

Less Is More: A Comprehensive Framework for the Number of Components of Ensemble Classifiers

Hamed Bonab¹ and Fazli Can²

Abstract—The number of component classifiers chosen for an ensemble greatly impacts the prediction ability. In this paper, we use a geometric framework for *a priori* determining the ensemble size, which is applicable to most of the existing batch and online ensemble classifiers. There are only a limited number of studies on the ensemble size examining majority voting (MV) and weighted MV (WMV). Almost all of them are designed for batch-mode, hardly addressing online environments. Big data dimensions and resource limitations, in terms of time and memory, make the determination of ensemble size crucial, especially for online environments. For the MV aggregation rule, our framework proves that the more strong components we add to the ensemble, the more accurate predictions we can achieve. For the WMV aggregation rule, our framework proves the existence of an ideal number of components, which is equal to the number of class labels, with the premise that components are completely independent of each other and strong enough. While giving the exact definition for a strong and independent classifier in the context of an ensemble is a challenging task, our proposed geometric framework provides a theoretical explanation of diversity and its impact on the accuracy of predictions. We conduct a series of experimental evaluations to show the practical value of our theorems and existing challenges.

Index Terms—Data stream, ensemble cardinality, ensemble size, law of diminishing returns, majority voting (MV), supervised learning, voting framework, weighted MV (WMV).

I. INTRODUCTION

OVER the last few years, the design of learning systems for mining the data generated from the real-world problems has encountered new challenges such as the high dimensionality of big data, as well as growth in volume, variety, velocity, and veracity—the four V’s of big data.¹ In the context of data dimensions, the volume is the amount of data, variety is the number of types of data, velocity is the speed of data, and veracity is the uncertainty of data; generated in real-world applications and processed by the

learning algorithm. The dynamic information processing and incremental adaptation of learning systems to the temporal changes are among the most demanding tasks in the literature for a long time [2], [3].

Ensemble classifiers are among the most successful and well-known solutions to supervised learning problems, particularly for online environments [4]–[6]. The main idea is to construct a collection of individual classifiers, even with weak learners, and combine their votes. The aim is to build a stronger classifier, compared with each individual component classifier [7]. The training mechanism of components and the vote aggregation method mostly characterize an ensemble classifier [8].

There are two main categories of vote aggregation methods for combining the votes of component classifiers: weighting methods and metalearning methods [8], [9]. Weighting methods assign a combining weight to each component and aggregate their votes based on these weights [e.g., majority voting (MV), performance weighting, and Bayesian combination]. They are useful when the individual classifiers perform the same task and have comparable success. Metalearning methods refer to learning from the classifiers and from the classifications of these classifiers on training data (e.g., stacking, arbiter trees, and grading). They are best suited for situations where certain classifiers consistently misclassify or correctly classify certain instances [8]. In this paper, we study the ensembles with the weighting combination rule. Metalearning methods are out of the scope of this paper.

An important aspect of ensemble methods is to determine how many component classifiers should be included in the final ensemble, known as the ensemble size or ensemble cardinality [1], [8], [10]–[14]. The impact of ensemble size on *efficiency* in terms of time and memory and *predictive performance* make its determination an important problem [15], [16]. Efficiency is especially important for online environments. In this paper, we extend our geometric framework [1] for predetermining the ensemble size, applicable to both batch and online ensembles.

Furthermore, diversity among component classifiers is an influential factor for having an accurate ensemble [8], [17]–[19]. Liu *et al.* [20] empirically studied ensemble size on diversity. Hu [11] explained that component diversity leads to uncorrelated votes, which in turn improves predictive performance. However, to the best of our knowledge, there is no explanatory theory revealing how and why diversity among components contributes toward overall ensemble accuracy [21]. Our proposed geometric framework

Manuscript received September 9, 2017; revised July 25, 2018 and November 8, 2018; accepted December 2, 2018. This work was supported in part by the Bilkent Information Retrieval Group and in part by the Center for Intelligent Information Retrieval. This paper is an extended version of the work presented at the Conference on Information and Knowledge Management (CIKM), 2016 [1]. (Corresponding author: Hamed Bonab.)

