

Ensemble Pruning Based on Objection Maximization With a General Distributed Framework

Yijun Bian, Yijun Wang^{1b}, Yaqiang Yao^{1b}, and Huanhuan Chen^{1b}, *Senior Member, IEEE*

Abstract—Ensemble pruning, selecting a subset of individual learners from an original ensemble, alleviates the deficiencies of ensemble learning on the cost of time and space. Accuracy and diversity serve as two crucial factors, while they usually conflict with each other. To balance both of them, we formalize the ensemble pruning problem as an objection maximization problem based on information entropy. Then we propose an ensemble pruning method, including a centralized version and a distributed version, in which the latter is to speed up the former. Finally, we extract a general distributed framework for ensemble pruning, which can be widely suitable for most of the existing ensemble pruning methods and achieve less time-consuming without much accuracy degradation. Experimental results validate the efficiency of our framework and methods, particularly concerning a remarkable improvement of the execution speed, accompanied by gratifying accuracy performance.

Index Terms—Composable core-sets, diversity, ensemble learning, ensemble pruning, information entropy.

I. INTRODUCTION

Thanks to its remarkable potential, ensemble learning has attracted much interest in the machine learning community [1] and has been applied widely in many real-world tasks such as object detection, object recognition, and object tracking [2]–[5]. As it is also known as committee-based learning, multiple classifier systems, or mixtures of experts [1], [6], [7], an ensemble is a set of learned models that make decisions collectively rather than relying on one single model. A variety of types of individual learners categorize an ensemble as heterogeneous ensembles and homogeneous ensembles. Most of the ensemble methods concentrate on the latter such as bagging [8] and boosting [9], [10].

The success of ensemble methods is commonly attributable to two key issues: the accuracy of individual learners and the diversity among them [11]. For classification problems, one classifier is accurate if its error rate is better than random guessing on new instances; two classifiers are diverse if they make different errors on new instances [11]. Unfortunately, there is still no consensus in the community on the definition or measurement for diversity, unlike the apparent accuracy. In addition, the diversity among individual classifiers usually decreases when these individuals approach a higher level of accuracy. Thus, how to handle the tradeoff between the two criteria is an essential issue in ensemble learning.

Although ensemble methods are effectual, one significant drawback here is that both the required memory and processing time increase visibly with the number of individual learners in the ensemble. To mitigate this shortcoming motivates ensemble pruning that aims to

Manuscript received January 7, 2019; revised June 23, 2019; accepted September 25, 2019. Date of publication November 5, 2019; date of current version September 1, 2020. This work was supported in part by the National Key Research and Development Program of China under Grant 2016YFB1000905 and in part by the National Natural Science Foundation of China under Grant 91846111 and Grant 91746209. (*Corresponding author: Huanhuan Chen.*)

The authors are with the School of Computer Science and Technology, University of Science and Technology of China (USTC), Hefei 230027, China (e-mail: yjbian@mail.ustc.edu.cn; wyjun@mail.ustc.edu.cn; yaoyaq@mail.ustc.edu.cn; hchen@ustc.edu.cn).

Color versions of one or more of the figures in this article are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNNLS.2019.2945116

2162-237X © 2019 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See <https://www.ieee.org/publications/rights/index.html> for more information.

select a subset of individual learners in an ensemble, named as ensemble selection or ensemble thinning as well [12]–[19]. It could even improve the generalization performance of an ensemble with a smaller size [20]. There has been a numerous progression on ensemble pruning methods in the past two decades. Most of the existing pruning methods, however, are centralized, in which all individual classifiers have to be stored and processed on one single machine. As the scale of data and an ensemble itself enlarge rapidly in the context of big data, the performance of centralized methods is becoming the bottleneck in execution time, which is why distributed approaches need to emerge.

To deal with ensemble pruning problems fast with balancing diversity and accuracy appropriately, we first treat ensemble pruning as an objection maximization problem using information entropy to reflect diversity and accuracy. The objective function that we aim to maximize is a tradeoff between diversity and accuracy from an information entropy perspective. Second, we transform this approach to one distributed version to speed up the execution, inspired by the emerging concept of “composable core-sets” in recent years. It adopts the same idea as a two-round divide-and-conquer strategy, which is particularly suitable for distributed settings. Third, we extract a general distributed framework for ensemble pruning from our method’s distributed version. It could be widely applicable to various existing methods for ensemble pruning and achieve less time-consuming without much accuracy degradation.

Our contribution in this brief is fourfold as follows.

- 1) We formalize the ensemble pruning problem as an objection maximization problem based on information entropy, in order to balance diversity and accuracy.
- 2) We propose an ensemble pruning method, including a centralized version and a distributed version, utilizing accuracy and diversity concurrently.
- 3) We propose a general distributed framework for ensemble pruning, which could be widely utilized and achieve less time consuming without much accuracy degradation.
- 4) We design detailed experiments to validate the effectiveness of our distributed framework and approaches.

II. RELATED WORK

In this section, we introduce diversity (i.e., a key concept in ensemble learning) and existing research on it first. Then we describe the difficulty and existing methods in ensemble pruning. Third, we introduce a concept of “composable core-sets” and its development, which sheds some light on our work. Finally, we explain the distinction between our proposed methods and other existing ensemble pruning method and summarize our contribution specifically.

A. Diversity in Ensemble Learning

Diversity, intuitively considered as the difference among individual learners in an ensemble, is a fundamental issue in ensemble methods [1], with several alternative names as dependence, orthogonality, or complementarity of learners [6]. Practically, individual classifiers are usually trained on the subsets of the same training

data, which drives them highly correlated, breaks the assumption about the independence of individual classifiers, and makes it hard to seek diversity. Numerous ensemble methods attempt to encourage diversity implicitly or heuristically [21]. For instance, boosting and bagging promote diversity by reweighting and subsampling existing training samples, respectively [8]–[10], [22], [23]; neural network (NN) ensembles also create diversity using different initial weights, different architectures of the networks, and different learning algorithms.

Unfortunately, researchers still have not reached a consensus yet on an official measure of diversity. Several measures have been proposed to represent diversity, which could be generally divided into pairwise and nonpairwise diversity measures [6], while no superior exists [7]. Based on coincident errors between a pair of individual classifiers, pairwise diversity represents the behavior whether both of them predict an instance identically or disagree with each other, including Q -statistic [24], κ -statistic [25], disagreement measure [26], [27], correlation coefficient [28], and double-fault [29]. By contrast, nonpairwise diversity is the average of all possible pairs, directly measuring a set of classifiers using the variance, entropy, or the proportion of individual classifiers that fail on instances chosen randomly. It includes interrater agreement [30], Kohavi–Wolpert variance [31], the entropy of the votes [32], [33], the difficulty index [6], [34], the generalized diversity [35], and the coincident failure diversity [35]. In addition, there are also two other measures that do not fall into these two mentioned categories previously: one is the correlation penalty function [36], measuring the diversity of each member against the entire ensemble in the negative correlation learning (NCL) [37], [38] and the other is ambiguity [39], measuring the average offset of each member against the entire ensemble output.

Moreover, few researchers could tell how diversity works exactly although the crucial role of diversity has been widely accepted in ensemble methods. In the past decade or so, Brown [40] claimed that from an information theoretic perspective, diversity within an ensemble existed indeed on numerous levels of interaction between the classifiers. His work inspired Zhou and Li [41] to propose that the mutual information should be maximized to minimize the prediction error of an ensemble from the view of multi-information. Subsequently, Yu *et al.* [21] claimed that the diversity among individual learners in a pairwise manner, used in their diversity regularized machine (DRM), could reduce the hypothesis space complexity, which implied that controlling diversity played the role of regularization in ensemble methods.

B. Ensemble Pruning

Ensemble pruning deals with the reduction of an ensemble while improving its efficiency and predictive performance [42]. Margineantu and Dietterich [43] showed the possibility to obtain nearly the same level of performance as the entire set by selecting a subset of learners from an ensemble in the first study on ensemble pruning. Zhou *et al.* [20] provided the bias-variance decomposition of error as the principal factor of the success of their approach named as “genetic algorithm based on selective ensemble (GASEN),” and claimed that pruning could lead to smaller ensembles with better generalization performance. It is difficult, however, to select the subensembles with the best generalization performance. One trouble is to estimate the generalization performance of a subensemble, and the other is that finding the optimal subset is a combinatorial search problem with exponential computational complexity [44]. Note that selecting the best combination of classifiers from an ensemble is NP-complete hard and even intractable to approximate [45].

