)$. These distributions are formed from the training data and accompanying ground truth class labels. For a terminal leaf node ℓ , the stored probability value for class j is given as:

$$q_j = Q(\ell)_j = \frac{\sum_n c_{n,j}}{N_\ell} \quad (1)$$

where N_ℓ is the total number of samples routed into leaf node ℓ and $c_{n,j}$ is the observed class label for sample n in class j :

$$c_{n,j} = \begin{cases} 1, & \text{if sample } n \text{ has class label } j \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

Typically, the decision nodes of the tree are chosen such that resulting distributions generated give the best split according to some purity measure (*i.e.* entropy or the Gini index [9]).

A conventional decision forest is an ensemble of \mathcal{N}_t number of decision trees. It delivers a final prediction for the input \mathbf{x} by routing \mathbf{x} through each decision tree in the ensemble to a leaf node and taking the average prediction delivered by each tree:

$$P_j = \frac{1}{\mathcal{N}_t} \sum_t Q'(\delta^t(\mathbf{x}; \Theta^t))_j \quad (3)$$

where Q' , δ^t and Θ^t are the respective distributions, routing functions and parameters of tree t in the ensemble.

4 Residual Likelihood Forests

First, we motivate our proposed approach using the following simple classification example involving two weak learners evaluating some input data. To keep our notation concise, we denote 1p_j and 2p_j to represent the probabilities assigned by each weak learner for the underlying class label j given the input. Given that the prior is uniform over classes, the correct approach to combining these two weak learners depends on the correlation between the assigned probabilities from each weak learner.

If 1p_j and 2p_j are independent of one another, the correct class probability should be given by the normalized product for the class j :

$$P_j = \frac{{}^1p_j {}^2p_j}{\sum_k^{\mathcal{N}_c} {}^1p_k {}^2p_k} \quad (4)$$

where \mathcal{N}_c is the number of classes. However, when the two weak learners are fully correlated, the assigned probabilities from each weak learner would be the same and ${}^1p_j = {}^2p_j$. In this instance, applying Eq. 4 would result in an incorrect estimate of the class probabilities as the normalized squared distribution and we should instead simply average their marginal distributions.

4.1 Weak Learners Generating Likelihoods

In practice, probability distributions ${}^1\mathbf{p}$ and ${}^2\mathbf{p}$ will rarely be fully independent or fully correlated. As such, deciding the amount each weak learner should contribute to the strong classifier's prediction is a non-trivial task. Typically, ensemble methods combine their weak learners by averaging their marginal distributions [1, 8, 18], relying on the law of large numbers to get a good estimate of the underlying distribution of the data and smooth out variance between learners.

The correlation between learners can be problematic if we wish to have a consistent approach for correctly combining weak learners within an ensemble since the approach that should be adopted depends on this correlation. Since it is difficult to determine whether weak learners are learning independent information or correlated information, we approach the problem from a different perspective: instead of deciding how to combine distributions from each weak learner, we can instead modify the stored p distributions of each weak learner in a manner which allows for a consistent approach towards combining each weak learner's contribution irrespective of their correlation.

Looking back at the simple classification example, if ${}^1p_j = {}^2p_j$, modifying the stored distributions p into information vectors Π by replacing 1p_j and 2p_j with their square roots and applying Eq. 4 gives the correct answer for ${}^1p_j = {}^2p_j$:

$$P_j = \frac{{}^1\Pi_j {}^2\Pi_j}{\sum_k {}^1\Pi_k {}^2\Pi_k} = \frac{\sqrt{{}^1p_j} \sqrt{{}^2p_j}}{\sum_k \sqrt{{}^1p_k} \sqrt{{}^2p_k}} = {}^1p_j = {}^2p_j \quad (5)$$