H. Bonab is with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA 01003 USA (e-mail: bonab@cs.umass.edu).

F. Can is with the Bilkent Information Retrieval Group, Computer Engineering Department, Bilkent University, 06800 Ankara, Turkey (e-mail: canf@cs.bilkent.edu.tr).

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

Digital Object Identifier 10.1109/TNNLS.2018.2886341

¹<http://www.ibmbigdatahub.com/infographic/four-vs-big-data>

introduces a theoretical explanation for the understanding of diversity in the context of ensemble classifiers. The main contributions of this paper are the following.

- 1) We present a brief comprehensive review of existing approaches for determining the number of component classifiers of ensembles.
- 2) We provide a spatial modeling for ensembles and use the linear least squares (LSQs) solution [22] for optimizing the weights of components of an ensemble classifier, applicable to both online and batch ensembles.
- 3) We exploit the geometric framework for the first time in the literature, for *a priori* determining the number of component classifiers of an ensemble.
- 4) We explain the impact of diversity among component classifiers of an ensemble on the predictive performance, from a theoretical perspective and for the first time in the literature.
- 5) We conduct a series of experimental evaluations on more than 16 different real-world and synthetic data streams and show the practical value of our theorems and existing challenges.

II. RELATED WORKS

The importance of ensemble size is discussed in several studies. There are two categories of approaches in the literature for determining ensemble size. Several ensembles *a priori* determine the ensemble size with a fixed value (such as bagging), while others try to determine the best ensemble size dynamically by checking the impact of adding new components to the ensemble [8]. Zhou *et al.* [23] analyzed the relationship between an ensemble and its components and concluded that aggregating *many* of the components may be the better approach. Through an empirical study, Liu *et al.* [20] showed that a subset of the components of a larger ensemble can perform comparably to the full ensemble, in terms of accuracy and diversity. Ulaş *et al.* [24] discussed approaches for incrementally constructing a batch-mode ensemble using different criteria including accuracy, significant improvement, diversity, correlation, and the role of search direction.

This led to the idea in ensemble construction, that it is sometimes useful, to let the ensemble extend unlimitedly and then prune the ensemble in order to get a more effective ensemble [8], [25]–[29]. Ensemble selection methods are developed as pruning strategies for ensembles. However, with today’s data dimensions and resource constraints, the idea seems impractical. Since the number of data instances grows exponentially, especially in online environments, there is a potential problem of approaching an infinite number of components for an ensemble. As a result, determining an upper bound for the number of components with a reasonable resource consumption is essential. As mentioned in [30], the errors cannot be arbitrarily reduced by increasing the ensemble size indefinitely.

There are a limited number of studies for batch-mode ensembles. Latinne *et al.* [10] proposed a simple empirical procedure for limiting the number of classifiers based on the McNemar nonparametric test of significance. Similar approaches [31], [32], suggested a range of 10–20 base

classifiers for bagging, depending on the base classifier and data set.

Oshiro *et al.* [12] cast the idea that there is an ideal number of component classifiers within an ensemble. They defined the ideal number as the ensemble size where exploiting more base classifiers brings no significant performance gain, and only increases computational costs. They showed this by using the weighted average area under the ROC curve, and some data set density metrics. Fumera *et al.* [32], [33] applied an existing analytical framework for the analysis of linearly combined classifiers of bagging, using misclassification probability. Hernández-Lobato *et al.* [13] suggested a statistical algorithm for determining the size of an ensemble, by estimating the required number of classifiers for obtaining stable aggregated predictions, using MV.

Pietruczuk *et al.* [34], [35] recently studied the automatic adjustment of ensemble size for online environments. Their approach determines whether a new component should be added to the ensemble by using a probability framework and defining a confidence level. However, the diversity impact of component classifiers is not taken into account, and there is a possibility of approaching the infinite number of components without reaching the confidence level. The assumption that the error distribution is i.i.d cannot be guaranteed, especially with a higher ensemble size; this reduces the improvements due to each extra classifier [30].

III. GEOMETRIC FRAMEWORK

In this section, we propose a geometric framework for studying the theoretical side of ensemble classifiers based on [1]. We mainly focus on online ensembles since they are more specific models compared with batch-mode ensembles. The main difference is that online ensembles are trained and tested over the course of incoming data while batch-mode ensembles are trained and tested once. As a result, batch-mode ensembles are also applicable to our framework, with a simpler declaration. We use this geometric framework in [36] and [37] for aggregating votes and proposing a novel online ensemble for evolving data stream.