Numerous ensemble pruning methods have been proposed to overcome shortcomings of ensemble learning over the past two

decades, which could be categorized into three general families: ranking-based, clustering-based, and optimization-based. Ranking-based pruning methods, the simplest conceptually, order the learners in the ensemble and select the first few of them according to different evaluation functions [42], including minimizing the error [e.g., orientation ordering (OO) [46]], maximizing the diversity [e.g., KL-divergence pruning and Kappa pruning (KP) [43]], or combining them both [e.g., diversity regularized ensemble pruning (DREP) [44]]. Clustering-based pruning methods employ a clustering algorithm to detect groups of learners that make similar predictions initially and then prune each cluster separately to increase the overall diversity of the ensemble [42]. Note that an intrinsic property that those methods could be executed in a parallel manner is ignored frequently in the second phase. Optimization-based pruning methods pose ensemble pruning as an optimization problem which is to find the subset of the original ensemble that optimizes a measure indicating its generalization performance. Searching exhaustively in the space of ensemble subsets is unfeasible even for a moderate ensemble size since this problem is NP-complete hard [44], [45]. Thus, various techniques are utilized to alleviate this predicament, including genetic algorithm [47], greedy algorithm [48], hill climbing [49], and biobjective evolutionary optimization [50].

C. Composable Core-Sets

Over the past few years, an effective technique, captured via the concept of “composable core-sets,” arises in order to solve optimization problems over large data sets in the distributed computing literature. Its effectiveness has been confirmed empirically for many machine learning applications, such as diverse nearest neighbor search [51], diversity maximization [52], and feature selection [53].

The notion of “composable core-sets” is introduced explicitly by Indyk *et al.* [51] for the very first time, while the notion of “core-sets” can be dated back to [54]. A core-set for an optimization problem, informally, is a subset (with a guaranteed approximation factor) of that data on which solving the underlying problem could yield an approximate solution for the original data. Composable core-sets are a collection of core-sets in which the union of them gives a core-set for the union of the original data subsets [51]. In addition, a composable core-set with α approximate factor yields a solution that is an approximation of the optimal solution for the optimization problem, and the approximation is guaranteed by a factor α , which is $1/12$ [52] and could be improved to $8/25$ [53].

D. Our Contribution

An essential distinction between our proposed methods and other existing ensemble pruning methods is that our methods could tackle the ensemble pruning task in a distributed way, accelerating the pruning process substantially. Instead of using any previous measure of diversity, our methods utilize the concept of the mutual information $I(\cdot; \cdot)$ [55], the normalized mutual information $MI(\cdot, \cdot)$, and the normalized variation information $VI(\cdot, \cdot)$ [53] from an information entropy perspective, to conduct an objective function and take diversity and accuracy into consideration implicitly and simultaneously. In addition, our proposed framework [ensemble pruning framework in a distributed setting (EPFD)] could be widely applied to various existing ensemble pruning methods to achieve less time consuming without much accuracy degradation, which is an impressive advantage of our method.

III. METHODOLOGY

In this section, we first elaborate our objection maximization based on information entropy for ensemble pruning in a centralized

way, then attain a distributed version by introducing the concept of composable core-sets, and finally extract a general distributed framework for ensemble pruning.

A. Objection Maximization Based on Information Entropy for Ensemble Pruning

Given a large data set \mathcal{D} with the size d of labeled instances obtained gradually from stream data, and their labels represented by a d -dimensional vector \mathbf{c} . A set of n trained individual classifiers $\mathcal{H} = \{h_i\}_{i=1}^n$ is considered as the original ensemble, in which each one maps the feature space of instances to the label space. The classification result vector of any individual classifier h_i from the ensemble \mathcal{H} on the data set \mathcal{D} , similar to the class label vector \mathbf{c} , is represented by a d -dimensional vector \mathbf{h}_i . The ensemble pruning task aims to find a compact subset of the original ensemble, which will predict the labels with high accuracy. These chosen individual classifiers need to be diverse and accurate simultaneously to achieve this goal. To this end, we select some diversified individual classifiers from the original ensemble, which are relevant to the vector of class labels. Hence, we strive to define a metric distance between individual classifiers in consideration of diversity and accuracy concurrently inspired by Zadeh *et al.* [53] so that the ensemble pruning problem would be reduced to an objection maximization problem.

Given two discrete random variables X and Y , Cover and Thomas [55] defined the mutual information $I(\cdot; \cdot)$ between them, that is,

$$\begin{aligned} I(X; Y) &= H(X) - H(X|Y) \\ &= \sum_{x \in X, y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)} \end{aligned} \quad (1)$$

and then Zadeh *et al.* [53] defined the normalized mutual information $MI(\cdot, \cdot)$ and the normalized variation of information $VI(\cdot, \cdot)$ of them, that is,

$$MI(X, Y) = \frac{I(X; Y)}{\sqrt{H(X)H(Y)}} \quad (2)$$

$$VI(X, Y) = 1 - \frac{I(X; Y)}{H(X, Y)} \quad (3)$$

wherein $p(\cdot, \cdot)$, $H(\cdot)$, and $H(\cdot, \cdot)$ are the joint probability, the entropy function, and the joint entropy function, respectively.

Consider the class label vector \mathbf{c} and two classification result vectors (\mathbf{h}_i and \mathbf{h}_j) generated over the data set \mathcal{D} by any two individual classifiers (h_i and h_j). In this case, the normalized mutual information $MI(\mathbf{h}_i, \mathbf{c})$ exhibits the relevance between this individual classifier h_i and the class label vector \mathbf{c} , implying the accuracy of this individual classifier on the training data set; the normalized variation of information $VI(\mathbf{h}_i, \mathbf{h}_j)$ reveals the redundancy between these two individual classifiers, indicating the diversity between them. Since class labels have already been discrete values and these values are only relevant to the number of classes in those used data sets, we do not need to discretize continuous variables to calculate the probabilities used in $MI(\cdot, \cdot)$ and $VI(\cdot, \cdot)$, while Zadeh *et al.* [53] have to deal with it.

In order to take both diversity and accuracy into consideration concurrently, the objective function between two individual classifiers [(a tradeoff between diversity and accuracy of two individual classifiers (TDAC))] is defined naturally as

$$\begin{aligned} \text{TDAC}(h_i, h_j) &= \begin{cases} \lambda VI(\mathbf{h}_i, \mathbf{h}_j) + (1 - \lambda) \frac{MI(\mathbf{h}_i, \mathbf{c}) + MI(\mathbf{h}_j, \mathbf{c})}{2}, & \text{if } h_i \neq h_j \\ 0, & \text{otherwise} \end{cases} \end{aligned} \quad (4)$$

where a regularization factor λ is introduced to balance between these two criteria, indicating their importance as well. The first criterion is to raise diversity by avoiding redundancy, and the second one is to promote accuracy by maximizing their relevance. Note that $VI(\cdot, \cdot)$ is metric [56] and $MI(\cdot, \mathbf{c})$ is nonnegative [53]. Consequently, $\text{TDAC}(\cdot, \cdot)$ is metric as well, which means $\text{TDAC}(h_i, h_j) + \text{TDAC}(h_j, h_k) \geq \text{TDAC}(h_i, h_k)$.