These information vectors ${}^1\Pi$ and ${}^2\Pi$ no longer represent probability distributions. Instead, each information vector can be learned and combined with other information vectors contributed by other weak learners from the ensemble. This forms the basis for our ensemble learning framework: for each decision tree, we construct the information vectors $\{\Pi_1, \dots, \Pi_{\mathcal{N}_l}\}$ for its \mathcal{N}_l leaf nodes. Each vector is conditioned upon the prior information contained by trees already existing in the ensemble and is constructed to contain any new residual information gained from adding the new tree. and henceforth, we refer to each vector Π as a *residual likelihood*. As we will demonstrate shortly, we can generate these residual likelihoods such that they can be treated as independent sources of information, forming a strong classifier distribution by taking their product and normalizing:

$$P_j = \frac{\prod_t^{\mathcal{N}_t} \Pi_{t,j}}{\sum_k \prod_t^{\mathcal{N}_t} \Pi_{t,k}} \quad (6)$$

4.2 Optimization Method

To construct our forest, we minimize a global loss function in a similar manner to gradient boosting approaches [10]. The key difference in our approach is that the importance allocated towards hard-to-classify samples during the training process is implicit rather than explicit (as in the case of gradient boosted trees) and is directly incorporated into our induction algorithm. Specifically, our method still chooses decisions from a random subset of the feature space but these decisions are selected with the objective of minimizing the global loss function. Incidentally, this removes the need for weighted training samples or specialized loss functions found in traditional boosting techniques and simplifies the model tuning process.

Deriving the Objective Loss For a new tree to be added to the ensemble, we can derive a solution which optimizes the residual likelihood its leaf node contributes. Consider the i^{th} tree to be added to the ensemble. We wish to find the residual likelihood for class j ($\Pi_{i,j}$) it contributes such that the objective loss is minimized when this tree is added to the ensemble. For a given leaf node, we define the following terms:

- $P_{n,j}^-$ is the prior value of class j for sample n , combining likelihoods $\{\Pi_{1,j}, \dots, \Pi_{i-1,j}\}$ from all existing weak learners in the ensemble (excluding the new tree):

$$P_{n,j}^- = \prod_t^{i-1} \Pi_{t,j} \quad (7)$$

- $\Pi_{i,j}$ is the stored value for class j of the residual likelihood to be contributed by the new tree i .
- $P_{n,j}^+$ is the posterior probability of class j for sample n after new tree i has contributed to the ensemble:

$$P_{n,j}^+ = \frac{P_{n,j}^- \Pi_{i,j}}{\sum_k P_{n,k}^- \Pi_{i,k}} \quad (8)$$

- $c_{n,j}$ is the ground truth class label for sample n taking the class label j , as defined in Eq. 2.

The cross entropy loss of the strong classifier after adding the new tree i is defined as:

$$\mathbb{L} = - \sum_n \sum_j c_{n,j} \log(P_{n,j}^+) \quad (9)$$

After applying Eq. 8, we obtain the following:

$$\begin{aligned} \mathbb{L} &= - \sum_n \sum_j c_{n,j} \log \left(\frac{P_{n,j}^- \Pi_{i,j}}{\sum_k P_{n,k}^- \Pi_{i,k}} \right) \\ &= - \sum_n \sum_j c_{n,j} \left(\log(P_{n,j}^- \Pi_{i,j}) - \log \left(\sum_k P_{n,k}^- \Pi_{i,k} \right) \right) \end{aligned} \quad (10)$$

For the i^{th} tree added to the strong classifier, we require a contributed residual likelihood $\Pi_{i,j}$ to the class j such that \mathbb{L} is minimized. For the class j , the first-order derivatives of Eq. 10 is given as:

$$\frac{\partial \mathbb{L}}{\partial \Pi_{i,j}} = - \sum_n c_{n,j} \frac{P_{n,j}^-}{P_{n,j}^- \Pi_{i,j}} + \sum_n \frac{P_{n,j}^-}{\sum_k P_{n,k}^- \Pi_{i,k}} \quad (11)$$