Suppose we have an ensemble of m component classifiers, $\xi = \{CS_1, CS_2, \dots, CS_m\}$. Due to resource limitations, we are only able to keep the n latest instances of an incoming data stream as an *instance window*, $I = \{I_1, I_2, \dots, I_n\}$, where I_n is the latest instance and all the true-class labels are available. We assume that our supervised learning problem has p class labels, $C = \{C_1, C_2, \dots, C_p\}$. For batch-mode ensembles, I can be considered as the whole set of training data. Table I presents the notation of symbols for our geometric framework.

Our framework uses a p -dimensional Euclidean space for modeling the components’ votes and true-class labels. For a given instance $I_i (1 \leq i \leq n)$, each component classifier $CS_j (1 \leq j \leq m)$ returns a *score-vector* as $s_{ij} = \langle S_{ij}^1, S_{ij}^2, \dots, S_{ij}^p \rangle$, where $\sum_{k=1}^p S_{ij}^k = 1$. Considering all the score-vectors in our p -dimensional space, the framework builds a polytope of votes, which we call the *score-polytope* of I_i . For the true-class label of I_i , we have $o_i = \langle O_i^1, O_i^2, \dots, O_i^p \rangle$ as the *ideal-point*. Note that o_i is




Fig. 1. Schema of the geometric framework (obtained from [1]). I_t with class label $y_t = C_1$ is fed to the ensemble. Each component classifier, CS_j , generates a score-vector, S_{tj} . These score-vectors construct a surface in the Euclidean space, called *score-polytope*.

TABLE I
SYMBOL NOTATIONS OF THE GEOMETRIC FRAMEWORK

Notation	Definition
$I = \{I_1, I_2, \dots, I_n\}$	Instance window, I_i ; ($1 \leq i \leq n$)
$\xi = \{CS_1, CS_2, \dots, CS_m\}$	Ensemble of m component classifiers, CS_j ; ($1 \leq j \leq m$ and $2 \leq m$)
$C = \{C_1, C_2, \dots, C_p\}$	Classification problem with p class labels, C_k ; ($1 \leq k \leq p$ and $2 \leq p$)
$s_{ij} = \langle S_{ij}^1, S_{ij}^2, \dots, S_{ij}^p \rangle$	Score-vector for I_i and CS_j
$o_i = \langle O_i^1, O_i^2, \dots, O_i^p \rangle$	Ideal-point for I_i , O_i^k ; ($1 \leq k \leq p$)
$a_i = \langle A_i^1, A_i^2, \dots, A_i^p \rangle$	Centroid-point for I_i , A_i^k ; ($1 \leq k \leq p$)
$w = \langle W_1, W_2, \dots, W_m \rangle$	Weight vector for ξ , W_j ; ($1 \leq j \leq m$)
$b_i = \langle B_i^1, B_i^2, \dots, B_i^p \rangle$	Weighted-centroid-point for I_i , B_i^k ; ($1 \leq k \leq p$)

a *one-hot* vector in this paper. However, there are other supervised problems this assumption is not true—e.g., multilabel classification [37]. Studying other variations of the supervised problem is out of the scope of this paper. A general schema of our geometric framework is presented in Fig. 1.

Example: Assume, we have a supervised problem with three class labels, $C = \{C_1, C_2, C_3\}$. For a given instance I_t , the true-class label is C_2 . The ideal-point would be $o_t = \langle 0, 1, 0 \rangle$.

One could presumably define many different algebraic rules for vote aggregation [38], [39]—minimum, maximum, sum, mean, product, median, and so on. While these vote aggregation rules can be expressed using our geometric framework, we study the MV and weighted MV (WMV) aggregation rules in this paper. In addition, individual vote scores can be aggregated based on two different voting schemes [40].

- 1) *Hard Voting:* The score-vector of a component classifier is first transformed into a one-hot vector, possibly using a hard-max function, and then combined.
- 2) *Soft Voting:* The score-vector is used for vote aggregation. We use soft voting for our framework.