Subsequently, for an ensemble \mathcal{H} that is a set composed of n individual classifiers, the objection [(a tradeoff between diversity and accuracy of a set of an ensemble (TDAS))] is defined naturally as

$$\text{TDAS}(\mathcal{H}) = \frac{1}{2} \sum_{h_i \in \mathcal{H}} \sum_{h_j \in \mathcal{H}} \text{TDAC}(h_i, h_j). \quad (5)$$

Note that $\text{TDAS}(\mathcal{H})$ in (5) could be reformulated as

$$\begin{aligned} \text{TDAS}(\mathcal{H}) &= \frac{1}{2} \lambda \sum_{h_i \in \mathcal{H}} \sum_{h_j \in \mathcal{H}} VI(\mathbf{h}_i, \mathbf{h}_j) \\ &\quad + \frac{n-1}{2} (1 - \lambda) \sum_{h_i \in \mathcal{H}} MI(\mathbf{h}_i, \mathbf{c}) \end{aligned} \quad (6)$$

where $VI(\cdot, \cdot)$ of two similar individual classifiers will be near to zero. The first term in (6) prevents to select similar individual classifiers, and the second term ensures that those selected individual classifiers are relevant to the class labels. Therefore, the ensemble pruning task is reformulated as an objective function maximization problem, which aims to find a subset $\mathcal{P} \subset \mathcal{H}$ with a specified condition $|\mathcal{P}| = k$ to restrict the size of the pruned subensemble

$$\max_{\substack{\mathcal{P} \subset \mathcal{H}, \\ |\mathcal{P}|=k}} \text{TDAS}(\mathcal{P}) = \max_{\substack{\mathcal{P} \subset \mathcal{H}, \\ |\mathcal{P}|=k}} \frac{1}{2} \sum_{h_i \in \mathcal{P}} \sum_{h_j \in \mathcal{P}} \text{TDAC}(h_i, h_j). \quad (7)$$

It is a biobjective optimization problem to combine diversity and accuracy concurrently for ensemble pruning tasks, but we could manage to reformulate it as a single-objective optimization problem by introducing $\text{TDAC}(\cdot, \cdot)$ in (4). That is one way to solve the problem by using the objective weighting [57]. Another way could be introducing a concept of domination [50], [58] to obtain a Pareto optimal solution, which we will leave it for future work. In addition, our solution manages to fix the size of the pruned subensemble while the latter solution cannot, which usually leads to oversize subensembles and affects the space cost.

Algorithm 1 COMEP

Require: Set of an original ensemble \mathcal{H} , threshold k as the size of the pruned sub-ensemble.

Ensure: Set of the pruned sub-ensemble \mathcal{P} satisfying that $\mathcal{P} \subset \mathcal{H}$ and $|\mathcal{P}| \leq k$.

- 1: $\mathcal{P} \leftarrow$ an arbitrary individual classifier $h_i \in \mathcal{H}$.
 - 2: **for** $2 \leq i \leq k$ **do**
 - 3: $h^* \leftarrow \operatorname{argmax}_{h_i \in \mathcal{H} \setminus \mathcal{P}} \sum_{h_j \in \mathcal{P}} \text{TDAC}(h_i, h_j)$.
 - 4: Move h^* from \mathcal{H} to \mathcal{P} .
 - 5: **end for**
-

Up until now, we have modeled the ensemble pruning task through an objective function maximization problem as shown in (7), which is enough to form a centralized algorithm to accomplish this goal of ensemble pruning. This centralized method, named as ‘‘centralized objection maximization for ensemble pruning (COMEP)’’ shown in Algorithm 1, selects greedily the current optimal classifier at each step and could achieve a 1/2 approximation factor for objective function maximization problem according to [59].

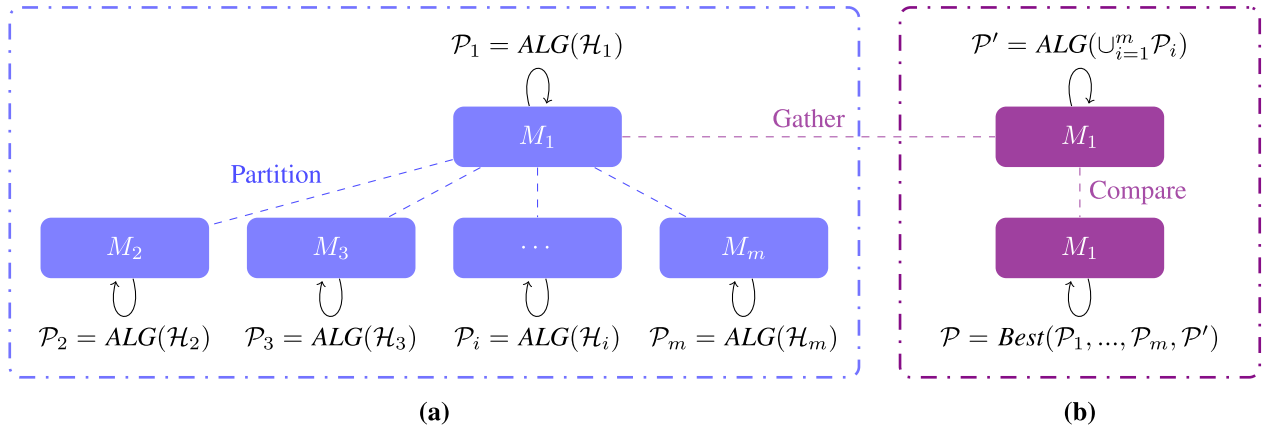


Fig. 1. Diagram of Algorithm 3: EPFD. (a) In the first phase, a primary machine (e.g., M_1) partitions all individual classifiers in the original ensemble into m groups $\{\mathcal{H}_i\}_{i=1}^m$ randomly and allocates them to different machines. For each i ($1 \leq i \leq m$), the machine M_i runs ALG on its allocated set \mathcal{H}_i independently and selected a subset \mathcal{P}_i from it in parallel. (b) In the second phase, the primary machine (e.g., M_1) gathers all subsets, runs ALG on their union $\cup_{i=1}^m \mathcal{P}_i$ to produce a subset \mathcal{P}' , and eventually outputs the best one of them by comparing \mathcal{P}' with \mathcal{P}_i ($1 \leq i \leq m$). Remark: (1) ALG could be any one of existing pruning method, including COMEP. (2) The final \mathcal{P} is chosen by $Best(\cdot)$ according to some certain criteria such as accuracy or TDAS(\cdot) in DOMEF.

B. Distributed Objection Maximization for Ensemble Pruning

“Distributed objection maximization for ensemble pruning (DOMEF),” the distributed version of COMEP shown in Algorithm 2, adopts a two-round divide-and-conquer strategy and composable core-sets [60] as guidelines, which are particularly suitable for distributed settings. It partitions a set of individual classifiers of an ensemble into smaller pieces, solves the ensemble pruning problem on each piece separately, and eventually obtains a subset from the union of these representative subsets for all pieces.

Algorithm 2 DOMEF

Require: Set of an original ensemble \mathcal{H} , threshold k as the size of the pruned sub-ensemble, number of machines m .

Ensure: Set of the pruned sub-ensemble \mathcal{P} meeting that $\mathcal{P} \subset \mathcal{H}$ and $|\mathcal{P}| \leq k$.

- 1: Partition \mathcal{H} randomly into m groups as equally as possible, i.e., $\mathcal{H}_1, \dots, \mathcal{H}_m$.
 - 2: **for** $1 \leq i \leq m$ **do**
 - 3: $\mathcal{P}_i \leftarrow \text{COMEP}(\mathcal{H}_i, k)$.
 - 4: **end for**
 - 5: $\mathcal{P}' \leftarrow \text{COMEP}(\cup_{i=1}^m \mathcal{P}_i, k)$.
 - 6: $\mathcal{P} \leftarrow \arg\max_{\mathcal{T} \in \{\mathcal{P}_1, \dots, \mathcal{P}_m, \mathcal{P}'\}} \text{TDAS}(\mathcal{T})$.
-

Consider a set of n trained individual classifiers $\mathcal{H} = \{h_i\}_{i=1}^n$ as the original ensemble. In the first phase, a primary machine partitions all individual classifiers in the original ensemble into m groups $\{\mathcal{H}_i\}_{i=1}^m$ randomly and allocates them to different machines. Note that $\cup_{i=1}^m \mathcal{H}_i = \mathcal{H}$, where m is the number of machines, and that the primary machine could be any one of these available machines. For each i ($1 \leq i \leq m$), machine i runs COMEP on its allocated set \mathcal{H}_i independently and selects a subset \mathcal{P}_i from it in parallel. In the second phase, the primary machine gathers all subsets, runs COMEP on their union $\cup_{i=1}^m \mathcal{P}_i$ to produce a subset \mathcal{P}' , and eventually outputs the best one of them by comparing \mathcal{P}' with \mathcal{P}_i ($1 \leq i \leq m$) according to (7). It suffices to output the satisfying subset \mathcal{P} after these two phases (Lines 1–5 in Algorithm 2) in practice and the additional comparison purposes to get a higher approximation factor which is $1/4$ theoretically and could even reach $8/25$ under some special conditions [53].