Where we set Eq. 11 to zero to yield a solution:

$$\begin{aligned} \frac{N_j}{\Pi_{i,j}} &= \sum_n \frac{P_{n,j}^-}{\sum_k P_{n,k}^- \Pi_{i,k}} \\ \therefore N_j &= \sum_n \frac{P_{n,j}^- \Pi_{i,j}}{\sum_k P_{n,k}^- \Pi_{i,k}} \\ \therefore N_j &= \sum_n P_{n,j}^+ \end{aligned} \quad (12)$$

Here, N_j is the total number of samples with class label j within the given leaf node. This is a convex optimization problem with a single global solution where the loss function is convex:

$$\begin{aligned} \frac{\partial^2 \mathbb{L}}{\partial \Pi_{i,j}^2} &= \sum_n c_{n,j} \frac{1}{\Pi_{i,j}^2} - \sum_n \frac{(P_{n,j}^-)^2}{(\sum_k P_{n,k}^- \Pi_{i,k})^2} \geq 0 \\ \therefore \sum_n c_{n,j} &\geq \sum_n \frac{(P_{n,j}^- \Pi_{i,j})^2}{(\sum_k P_{n,k}^- \Pi_{i,k})^2} = \sum_n (P_{n,j}^+)^2 \end{aligned} \quad (13)$$

Eq. 13 holds true for each sample n where $0 < \sum_k (P_{n,k}^+)^2 \leq 1$ and $\sum_k c_{n,k} = 1$, which satisfies the criteria for convexity.

Computing Residual Likelihoods Eq. 12 indicates that for a leaf node and samples routed to that leaf node, the solution for the residual likelihood Π exists when the sum of strong classifier probabilities across samples is equal to the sum of class labels associated with those samples for a given class. For computing a Π that fulfills this criteria, we adopt a straightforward strategy: using Eq. 12, we first initialize the residual likelihoods of the new tree $\Pi = \mathbf{1}$ and solve for the updated residual:

$$\text{Initialize: } \Pi_j \leftarrow 1 \quad (14)$$

$$\text{Iterate: } \Pi_j \leftarrow \Pi_j \frac{N_j}{\sum_n P_{n,j}^+} \quad (15)$$

which converges quickly and stabilizes when Eq. 12 holds, solving for Π_j for all j . In theory, the estimate for Π_j at the solution can be improved arbitrarily by continuously iterating on Eq. 15. However, in practice, we found that only a single iteration was required for Eq. 12 to hold. In Fig 1, we outline our approach.

5 Experiments

Experiment Settings For all our experiments, we follow the settings described in [20]. We set the default number of trees in an ensemble to 100, keeping in line with the experiments

Algorithm 1 RLF Optimization

Require: \mathcal{N} : training set
Require: Number of trees in forest \mathcal{N}_t
Require: Feature pool size S
Require: Maximum tree depth D_{max}

- 1: **for all** $t \in \{1, \dots, \mathcal{N}_t\}$ **do**
- 2: **for all** $m \in \{1, \dots, D_{max}\}$ **do**
- 3: In parallel:
- 4: **for all** $s \in \{1, \dots, 2^{D_{max}-1}\}$ **do**
- 5: In parallel:
- 6: **for all** $p \in \{1, \dots, S\}$ **do**
- 7: Choose a random θ_p
- 8: Calculate N_j and $\sum_i P_{n,j}^-$
- 9: Calculate leaf residuals Π_j
- 10: Calculate \mathbb{L}
- 11: **end for**
- 12: Choose θ_p for node d that minimizes \mathbb{L}
- 13: **end for**
- 14: **end for**
- 15: Copy the winning Π_j into the leaves of the tree t
- 16: **end for**

in [14, 20]. We use the XGBoost [2] implementation of Gradient Boosted Trees (GBT), using the prescribed default settings. For completeness, we also trained an additional GBT model (XGBoost-Large) in which we match the number of trees and tree depth of our RLF model.