The Euclidean norm is used as the loss function, $\text{loss}(\cdot, \cdot)$, for optimization purposes [22]. The Euclidean distance of any score-vector and ideal-point expresses the effectiveness of that component for the given instance. Using aggregation rules, we aim to define a mapping function from a score-polytope into a single vector and measure the effectiveness of our ensemble. Wu and Crestani [41] applied a similar geometric framework for data fusion of information retrieval systems. In this paper, some of our theorems are obtained and adapted to ensemble learning from their framework.

A. Majority Voting (MV)

The mapping of a given score-polytope into its centroid can be expressed as the MV aggregation—plurality voting or averaging. For a given instance, I_t , we have the following mapping to the *centroid-point*, $a_t = \langle A_t^1, A_t^2, \dots, A_t^p \rangle$

$$A_t^k = \frac{1}{m} \sum_{j=1}^m S_{tj}^k \quad (1 \leq k \leq p). \quad (1)$$

Theorem 1: For I_t , the loss between the centroid-point a_t and ideal-point o_t is not greater than the average loss between m score-vectors and o_t , that is to say

$$\text{loss}(a_t, o_t) \leq \frac{1}{m} \sum_{j=1}^m \text{loss}(s_{tj}, o_t). \quad (2)$$

Proof: Based on Minkowski's inequality for sums [42]

$$\sqrt[p]{\sum_{k=1}^p \left(\sum_{j=1}^m \theta_j^k \right)^2} \leq \sum_{j=1}^m \sqrt[p]{\sum_{k=1}^p (\theta_j^k)^2}.$$

Letting $\theta_j^k = S_{tj}^k - O_t^k$ and substituting results

$$\sqrt[p]{\sum_{k=1}^p \left(m \left(\frac{1}{m} \sum_{j=1}^m (S_{tj}^k - O_t^k) \right) \right)^2} \leq \sum_{j=1}^m \sqrt[p]{\sum_{k=1}^p (S_{tj}^k - O_t^k)^2}.$$

Since $m > 0$, we have the following

$$\sqrt[p]{\sum_{k=1}^p \left(\frac{1}{m} \sum_{j=1}^m S_{tj}^k - O_t^k \right)^2} \leq \frac{1}{m} \sum_{j=1}^m \sqrt[p]{\sum_{k=1}^p (S_{tj}^k - O_t^k)^2}.$$

Using (1) and loss definition, (2) can be achieved. \square

Discussion: Theorem 1 shows that the performance of an ensemble with the MV aggregation rule is at least equal to the average performance of all individual components.

Theorem 2: For I_t , let $\xi_l = \xi - \{CS_l\}$ ($1 \leq l \leq m$) be a subset of ensemble ξ without CS_l . Each ξ_l has a_{tl} as its centroid-point. We have

$$\text{loss}(a_t, o_t) \leq \frac{1}{m} \sum_{l=1}^m \text{loss}(a_{tl}, o_t). \quad (3)$$

Proof: a_t is the centroid-point of all a_{tl} points according to the definition. Assume ξ_l as an individual classifier

with score-vector of a_{tl} . Theorem 1 for every ζ_l results (3) directly. \square

Discussion: Theorem 2 can be generalized for any subset definition with $(m - f)$ components, $1 \leq f \leq (m - 2)$. This shows that an ensemble with m components performs better (or at least equal) compared with the average performance of ensembles with $m - f$ components. It can be concluded that better performance can be achieved if we aggregate more component classifiers. However, if we keep adding poor components to the ensemble, it can diminish overall prediction accuracy by increasing the upper bound in (2). This is in agreement with the result of the Bayes error reduction analysis [43], [44]. Setting a threshold, as expressed in [30], [34], [35], and [44], can give us the ideal number of components for a specific problem.