Finally, there might be some confusion over how to partition n individual classifiers in \mathcal{H} randomly into m groups as equally

as possible. The first step is to shuffle these individual classifiers with a different random order, and the second step is to arrange them to get m groups. There would be no doubt when n is a multiple of m in which each group would contain n/m individual classifiers. However, if n is not a multiple of m , there would be $(n \bmod m)$ groups of them that each contains $\lceil n/m \rceil$ individual classifiers and $(n \bmod m)$ groups of them that each contains $\lfloor n/m \rfloor$ individual classifiers. Note that $n \bmod m = m \lceil n/m \rceil - n$ according to that in [61].

C. General Distributed Framework for Ensemble Pruning

A general distributed framework is extracted from DOMEF, named as “ensemble pruning framework in a distributed setting (EPFD)” shown in Algorithm 3 and Fig. 1, which likewise adopts the two-round divide-and-conquer strategy and composable core-sets [60]. It enables the ensemble pruning problem to be solved fast in a distributed way. Ensemble pruning is usually described as a process to acquire the optimum subset from the original ensemble. Let ALG denote an arbitrary algorithm to perform this task and \mathcal{H} denote the original ensemble. EPFD consists of two main phases just like DOMEF that could be regarded as a special case of EPFD where COMEP is chosen as the used pruning method (Lines 3 and 5 in Algorithm 3). Another key difference is that the criterion (Line 6 in Algorithm 3) here is not limited to (7). For instance, it could use accuracy or any other measures corresponding to data to compare different subsets. EPFD is a simple yet powerful tool to accelerate the original methods for ensemble pruning without much performance degradation, which is elaborated in Section IV-C.

Algorithm 3 EPFD

Require: Set of an original ensemble \mathcal{H} , number of machines m , a pruning method ALG.

Ensure: Set of the pruned sub-ensemble \mathcal{P} meeting that $\mathcal{P} \subset \mathcal{H}$.

- 1: Partition \mathcal{H} into $\{\mathcal{H}_i\}_{i=1}^m$ randomly.
 - 2: **for** $1 \leq i \leq m$ **do**
 - 3: $\mathcal{P}_i \leftarrow$ output from any pruning method ALG on \mathcal{H}_i .
 - 4: **end for**
 - 5: $\mathcal{P}' \leftarrow$ output from ALG on $\cup_{i=1}^m \mathcal{P}_i$.
 - 6: $\mathcal{P} \leftarrow$ the best one among $\mathcal{P}_1, \dots, \mathcal{P}_m$, and \mathcal{P}' according to some certain criteria such as accuracy.
-

TABLE I
COMPARISON OF THE STATE-OF-THE-ART METHODS WITH COMEP AND DOME P USING BAGGING TO
PRODUCE AN ENSEMBLE WITH DTs AS INDIVIDUAL CLASSIFIERS

Dataset	KL	KP	OO	RE	DREP	SEP	OEP	PEP	COMEP	DOME P
Iono	89.43±3.44	90.29±3.41	90.57±4.69	88.57±3.03	89.71±3.70	90.86±2.96	89.71±4.33	91.71±3.41	91.14±3.41	91.71±3.56
Liver	61.45±8.24	62.32±5.12	64.06±7.35	61.16±4.96‡	56.23±8.41	60.87±4.35‡	64.64±4.54	62.32±4.23	62.90±4.18	64.64±4.30
Spam	93.28±1.13	93.21±0.93	93.30±1.43	93.54±1.16	92.08±1.48	93.43±0.82	93.56±1.24	93.38±0.80	93.41±0.90	93.45±1.00
Wisconsin	94.37±3.58	94.67±3.25‡	95.56±4.06	95.41±3.25	93.93±3.68	94.96±3.86‡	95.26±3.50	95.41±3.79	95.56±3.70	95.70±3.45
Credit	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00	100.00±0.00
Landsat	97.34±0.45	97.22±0.66	97.29±0.37	97.31±0.28	96.08±1.06‡	97.26±0.26	97.34±0.35	97.29±0.56	97.15±0.43	97.42±0.34
Wilt	98.24±0.32	98.24±0.26	98.37±0.42	98.28±0.32	98.01±0.51	98.37±0.51	98.16±0.18	98.39±0.53	98.41±0.33	98.35±0.38
Shuttle	99.96±0.03	99.96±0.03	99.97±0.02	99.97±0.02	99.97±0.02	99.96±0.03	99.96±0.03	99.97±0.02	99.96±0.03	99.97±0.03
Ecoli	94.55±3.14	94.85±2.03	93.64±2.71	95.15±2.49	91.52±3.49‡	94.85±2.75	93.03±2.30	94.24±2.71	95.15±2.49	95.15±1.98
SensorReadings	99.25±0.40	99.19±0.50	99.41±0.47	99.52±0.24	99.28±0.25	99.45±0.23	99.50±0.19	99.36±0.43	99.43±0.33	99.38±0.41
<i>t</i> -Test (W/T/L)	0/10/0	1/9/0	0/10/0	1/9/0	2/8/0	2/8/0	0/10/0	0/10/0	0/10/0	—
Average Rank	7.15	7.20	4.80	4.60	8.40	5.55	5.20	4.75	4.40	2.95

¹ The reported results are the average test accuracy of each method and the corresponding standard deviation under 5-fold cross-validation on each data set. The best results with higher accuracy and lower standard deviation are highlighted in boldface.

² By two-tailed paired *t*-test at 5% significance level, ‡ and † denote that the performance of *DOME P* is superior to and inferior to that of the comparative method, respectively.

³ The last two rows show the results of *t*-test and average rank, respectively. The “W/T/L” in *t*-test indicates that *DOME P* is superior to, not significantly different from, or inferior to the corresponding comparative methods. The average rank is calculated according to the Friedman test [64].

IV. EXPERIMENTS

In order to evaluate our proposed methods, in this section, we elaborate our experiments on 17 binary and 12 multiclass data sets, including an image data set with 12500 pictures (Dogs vs. Cats¹) and 28 data sets from UCI repository [62]. Standard fivefold cross validation is used in these experiments where the entire data set is split into three parts in each iteration, with 60% as the training set, 20% as the validation set, and 20% as the test set. In addition, we construct homogeneous ensembles using Bagging on various types of classifiers including decision trees (DT), naive Bayesian (NB) classifiers, *k*-nearest neighbors (KNN) classifiers, linear model (LM) classifiers, and linear support vector machines (LSVMs). An ensemble is first trained on the training set, then pruned by a pruning method on the validation set, and finally tested on the test set. The baselines that we consider are a variety of ranking-based methods as well as optimization-based methods. These ranking-based methods include KL-divergence (KL) pruning, KP [43], OO pruning [46], reduce-error (RE) pruning [63], DREP [44], and ordering-based ensemble pruning (OEP); These optimization-based methods are single-objective ensemble pruning (SEP) and Pareto ensemble pruning (PEP) [50]. Note that several methods cannot fix the number of learners after ensemble pruning (such as OO, DREP, SEP, OEP, and PEP), while others could fix it by giving a pruning rate, which is the up limit of the percentage of those discarded individual classifiers in the original ensemble. Those methods that cannot fix the size might lead to oversize or undersize subensembles and affect their space cost. Due to space constraints, we only report the comparisons of the time cost and the test accuracy hereinafter.

A. Comparison of COMEP and DOME P to the State-of-the-Art Ensemble Pruning Methods

In this section, we compare the quality of various ensemble pruning methods (the original centralized version) including KL, KP, OO, RE, DREP, SEP, OEP, and PEP with our proposed centralized (COMEP) and distributed (DOME P) methods.