For each dataset, we test a number of random features equal to the square root of the feature dimensions, as recommended in [9], allocating 10 random thresholds per feature. In each case, we report mean error and standard deviation across 10 separate runs to account for variance due to randomness during training, except for the *G50c* dataset which we report the mean error and standard deviation across 250 separate runs as was done in [20].

Due to our training scheme differing from the conventional entropy minimization approaches of competing approaches, we do not specify an early node termination policy. Instead, we rely on the efficiency of our decision node split choices and construct considerably shallower models. Hence, for practical purposes, we construct trees up to a maximum specified depth of 15. For the overall results shown in Table 1, we follow the tree depth settings in [20], except for the *Chars74k* dataset in which we set a tree depth of 15 (as opposed to 25).

5.1 Chars74k & MNIST Ablations

Chars74k We vary hyperparameters of RLFs and competing baselines for the *Chars74k* dataset, observing the classification error on the test data. The results of these experiments are shown in Fig. 1. Figs. 1(a), 1(b) and 1(c) show the test classification error on the *Chars74k* dataset when maximum tree depth is restricted to 5, 10 and 15 respectively. We can see that our RLF approach clearly outperforms all competing baselines, particularly with

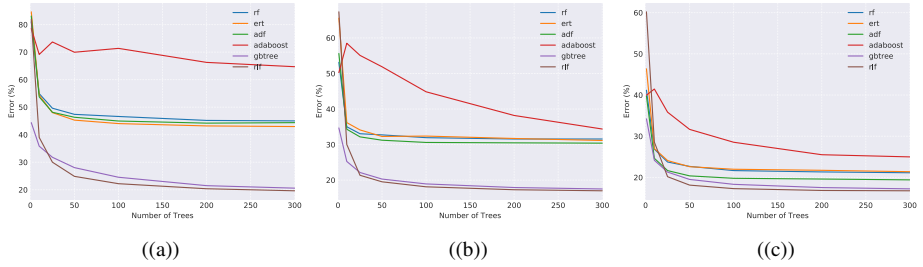


Figure 1: Classification error of RLF versus competing baselines, where the number of trees is varied from $[1, 300]$ for a fixed maximum tree depth of (a) 5, (b) 10 and (c) 15 on the *Chars74k* test data.




Figure 2: Classification error of RLF versus competing baselines, where the number of trees is varied from $[1, 300]$ for a fixed maximum tree depth of (a) 5, (b) 10 and (c) 15 on the *MNIST* test data.

shallow trees. Only Gradient Boosted Trees (GBT) offers a competitive result, although we note that our method far easier to train than GBT, which we found was extremely sensitive to hyperparameters such as the learning rate.

MNIST Furthermore, we compare RLFs to competing baselines for the *MNIST* dataset and observe classification error on the test data. These ablations are shown in Fig. 2. In Figs. 2(b), 2(b) and 2(c), we show the test classification error on the *MNIST* dataset when maximum tree depth is restricted to 5, 10 and 15 respectively. We can see again our RLF approach clearly outperforms all competing baselines, particularly with shallow trees. Similarly, Gradient Boosted Trees (GBT) offers a competitive result although this model was sensitive to hyperparameters when compared with our RLF models.

Consistency of Performance We also note the performance inconsistencies between competing baselines across datasets (*e.g.* Adaboost significantly underperforms on the *Chars74k* dataset, whilst offering one of the more competitive results on the *MNIST* dataset). This again highlights the sensitivity of these competing models compared to RLFs which remain relatively robust.