B. Weighted Majority Voting (WMV)

For this aggregation rule, a weight vector $w = \langle W_1, W_2, \dots, W_m \rangle$ for components of ensemble is defined, $W_j \geq 0$ and $\sum W_j = 1$ for $1 \leq j \leq m$. For a given instance, I_t , we have the following mapping to the *weighted-centroid-point*, $b_t = \langle B_t^1, B_t^2, \dots, B_t^p \rangle$

$$B_t^k = \sum_{j=1}^m W_j S_{ij}^k \quad (1 \leq k \leq p). \quad (4)$$

Note that giving equal weights to all the components will result in the MV aggregation rule. WMV presents a flexible aggregation rule. No matter how poor a component classifier is, with a proper weight vector we can cancel its effect on the aggregated results. However, as discussed earlier, this is not true for the MV rule. In the following, we give the formal definition of the optimum weight vector, which we aim to find.

Definition 1 (Optimum Weight Vector): For an ensemble, ζ , and a given instance, I_t , weight vector w_o with the weighted-centroid-point b_o is the optimum weight vector where for any w_x with weighted-centroid-point b_x the following is true; $\text{loss}(b_o, o_t) \leq \text{loss}(b_x, o_t)$.

Theorem 3: For a given instance, I_t , let the optimum weight vector, w_o , and the weighted-centroid-point b_t . The following must hold

$$\text{loss}(b_t, o_t) \leq \min\{\text{loss}(s_{t1}, o_t), \dots, \text{loss}(s_{tm}, o_t)\}. \quad (5)$$

Proof: Assume that the least loss belongs to component j , among m score-vectors. We have the following two cases.

1) *Maintaining the Performance:* Simply giving a weight of 1 to j 's component and 0 for the remaining components result in the equality case; $\text{loss}(b_t, o_t) = \text{loss}(s_{tj}, o_t)$.

2) *Improving the Performance:* Using a linear combination of j and other components with proper weights result in a weighted-centroid-point closer to the ideal point. We can always find such a weight vector in the Euclidean space if other components are not the exact same as j . \square

Using the squared Euclidean norm as the measure of closeness for the linear LSQ problem [22] results

$$\min_w \|o - wS\|_2^2 \quad (6)$$

where for each instance I_i in the instance window, $S \in \mathbb{R}^{m \times p}$ is the matrix with score-vectors s_{ij} in each row corresponding to the component classifier j , $w \in \mathbb{R}^m$ is the vector of weights to be determined, and $o \in \mathbb{R}^p$ is the vector of the ideal-point. We use the following function for our optimization solution

$$f(W_1, W_2, \dots, W_m) = \sum_{k=1}^p \left(\sum_{j=1}^m (W_j S_{ij}^k) - O_i^k \right)^2. \quad (7)$$

Taking a partial derivation over W_q ($1 \leq q \leq m$), setting the gradient to zero, $\nabla f = 0$, and finding optimum points give us the optimum weight vector. Letting the following summations as λ_{qj} and γ_q

$$\lambda_{qj} = \sum_{k=1}^p S_{iq}^k S_{ij}^k, \quad (1 \leq q, j \leq m) \quad (8)$$

$$\gamma_q = \sum_{k=1}^p O_i^k S_{iq}^k, \quad (1 \leq q \leq m) \quad (9)$$

lead to m linear equations with m variables (weights). Briefly, $\Lambda w = \gamma$, where $\Lambda \in \mathbb{R}^{m \times m}$ is the coefficients matrix and $\gamma \in \mathbb{R}^m$ is the remainders vector—using (8) and (9), respectively. The proper weights in the matrix equation are our intended optimum weight vector. Here, we only use a single instance of the instance window for simplicity of equations, however, a summation on all the instances of the instance window can give us the optimal weights. For a more detailed explanation, see [36].

Discussion: According to (8), Λ is a symmetric square matrix. If Λ has full rank, our problem has a unique solution. On the other hand, in the sense of the LSQ solution [22], it is probable that Λ is rank-deficient, and we may not have a unique solution. Studying the properties of this matrix lead us to the following theorem.

Theorem 4: If the number of component classifiers is not equal to the number of class labels, $m \neq p$, then the coefficient matrix would be rank-deficient, $\det \Lambda = 0$.