Experimental results reported in Table I contain the average test accuracy of each method and the corresponding standard deviation under fivefold cross validation on each data set. Each row in Table I




Fig. 2. Comparison of the state-of-the-art methods with COMEP and DOME P on the test accuracy. (a) Friedman test chart (nonoverlapping means significant difference) [64]. (b) Aggregated rank for each method (the smaller the better) [50].

compares the classification accuracy using bagging with the same type of individual classifiers. The results with higher accuracy and lower standard deviation are indicated with bold fonts for each data set (row). In addition, we examine the significance of the difference in the accuracy performance between two ensemble pruning methods by two-tailed paired *t*-test at 5% significance level to tell whether these two methods have significantly different results. Two methods end in a tie if there is no significant statistical difference; otherwise, the one with higher values of accuracy will win. In the last two rows of Table I, the average rank of each method is calculated based on the Friedman test [64], and the performance of each method is compared with DOME P in terms of the number of data sets that DOME P has won, tied, or lost, respectively. It can be inferred that DOME P does not underperform centralized methods in many data sets even though it only utilizes local information unlike others, which confirms the reasonableness of DOME P (and COMEP) utilizing accuracy and diversity simultaneously. Despite slightly lower values of accuracy in some cases, DOME P remains acceptable results. Similar results are reported in Tables II and III using different individual classifiers. Moreover, Fig. 2 reports the comparison of the state-of-the-art methods with COMEP and DOME P on the test accuracy using statistical test methods [50], [64]. Fig. 2(a) shows that COMEP and DOME P have significant superiority over other compared centralized methods, utilizing the Friedman test chart [64];

¹<http://www.kaggle.com/c/dogs-vs-cats>

TABLE II
COMPARISON OF THE STATE-OF-THE-ART METHODS WITH COMEP AND DOMEPE USING BAGGING TO PRODUCE AN ENSEMBLE WITH SUPPORT VECTOR MACHINES (SVMs) AS INDIVIDUAL CLASSIFIERS

Dataset	KL	KP	OO	RE	DREP	SEP	OEP	PEP	COMEP	DOMEPE
Liver	58.84±1.65	58.84±1.65	58.26±1.21	58.84±1.65	58.84±1.65	58.55±1.65	58.84±1.65	58.55±1.65	58.84±1.65	58.84±1.65
Ringnorm	98.43±0.37	98.53±0.30†	98.51±0.36†	98.47±0.29	98.43±0.29	98.47±0.36	98.39±0.25	98.44±0.31	98.43±0.31	98.43±0.31
Waveform	91.45±1.60	91.23±1.50	91.35±1.58	91.17±1.67	90.83±1.52‡	91.27±1.53	91.13±1.61	91.37±1.63	91.45±1.49	91.23±1.63
Credit	78.02±0.07	78.01±0.11	78.03±0.08	78.02±0.08	77.96±0.12	78.04±0.12	78.02±0.11	78.02±0.13	78.00±0.09	78.01±0.11
Landsat	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00	65.24±0.00
Page	91.26±0.41	91.26±0.45	91.28±0.44	91.24±0.46	91.28±0.52	91.22±0.32	91.24±0.44	91.30±0.53	91.26±0.45	91.22±0.44
Wilt	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09	94.68±0.09
SensorReadings	89.17±0.78	89.24±0.88	89.28±0.52	89.35±0.51	88.91±0.60	89.11±0.91‡	89.26±0.50	89.51±0.94	89.42±0.83	89.37±1.04
EEGEyeState	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00	55.13±0.00
WaveformNoise	86.35±0.74	86.11±0.83	86.37±0.79	86.21±0.94	85.99±0.83	86.49±1.02	86.11±0.98	86.29±0.77	86.55±0.68	86.43±0.81
<i>t</i> -Test (W/T/L)	0/10/0	0/9/1	0/9/1	0/10/0	1/9/0	1/9/0	0/10/0	0/10/0	0/10/0	—
Average Rank	5.20	5.60	4.60	5.50	7.05	5.50	6.60	4.55	4.65	5.75

TABLE III
COMPARISON OF THE STATE-OF-THE-ART METHODS WITH COMEP AND DOMEPE USING BAGGING TO PRODUCE AN ENSEMBLE WITH KNNs AS INDIVIDUAL CLASSIFIERS

Dataset	KL	KP	OO	RE	DREP	SEP	OEP	PEP	COMEP	DOMEPE
Ames	58.83±5.48	58.54±7.20	60.29±7.45	57.81±7.14	57.52±7.29	59.71±6.18	59.56±6.96	58.69±5.90	60.15±6.90	59.12±5.95
Card	67.59±2.10	66.86±1.68	66.72±1.68	67.01±2.27	64.38±2.27	67.88±2.87	65.99±1.22	65.55±2.39‡	68.32±3.33	68.03±3.44
Sonar	80.98±7.98	80.00±8.52	81.95±7.03	81.95±5.62	79.02±5.62	81.46±5.87	80.00±5.29	78.05±5.72	82.44±7.19	80.98±8.69
Page	95.87±0.35	95.87±0.49	95.81±0.45	95.76±0.35	95.76±0.39	95.80±0.57	95.87±0.55	95.81±0.46	95.87±0.64	95.89±0.64
Wilt	97.93±0.28	97.89±0.32	98.01±0.24	97.99±0.20	97.77±0.36	97.93±0.35	97.89±0.40	98.01±0.32	97.97±0.19	98.01±0.26
Landsat	90.28±0.88	90.22±1.18	90.30±0.95	90.40±0.83	90.31±0.94	90.47±0.78	90.37±0.85	90.37±0.76	90.47±1.21	90.44±0.87
Shuttle	99.82±0.03	99.81±0.03	99.82±0.04	99.82±0.02	99.82±0.04	99.81±0.04	99.81±0.03	99.81±0.03	99.80±0.03‡	99.82±0.04
Ecoli	95.15±2.91	94.55±3.14	95.76±2.91	95.45±3.03	93.33±3.14	94.85±2.95	94.24±3.11	94.85±3.65	95.45±3.03	95.76±2.91
WaveformNoise	84.34±1.05	84.20±1.46	84.48±2.09	84.36±2.24	83.54±0.74	84.36±1.33	84.48±2.24	84.16±1.58	84.48±1.81	84.62±1.77
EEGEyeState	95.99±0.55	96.06±0.52	96.16±0.50	95.98±0.59	95.03±0.37‡	96.11±0.50	96.08±0.66	96.28±0.55	96.19±0.62	96.01±0.46
SensorReadings	96.29±0.47	96.11±0.39	96.33±0.62	96.37±0.53	96.00±0.51	96.29±0.64	96.39±0.42	96.28±0.60	96.42±0.64	96.31±0.31
<i>t</i> -Test (W/T/L)	0/11/0	0/11/0	0/11/0	0/11/0	1/10/0	0/11/0	0/11/0	1/10/0	1/10/0	—
Average Rank	5.82	7.45	3.77	5.27	8.95	5.09	5.77	6.55	3.05	3.27

Fig. 2(b) represents the aggregated rank for each method, depicting the same results.

B. DOMEPE Versus COMEP Over the Time Cost

In this section, we give the complexity analysis of the proposed methods and present the corresponding experimental results. According to Algorithm 1, the computational complexity of COMEP is analyzed as follows.

- 1) First, the complexity of an essential term of TDAC(h_i, h_j) is $\mathcal{O}(d^2n_c^2)$, where d and n_c are the number of instances and labels, respectively.
- 2) Second, the chosen individual classifier is obtained with the complexity of $\mathcal{O}(-(1/3)k^3 + (1/2)nk^2)$ times that of the TDAC(h_i, h_j) term since $\sum_{i=2}^k (i-1)(n-i+2) = -(1/3)k^3 + (n+2/2)k^2 - (3n+4/6)k$, where k is the size of the pruned subensemble.

Therefore, the overall computational complexity of COMEP is $\mathcal{O}((-(1/3)k^3 + (1/2)nk^2)d^2n_c^2)$. In addition, since \mathcal{H} is partitioned into m groups in DOMEPE, the overall computational complexity of DOMEPE is $\mathcal{O}((-(1/3)k^3 + (n/2m)k^2)d^2n_c^2)$.

In this experiment, we employ different numbers of machines in DOMEPE in order to test its speedup in comparison with COMEP. Under the ideal conditions, Zadeh *et al.* [53] pointed that the speedup between the distributed and centralized version was almost linear in terms of the number of used machines since there was no overhead of information-sharing between those machines. To verify whether DOMEPE could achieve competitive performance faster

than COMEP or not, we introduce two performance indicators from parallel processing, i.e., speedup and efficiency,² where speedup is defined as a quotient of the execution time of COMEP and that of DOMEPE, to report how much speedup of DOMEPE over COMEP. Constrained by the capability of machines on which we test, several experiments for ensemble pruning conducted on two or three machines are used as a typical example to present the performance of DOMEPE. The results using various settings of this experiment are summarized in Fig. 3, which indicates that DOMEPE runs faster than COMEP even reaching super-linear speedup³ on a tiny minority of experiments.