	<i>G50c</i>	<i>Letter</i>	<i>USPS</i>	<i>MNIST</i>	<i>Chars74k</i>
Adaboost	36.56±1.30	3.98±0.11	5.89±0.22	3.20±0.08	24.99±0.23
Random Forests	18.91±1.33	5.16±0.14	5.86±0.17	3.18±0.06	21.14±0.15
Extremely Randomized Forests	18.91±1.33	4.77±0.12	5.67±0.18	3.36±0.09	21.39±0.18
Gradient Boosted Trees	18.90±1.27	4.70±0.18	5.93±0.27	3.15±0.05	17.59±0.29
XGBoost-Large	26.40±1.22	3.95±0.14	7.22±0.32	2.98±0.06	18.32±0.23
Alternating Decision Forests	18.71±1.27	3.52±0.12	5.59±0.16	2.71±0.10	16.67±0.21
Residual Likelihood Forests (Ours)	17.75±1.20	2.59±0.06	4.89±0.09	1.81±0.03	16.33±0.25

Table 1: Overall performance of our method when compared with its main competitors on 5 datasets. The best performing methods are bolded. We train our RLF model using the same settings as described by [20] for each of the 5 datasets shown.

Dataset	Performance Error (%)				Compression Ratio			
	Refined-A	Refined-E	RLF-D	RLF-S	Refined-A	Refined-E	RLF-D	RLF-S
<i>USPS</i>	5.10±0.10	5.69±0.15	5.01±0.04	5.46±0.16	2.86	15.14	5.58	22.32
<i>Letter</i>	2.98±0.15	4.33±0.08	2.68±0.06	4.30±0.22	2.33	30.32	3.49	446.42
<i>MNIST</i>	2.05±0.02	2.95±0.03	1.81±0.03	2.41±0.05	6.29	76.92	6.98	111.61
<i>Chars74k</i>	15.40±0.1	18.00±0.09	16.33±0.25	18.51±0.17	1.70	37.04	5.75	368.30

Table 2: Residual Likelihood Forests compared with [14] for performance/model compression trade off on 4 datasets. For a fair comparison, the best performing method for each dataset comparing across Refined-A and RLF-D is **bolded**, and the best performing method comparing across Refined-E and RLF-S is **shaded in blue**. We note that when compared to the best performing accurate model of [14], our method (RLF-D) achieves higher performance on 3 out of the 4 datasets and achieves competitive results on the *Chars74k* dataset, using a much smaller fraction of the model capacity of [14]. When compared to the economic model of [14], our method is able to achieve better than or competitive performance whilst further reducing model capacity by more than an order of magnitude.

5.2 Overall Results

In Table 1, we list the overall empirical results which compare our RLF method against various ensemble baselines. We can see that RLF outperforms all other baselines across all 5 datasets. Each of the methods uses an ensemble of 100 trees of 15 tree depth, except for Gradient Boosted Trees, where we found that changing the default settings in the XGBoost [2] implementation from the default settings of 100 trees at a tree depth of 3 resulted in lower performance. For completeness, we list the results of a GBT model with 100 trees and 15 depth (XGBoost-Large) as a comparison.

5.3 Comparison with Global Refined Forests

A recent work that tries to infuse global information into the RF approach is the method of [13], where global leaf refining along with leaf pruning enables mutual information to be accounted for as well as reducing model size. However, this approach still requires the training of the full RFs before iterative refinement and pruning can occur as global information is not injected until after the initial forest construction stage has completed. We compare the accurate refined model (Refined-A) and economic refined model (Refined-E) from [13] with a shallow and deep variant of RLF (RLF-S and RLF-D respectively as shown in Table 3).

Dataset	Maximum Tree Depth			
	USPS	Letter	MNIST	Chars74k
RF/ADF	10	15	15	25
RLF-D	7	13	12	15
RLF-S	5	6	7	10

Table 3: Comparison of maximum tree depths.