Proof: Since we have p dimensions in our Euclidean space, p independent score-vectors would be needed for the basis spanning set. Any number of vectors, m , more than p is dependent on the basis spanning set, and any number of vectors, m , less than p is insufficient for constructing the basis spanning set. \square

Discussion: The above-mentioned theorem excludes some cases in which we cannot find optimum weights for aggregating votes. There are several numerical solutions for solving rank-deficient LSQs problems (e.g., QR factorization and singular value decomposition), however, the resulting solution is relatively expensive, may not be unique, and optimality is not guaranteed. Theorem 4's outcome is that the number of independent components for an ensemble is crucial for providing a full-rank coefficient matrix, in the aim of an optimal weight vector solution.

C. Diversity Among Components

Theorem 4 shows that for weight vector optimality, $m = p$ should be true. However, the reverse cannot be guaranteed

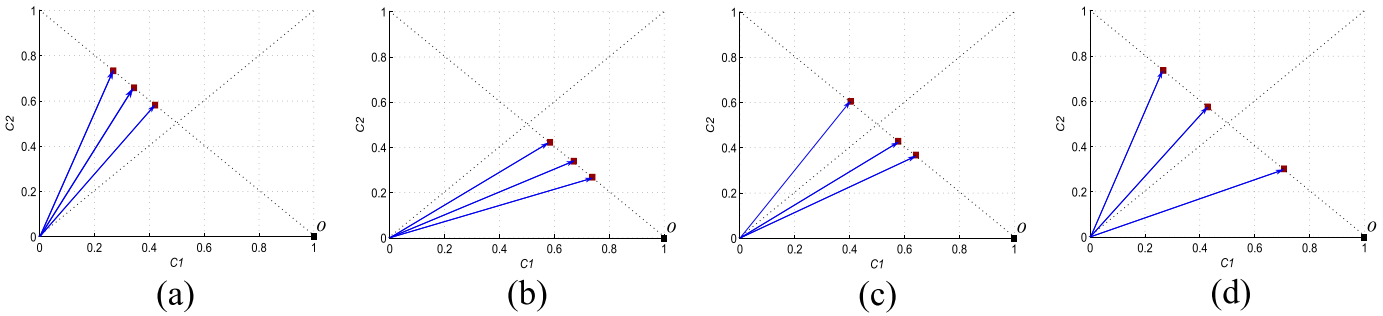


Fig. 2. Four score-vector possibilities of an ensemble with size three. The true-class label of the instance is $C1$ for a binary classification problem. If two of these score-vectors exactly match each other for several data instances, we cannot consider them to be independent and diverse enough components.

in general. Assuming $m = p$ and letting $\det \Lambda = 0$ for the parametric coefficient matrix results in some conditions where we have vote agreement, and cannot find a unique optimum weight vector. As an example, suppose we have two component classifiers for a binary classification task, $m = p = 2$. Letting $\det \Lambda = 0$, results the following equations; $S_{11}^1 + S_{12}^2 = 1$ or $S_{11}^2 + S_{12}^1 = 1$, meaning the agreement of component classifiers—i.e., the exact same vote vectors. More specifically, this suggests another condition for weight vector optimality: the importance of diversity among component classifiers.

Fig. 2 presents four mainly different score-vector possibilities for an ensemble with size three. The true-class label of the examined instance is $C1$ for a binary classification problem. All score-vectors are normalized and placed on the main diagonal of the spatial environment. The counter-diagonal line divides the decision boundary for the class label determination based on the probability values. If the component's score-vector is in the lower triangular, it is classified $C1$ and similarly, if it is in the left triangular part it is classified $C2$. Fig. 2(a) and (b) shows the misclassification and true classification situations, respectively. Fig. 2(c) and (d) shows disagreement among components of the ensemble.

If for several instances, in a sequence of data, the score-vectors of two components are equal (or act predictably similar), they are considered dependent components. There are several measurements for quantifying this dependence for ensemble classifiers (e.g., Q-statistic) [19]. However, most of the measurements in practice use the oracle output of components (i.e., only predicted class labels) [19]. Our geometric framework shows the potential importance of using score-vectors for diversity measurements. It is out of the scope of this paper to propose a diversity measurement using score-vectors and we leave it as a future work.