C. Comparison Between the State-of-the-Art Ensemble Pruning Methods and Their Corresponding Distributed Versions Generated With EPFD

In this section, to verify the effectiveness of EPFD, we compare the quality of various centralized ensemble pruning methods and their respective distributed versions that are generated using EPFD in terms of accuracy and time cost. To test the quality of the selected subensembles of each method, we control them under the same conditions (including the employed ensemble methods or types of

²Efficiency is a measure of how effectively parallel computing could be used to solve a particular problem. A parallel algorithm is considered cost efficient if its asymptotic running time multiplied by the number of processing units involved in the computation is comparable to the running time of the best sequential algorithm [65].

³Sometimes, a speedup of more than m when using m processors is observed in parallel computing, which is called super-linear speedup [65].

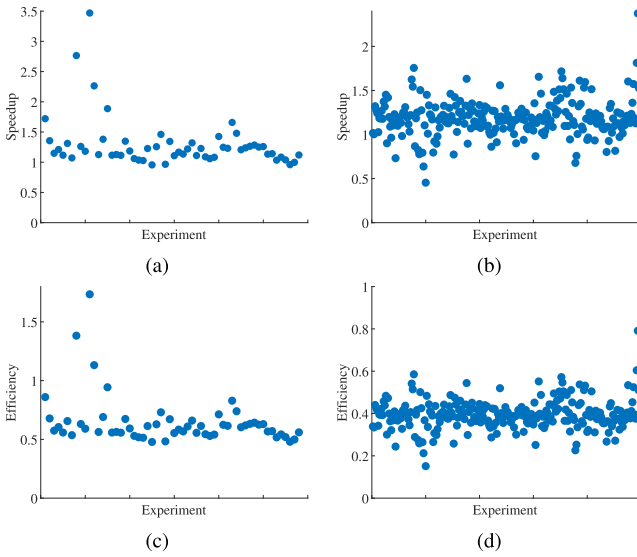


Fig. 3. Comparison of speedup and efficiency between COMEP and DOMEPE. (a) Speedup with two machines. (b) Speedup with three machines. (c) Efficiency with two machines. (d) Efficiency with three machines.




Fig. 4. Comparison between the state-of-the-art ensemble pruning methods and their distributed versions on the SensorReadings data set using Bagging with DTs as individual classifiers for binary classification. (a) Accuracy. (b) Time cost.




Fig. 5. Comparison between the state-of-the-art ensemble pruning methods and their distributed versions on the SensorReadings data set using Bagging with KNNs as individual classifiers for binary classification. (a) Accuracy. (b) Time cost.

individual classifiers) in each experiment. Fig. 4 shows the comparison results when individual classifiers are designated as DTs and assembled by bagging for binary classification. It can be inferred that, for each pruning method (each group on the horizontal axis), the accuracy of the distributed version is superior or equal to that of its corresponding centralized version. In consideration of the less time cost it takes, we believe that the distributed version of each method outperforms its original centralized version. Besides, we can tell that the effectiveness of EPFD is exceptionally evident on PEP, i.e., a complicated method utilizing an evolutionary Pareto optimization combined with a local search subroutine. Moreover, Fig. 5 reports the similar results of that with KNNs as individual classifiers for binary classification; Figs. 6 and 7 report similar results of that with




Fig. 6. Comparison between the state-of-the-art ensemble pruning methods and their distributed versions on the SensorReadings data set using Bagging with DTs as individual classifiers for multi-class classification. (a) Accuracy. (b) Time cost.




Fig. 7. Comparison between the state-of-the-art ensemble pruning methods and their distributed versions on the SensorReadings data set using Bagging with KNNs as individual classifiers for multiclass classification. (a) Accuracy. (b) Time cost.

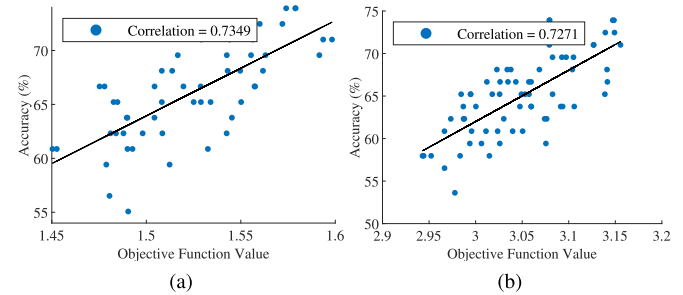


Fig. 8. Relation of binary classification accuracy and objective function value for 3-combinations and 4-combinations in the Ames data set using Bagging with DTs as individual classifiers for binary classification. (a) 3-combinations. (b) 4-combinations.

DTs and KNNs as individual classifiers for multiclass classification, respectively.

D. Validating the Objective Function

Regarding a new objective function, its relationship with classification accuracy is one of the fundamental questions. We select two small-sized ensembles (small in the number of individual classifiers) and evaluate all possible combinations of these individual classifiers in order to test this issue. In this experiment, we compare the classification accuracy for all the 3-combinations and 4-combinations of individual classifiers in the original ensemble against their corresponding objective value with the λ parameter equal to 0.5, which means that two criteria in (4) are equally important. Each small blue dot in Fig. 8 represents the classification accuracy on a 3-combination or 4-combination of the individual classifiers with the size 8 of an ensemble in the Ames data set for binary classification, and the line is the regression line. Similar results are shown in Figs. 9 and 10 in the Waveform data set for binary classification and multi-class classification, respectively. We observe that the objective value and the classification accuracy are highly correlated from Figs. 8–10,

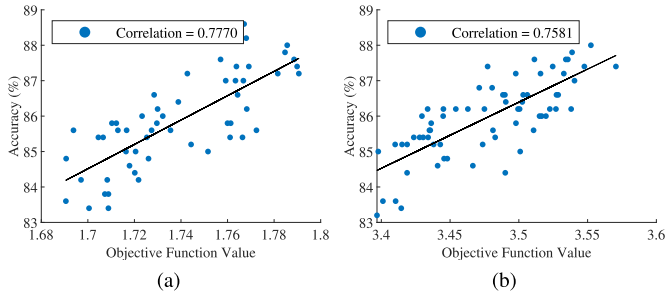


Fig. 9. Relation of binary classification accuracy and objective function value for 3-combinations and 4-combinations in the Waveform data set using Bagging with DTs as individual classifiers for binary classification. (a) 3-combinations. (b) 4-combinations.

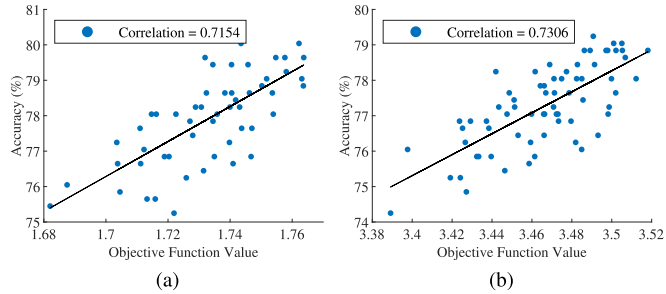


Fig. 10. Relation of multiclass classification accuracy and objective function value for 3-combinations and 4-combinations in the Waveform data set using Bagging with DTs as individual classifiers for multi-classification. (a) 3-combinations. (b) 4-combinations.

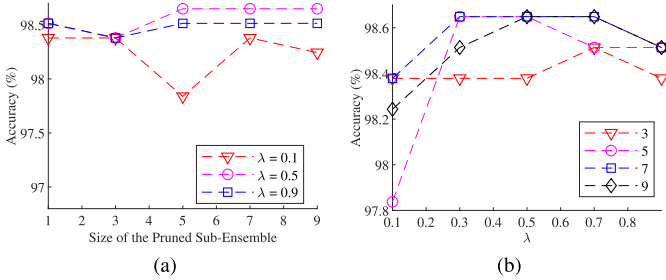


Fig. 11. Effect of λ value on the classification accuracy in the Ringnorm data set using Bagging with SVMs as individual classifiers for binary classification. (a) Accuracy of each criterion individually. (b) Slight differences of λ value while selecting the different size of the pruned subensemble (3, 5, 7, 9).

which means that maximizing this objective function leads to our target (i.e., the highly accurate subensembles).