In Table 2, we compare on model performance in classification error, as well as model compactness using the compression ratio defined in [10] between our approach and the approach in [10]. Similar to [10], the compression ratio is defined relative to model capacity of the original RF [9] and ADF [20] which have similar model sizes. In terms of classification error, our deep model (RLF-D) handily outperforms the accurate version of the refined model in [10] on 3 out of the 4 datasets compared against. It is worth noting that even on the *Chars74k* dataset, we achieve competitive results, despite limiting our maximum tree depth to 15 (compared to the 25 depth trees used in [10]). We note that our model demonstrates a drastic decrease in model capacity whilst only giving a minor trade off in accuracy. These results further highlight the efficiency of our RLF method in constructing compact models which offer strong classification performance.

6 Conclusion

We have proposed Residual Likelihood Forests which offers a new approach for optimizing the combining of weak learners in an ensemble learning framework. Our method shows that weak learners in an ensemble can be constructed to optimize a global loss in a complementary manner through the generation of residual likelihoods instead of probability distributions in the base tree weak learner. Empirically, we show that this allows for construction of much shallower and more compact models whilst yielding higher classification performance.

Acknowledgments The authors would like to thank the anonymous reviewers for their useful and constructive comments. This work was supported by the Australian Research Council Centre of Excellence for Robotic Vision (project number CE1401000016).

References

- [1] Anna Bosch, Andrew Zisserman, and Xavier Munoz. Image classification using random forests and ferns. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007.
- [2] Leo Breiman. Bagging predictors. *Machine learning*, 24(2):123–140, 1996.
- [3] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001.
- [4] Leo Breiman, Jerome Friedman, Richard Olshen, and Charles Stone. Classification and regression trees. wadsworth int. *Group*, 37(15):237–251, 1984.
- [5] Leo Breiman et al. Arcing classifier (with discussion and a rejoinder by the author). *The annals of statistics*, 26(3):801–849, 1998.

- [6] 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.
- [7] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794. ACM, 2016.
- [8] Yoav Freund. Boosting a weak learning algorithm by majority. In *COLT*, volume 90, pages 202–216, 1990.
- [9] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. In *European conference on computational learning theory*, pages 23–37. Springer, 1995.
- [10] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, NY, USA.:, 2001.
- [11] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [12] Juergen Gall and Victor Lempitsky. Class-specific hough forests for object detection. In *Decision forests for computer vision and medical image analysis*, pages 143–157. Springer, 2013.
- [13] Pierre Geurts, Damien Ernst, and Louis Wehenkel. Extremely randomized trees. *Machine learning*, 63(1):3–42, 2006.
- [14] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media, 2009.
- [15] Tin Kam Ho. Random decision forests. In *Document Analysis and Recognition, 1995., Proceedings of the Third International Conference on*, volume 1, pages 278–282. IEEE, 1995.
- [16] Frank Moosmann, Bill Triggs, and Frederic Jurie. Fast discriminative visual codebooks using randomized clustering forests. In *Twentieth Annual Conference on Neural Information Processing Systems (NIPS'06)*, pages 985–992. MIT Press, 2006.
- [17] Shaoqing Ren, Xudong Cao, Yichen Wei, and Jian Sun. Global refinement of random forest. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 723–730, 2015.
- [18] Robert E Schapire. The strength of weak learnability. *Machine learning*, 5(2):197–227, 1990.
- [19] Robert E Schapire. A brief introduction to boosting. In *Ijcai*, volume 99, pages 1401–1406, 1999.
- [20] Samuel Schulter, Paul Wohlhart, Christian Leistner, Amir Saffari, Peter M Roth, and Horst Bischof. Alternating decision forests. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 508–515, 2013.

- [21] Jamie Shotton, John Winn, Carsten Rother, and Antonio Criminisi. Textonboost: Joint appearance, shape and context modeling for multi-class object recognition and segmentation. In *European conference on computer vision*, pages 1–15. Springer, 2006.
- [22] Jamie Shotton, Matthew Johnson, and Roberto Cipolla. Semantic texton forests for image categorization and segmentation. In *Computer vision and pattern recognition, 2008. CVPR 2008. IEEE Conference on*, pages 1–8. IEEE, 2008.