To the best of our knowledge, there is no explanatory theory in the literature revealing why and how diversity among components contribute toward overall ensemble accuracy [17], [21]. Our geometric modeling of ensemble's score-vectors and the optimum weight vector solution provide a theoretical insight for the commonly agreed upon idea that “the classifiers should be different from each other, otherwise the overall decision will not be better than the individual decisions” [19]. Optimum weights can be reached when we have the same number of independent and diverse component

classifiers as class labels. Diversity has a great impact on the coefficient matrix that consequently impacts the accurate predictions of an ensemble. For the case of MV, adding more dependent classifiers will dominate the decision of other components.

Discussion: Our geometric framework supports the idea that there is an ideal number of component classifiers for an ensemble, with which we can reach the most accurate results. Increasing or decreasing the number of classifiers from this ideal point may deteriorate predictions, or bring no gain to the overall performance of the ensemble. Having more components than the ideal number of classifiers can mislead the aggregation rule, especially for MV. On the other hand, having fewer is insufficient for constructing an ensemble which is stronger than the single classifier. We refer to this situation as “the law of diminishing returns in ensemble construction.”

Our framework suggests that the number of class labels of a data set as the ideal number of component classifiers, with the premise that they generate independent scores and aggregated with optimum weights. However, real-world data sets and existing ensemble classifiers do not guarantee this premise most of the time. Determining the exact value of this ideal point for a given ensemble classifier, over real-world data, is still a challenging problem due to the different complexities of data sets.

IV. EXPERIMENTAL EVALUATION

The experiments conducted in [1] showed that for ensembles trained with a specific data set, we have an ideal number of components in which having more will deteriorate or at least provide no benefit to our prediction ability. Our extensive experiments in [36] show the practical value of this geometric framework for aggregating votes, Geometrically Optimum and Online-Weighted Ensemble (GOOWE). An adaption of the framework for the multilabel classification task, called GOOWE-ML, is introduced in [37]. The theoretical complexity analysis of the optimal weight calculation is presented in [37].

Here, through a series of experiments, we first investigate the impact of the number of class labels and the number of component classifiers for MV and WMV using a synthetic data set generator. Then, we study the impact of miscellaneous data streams using several real-world and