E. Effect of λ Value

Crucial as other issues, the relation of two criteria needs to be investigated in the defined objective function. To reveal how the classification results are affected with the regularization factor λ , different λ values (from 0.1 to 0.9 with 0.2 steps) are tested in the experiments of this part. Fig. 11 exemplifies the effect of λ on the Ringnorm data set. Fig. 11(a) illustrates that the linear combination concurrently considering them both performs better than focusing more on MI term ($\lambda = 0.1$) or VI term ($\lambda = 0.9$) in (6) individually, although finding the optimal value of λ is another challenge. Fig. 11(b) presents that a global maximum around the optimal λ exists, indeed, regardless of the size of the pruned subensemble, which suggests that it might be related to the intrinsic properties of the data set. Similar results are reported in Figs. 12 and 13 for binary classification and multiclass classification, respectively. Although proper results for all data sets have been brought with λ being set to 0.5 (in Tables I–III) for convenience, DOMEF would achieve a better performance in practice when the λ is adjusted for each data set separately.

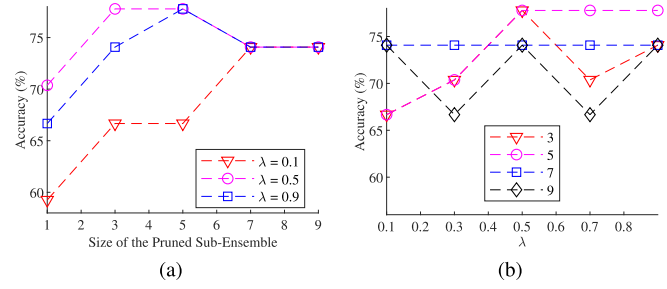


Fig. 12. Effect of λ value on the classification accuracy in the Heart data set using Bagging with KNNs as individual classifiers for binary classification. (a) Accuracy of each criterion individually. (b) Slight differences of λ value while selecting the different size of the pruned subensemble (3, 5, 7, 9).

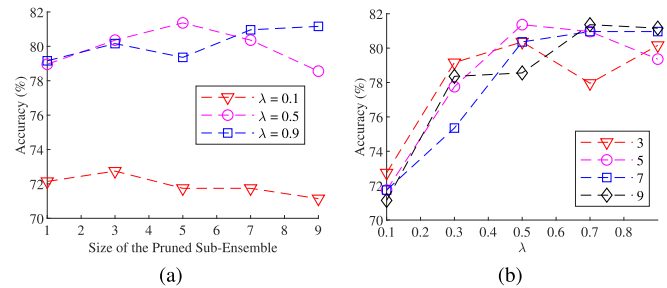


Fig. 13. Effect of λ value on the classification accuracy in the Waveform data set using Bagging with DTs as individual classifiers for multi-class classification. (a) Accuracy of each criterion individually. (b) Slight differences of λ value while selecting the different size of the pruned subensemble (3, 5, 7, 9).

V. CONCLUSION

In this brief, we formalize the ensemble pruning problem as an objection maximization problem based on information entropy to consider diversity and accuracy simultaneously. Then, we propose an ensemble pruning method according to this objection maximization problem (including two versions, COMEP and DOMEF) for ensemble pruning. We also present that our methods (COMEP and DOMEF) are consistently competitive with various existing methods for ensemble pruning, which could handle large-scale ensembles fast yet efficiently through handling the accuracy and diversity of the ensembles properly. Finally, we propose a general distributed framework (EPFD) for ensemble pruning, which could be widely applied to various existing methods for ensemble pruning, to achieve less time consuming without much accuracy degradation. The remarkable effectiveness of EPFD is positively valuable for enormous data in the real world. For future work, it seems like a promising direction to explore the deeper theoretical basis and to try other objective functions to achieve better performance.

REFERENCES

- [1] Z.-H. Zhou, *Ensemble Methods: Foundations and Algorithms*. Boca Raton, FL, USA: CRC Press, 2012.
- [2] R. Girshick, J. Donahue, T. Darrell, and J. Malik, “Rich feature hierarchies for accurate object detection and semantic segmentation,” in *Proc. CVPR*, Jun. 2014, pp. 580–587.
- [3] J. Wang, Z. Liu, Y. Wu, and J. Yuan, “Mining actionlet ensemble for action recognition with depth cameras,” in *Proc. CVPR*, Jun. 2012, pp. 1290–1297.
- [4] X. Zhou, L. Xie, P. Zhang, and Y. Zhang, “An ensemble of deep neural networks for object tracking,” in *Proc. ICIP*, Oct. 2014, pp. 843–847.
- [5] H. Ykhlef and D. Bouchaffra, “An efficient ensemble pruning approach based on simple coalitional games,” *Inf. Fusion*, vol. 34, pp. 28–42, Mar. 2017.
- [6] L. I. Kuncheva and C. J. Whitaker, “Measures of diversity in classifier ensembles and their relationship with the ensemble accuracy,” *Mach. Learn.*, vol. 51, no. 2, pp. 181–207, 2003.

- [7] E. K. Tang, P. N. Suganthan, and X. Yao, "An analysis of diversity measures," *Mach. Learn.*, vol. 65, no. 1, pp. 247–271, 2006.
- [8] L. Breiman, "Bagging predictors," *Mach. Learn.*, vol. 24, no. 2, pp. 123–140, 1996.
- [9] Y. Freund, "Boosting a weak learning algorithm by majority," *Inf. Comput.*, vol. 121, no. 2, pp. 256–285, Sep. 1995.
- [10] Y. Freund and R. E. Schapire, "Experiments with a new boosting algorithm," in *Proc. Int. Conf. Mach. Learn.*, Bari, Italy, vol. 96, 1996, pp. 148–156.
- [11] T. G. Dietterich, "Ensemble methods in machine learning," in *Proc. 1st Int. Workshop Multiple Classifier Syst.*, in Lecture Notes in Computer Science, vol. 1857. Berlin, Germany: Springer, 2000, pp. 1–15.
- [12] Z. Liu, Q. Dai, and N. Liu, "Ensemble selection by GRASP," *Appl. Intell.*, vol. 41, no. 1, pp. 128–144, 2014.
- [13] Y. Kokkinos and K. G. Margaritis, "Confidence ratio affinity propagation in ensemble selection of neural network classifiers for distributed privacy-preserving data mining," *Neurocomputing*, vol. 150, pp. 513–528, Feb. 2015.
- [14] C.-X. Zhang, J.-S. Zhang, and Q.-Y. Yin, "A ranking-based strategy to prune variable selection ensembles," *Knowl.-Based Syst.*, vol. 125, pp. 13–25, Jun. 2017.
- [15] R. Lysiak, M. Kurzynski, and T. Woloszynski, "Optimal selection of ensemble classifiers using measures of competence and diversity of base classifiers," *Neurocomputing*, vol. 126, pp. 29–35, Feb. 2014.
- [16] I. Partalas, G. Tsoumakas, and I. P. Vlahavas, "Focused ensemble selection: A diversity-based method for greedy ensemble selection," in *Proc. ECAI*, 2008, pp. 117–121.
- [17] R. Caruana, A. Niculescu-Mizil, G. Crew, and A. Ksikes, "Ensemble selection from libraries of models," in *Proc. ICML*, 2004, p. 18.
- [18] R. E. Banfield, L. O. Hall, K. W. Bowyer, and W. P. Kegelmeyer, "Ensemble diversity measures and their application to thinning," *Inf. Fusion*, vol. 6, no. 1, pp. 49–62, 2005.
- [19] H. Chen, P. Tiño, and X. Yao, "Predictive ensemble pruning by expectation propagation," *IEEE Trans. Knowl. Data Eng.*, vol. 21, no. 7, pp. 999–1013, Jul. 2009.
- [20] Z.-H. Zhou, J. Wu, and W. Tang, "Ensembling neural networks: Many could be better than all," *Artif. Intell.*, vol. 137, no. 1, pp. 239–263, 2002.
- [21] Y. Yu, Y.-F. Li, and Z.-H. Zhou, "Diversity regularized machine," in *Proc. IJCAI*, 2011, p. 1603.
- [22] P. Melville and R. J. Mooney, "Constructing diverse classifier ensembles using artificial training examples," in *Proc. IJCAI*, 2003, pp. 505–510.
- [23] R. G. F. Soares, H. Chen, and X. Yao, "A cluster-based semisupervised ensemble for multiclass classification," *IEEE Trans. Emerg. Topics Comput. Intell.*, vol. 1, no. 6, pp. 408–420, Dec. 2017.
- [24] G. U. Yule, "On the association of attributes in statistics: With illustrations from the material of the childhood society, &c," *Philos. Trans. Roy. Soc. London A, Math. Phys. Sci.*, vol. 194, no. 1900, pp. 257–319, 1900.
- [25] J. Cohen, "A coefficient of agreement for nominal scales," *Educ. Psychol. Meas.*, vol. 20, no. 1, pp. 37–46, 1960.
- [26] D. B. Skalak, "The sources of increased accuracy for two proposed boosting algorithms," in *Proc. AAAI, Integrating Multiple Learned Models Workshop*, Portland, OR, USA, 1996.
- [27] T. K. Ho, "The random subspace method for constructing decision forests," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 20, no. 8, pp. 832–844, Aug. 1998.
- [28] P. Sneath and R. Sokal, *Numerical Taxonomy: The Principles and Practice of Numerical Classification*. San Francisco, CA, USA: Freeman, 1973.
- [29] G. Giacinto and F. Roli, "Design of effective neural network ensembles for image classification purposes," *Image Vis. Comput.*, vol. 19, nos. 9–10, pp. 699–707, Aug. 2001.
- [30] J. L. Fleiss, B. Levin, and M. C. Paik, *Statistical Methods for Rates and Proportions*. Hoboken, NJ, USA: Wiley, 2013.
- [31] R. Kohavi and D. H. Wolpert, "Bias plus variance decomposition for zero-one loss functions," in *Proc. ICML*, vol. 96, 1996, pp. 275–283.
- [32] P. Cunningham and J. Carney, "Diversity versus quality in classification ensembles based on feature selection," in *Proc. ECML*. Berlin, Germany: Springer, 2000, pp. 109–116.
- [33] C. A. Shipp and L. I. Kuncheva, "Relationships between combination methods and measures of diversity in combining classifiers," *Inf. Fusion*, vol. 3, no. 2, pp. 135–148, 2002.
- [34] L. K. Hansen and P. Salamon, "Neural network ensembles," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 12, no. 10, pp. 993–1001, Oct. 1990.
- [35] D. Partridge and W. Krzanowski, "Software diversity: Practical statistics for its measurement and exploitation," *Inf. Softw. Technol.*, vol. 39, no. 10, pp. 707–717, 1997.
- [36] Y. Liu and X. Yao, "Ensemble learning via negative correlation," *Neural Netw.*, vol. 12, no. 10, pp. 1399–1404, Dec. 1999.
- [37] H. Chen and X. Yao, "Regularized negative correlation learning for neural network ensembles," *IEEE Trans. Neural Netw.*, vol. 20, no. 12, pp. 1962–1979, Dec. 2009.
- [38] H. Chen and X. Yao, "Multiobjective neural network ensembles based on regularized negative correlation learning," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 12, pp. 1738–1751, Dec. 2010.
- [39] G. Zenobi and P. Cunningham, "Using diversity in preparing ensembles of classifiers based on different feature subsets to minimize generalization error," in *Proc. ECML*. Berlin, Germany: Springer, 2001, pp. 576–587.
- [40] G. Brown, "An information theoretic perspective on multiple classifier systems," in *Proc. MCS*, 2009, pp. 344–353.
- [41] Z.-H. Zhou and N. Li, "Multi-information ensemble diversity," in *Proc. MCS*, 2010, pp. 134–144.
- [42] G. Tsoumakas, I. Partalas, and I. Vlahavas, "An ensemble pruning primer," in *Proc. SUEMA*. Berlin, Germany: Springer, 2009, pp. 1–13.
- [43] D. D. Margineantu and T. G. Dietterich, "Pruning adaptive boosting," in *Proc. Int. Conf. Mach. Learn. (ICML)*, vol. 97, 1997, pp. 211–218.
- [44] N. Li, Y. Yu, and Z.-H. Zhou, "Diversity regularized ensemble pruning," in *Proc. ECML PKDD*, 2012, pp. 330–345.
- [45] G. Martínez-Muñoz and A. Suárez, "Using boosting to prune bagging ensembles," *Pattern Recognit. Lett.*, vol. 28, no. 1, pp. 156–165, 2007.
- [46] G. Martínez-Muñoz and A. Suárez, "Pruning in ordered bagging ensembles," in *Proc. ICML*, 2006, pp. 609–616.
- [47] Z.-H. Zhou and W. Tang, "Selective ensemble of decision trees," in *Proc. RSFDGrC*, 2003, pp. 476–483.
- [48] I. Partalas, G. Tsoumakas, and I. Vlahavas, "A study on greedy algorithms for ensemble pruning," Dept. Inform., Aristotle Univ. Thessaloniki, Thessaloniki, Greece, Tech. Rep. TR-LPIS-360-12, 2012.
- [49] H. Guo, F. Sun, J. Cheng, Y. Li, and M. Xu, "A novel margin-based measure for directed hill climbing ensemble pruning," *Math. Problems Eng.*, vol. 2016, pp. 1–11, 2016, Art. no. 3845131, doi: 10.1155/2016/3845131.
- [50] C. Qian, Y. Yu, and Z.-H. Zhou, "Pareto ensemble pruning," in *Proc. AAAI*, 2015, pp. 2935–2941.
- [51] P. Indyk, S. Mahabadi, M. Mahdian, and V. S. Mirrokni, "Composable core-sets for diversity and coverage maximization," in *Proc. PODS*, 2014, pp. 100–108.
- [52] S. Aghamolaei, M. Farhadi, and H. Zarrabi-Zadeh, "Diversity maximization via composable coresets," in *Proc. CCGG*, 2015, p. 43.
- [53] S. Zadeh, M. Ghadiri, V. Mirrokni, and M. Zadimoghaddam, "Scalable feature selection via distributed diversity maximization," in *Proc. AAAI*, 2017, pp. 2876–2883.
- [54] P. Agarwal, S. Har-Peled, and K. Varadarajan, "Approximating extent measures of points," *J. ACM*, vol. 51, no. 4, pp. 606–635, 2004.
- [55] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 2012.
- [56] N. X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance," *J. Mach. Learn. Res.*, vol. 11, pp. 2837–2854, Jan. 2010.
- [57] N. Srinivas and K. Deb, "Multiobjective optimization using nondominated sorting in genetic algorithms," *J. Evol. Comput.*, vol. 2, no. 3, pp. 221–248, 1994.
- [58] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evol. Comput.*, vol. 6, no. 2, pp. 182–197, Apr. 2002.
- [59] B. Birnbaum and K. J. Goldman, "An improved analysis for a greedy remote-clique algorithm using factor-revealing LPs," *Algorithmica*, vol. 55, no. 1, pp. 42–59, 2009.
- [60] V. Mirrokni and M. Zadimoghaddam, "Randomized composable coresets for distributed submodular maximization," in *Proc. STOC*, 2015, pp. 153–162.
- [61] R. Graham, D. Knuth, and O. Patashnik, *Concrete Mathematics: A Foundation for Computer Science*. Reading, MA, USA: Addison-Wesley, 2012.
- [62] M. Lichman. (2013). *UCI Machine Learning Repository*. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [63] G. Martínez-Muñoz, D. Hernández-Lobato, and A. Suárez, "An analysis of ensemble pruning techniques based on ordered aggregation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 2, pp. 245–259, Feb. 2009.
- [64] J. Demšar, "Statistical comparisons of classifiers over multiple data sets," *J. Mach. Learn. Res.*, vol. 7, pp. 1–30, Jan. 2006.
- [65] Wikipedia. *Speedup*. Accessed: Jun. 14, 2019. [Online]. Available: <https://en.wikipedia.org/wiki/Speedup>