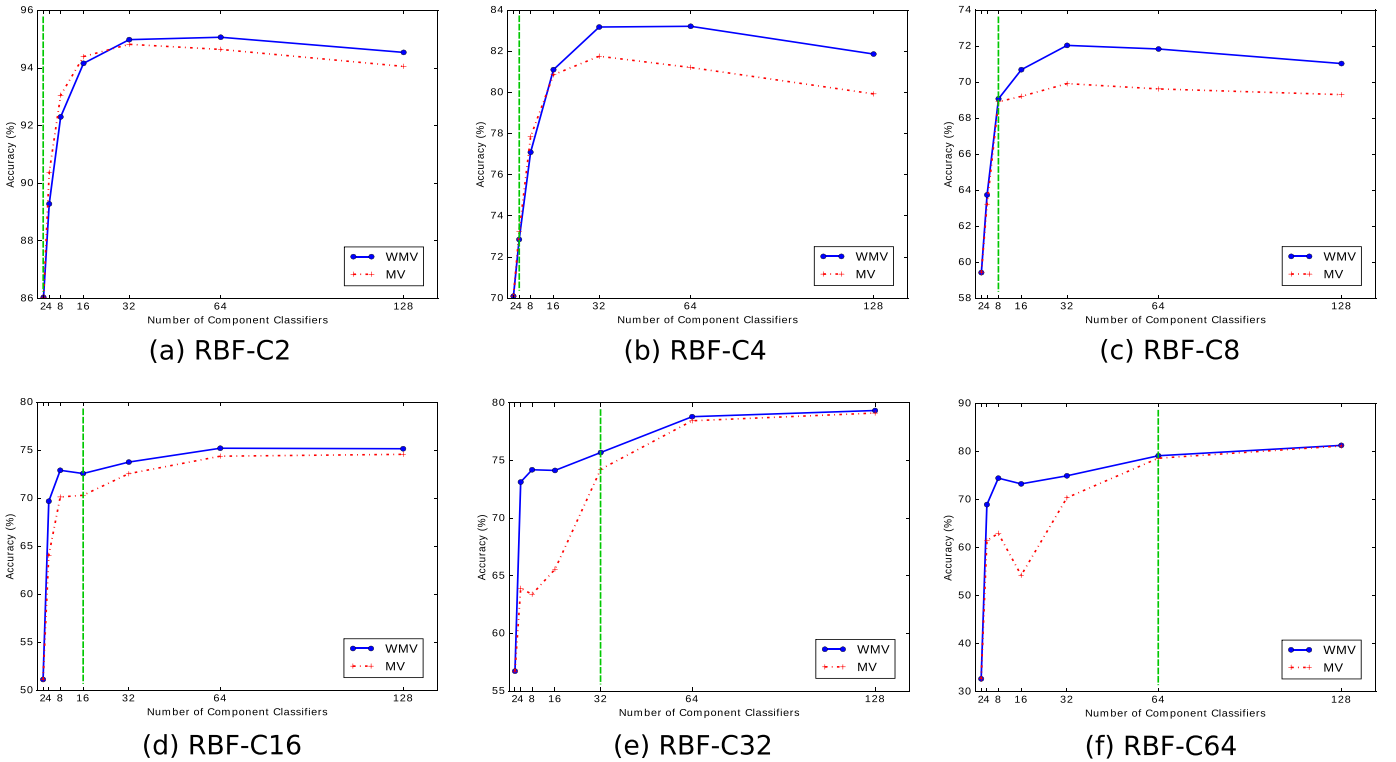


Fig. 3. Prediction behavior of WMV and MV aggregation rules, in terms of accuracy, for RBF-C data sets with increasing both the number of component classifiers, m , and the number of class labels, p . The equality case, $m = p$, is shown on each plot using a green dashed vertical line.

synthetic data sets. Finally, we explore the outcome of our theorems on the diversity of component classifiers and the practical value of our study. All the experiments are implemented using the Massive On-line Analysis (MOA) framework [45] and interleaved-test-then-train is used for accurate measurements. An instance window, with a length of 500 instances, is used for keeping the latest instances.

A. Impact of Number of Class Labels

1) *Setup*: To investigate the sole impact of the number of class labels of the data set, i.e., the p value, on the accuracy of an ensemble, we use the GOOWE [36] method. It uses our optimum weight vector calculation for vote aggregation using WMV. The Hoeffding Tree (HT) [46] is used as the component classifier, due to its high adaptivity to data stream classification. For a fair comparison, we modify GOOWE for having MV aggregation rule by simply providing equal weights to the components. These two variations of GOOWE, i.e., WMV and MV, are used for this experiment. Each of these variations trained and tested using different ensemble size values, starting from only two components and doubling at each step—i.e., our investigated ensemble sizes, m values, are 2, 4, ..., 128.

2) *Data Set*: Since existing real-world data sets are not consistent, in terms of classification complexity, we are only able to use synthetic data for this experiment in order to have reasonable comparisons. We choose the popular random radial basis function (RBF) generator since it is capable of generating data streams with an arbitrary number of features

and class labels [47]. Using this generator, implemented in the MOA framework [45], we prepare six data sets, each containing 1 million instances with 20 attributes, with the default parameter settings of the RBF generator. The only difference is the number of class labels among data sets which are 2, 4, 8, 16, 32, and 64. We reflect this in data set naming as RBF-C2, RBF-C4, ..., respectively.

3) *Results*: Fig. 3 presents prediction accuracy for WMV and MV with increasing component counts, m , on each data set. To mark the equality of m and p , we use a green dashed vertical line. We can make the following interesting observations.

- 1) A weighted aggregation rule becomes more vital with an increasing number of component classifiers.
- 2) WMV performs more resiliently in multiclass problems, compared with binary classification problems, when compared with MV. The gap between WMV and MV seems to increase with greater numbers of class labels.
- 3) There is a peak point in the accuracy value, and it is dependent on the number of class labels. This can be seen by comparing RBF-C2, RBF-C4, and RBF-C8 (the first row in Fig. 3) with RBF-C16, RBF-C32, and RBF-C64 (the second row in Fig. 3) plots. In the former set, we see that after a peak point, the accuracy starts to drop. However, in the latter set, we see that the peak points are with $m = 128$.
- 4) The theoretical vertical line, i.e., the equality case $m = p$, seems to precede the peak point on each plot. We suspect that this might be due to Theorem 4's

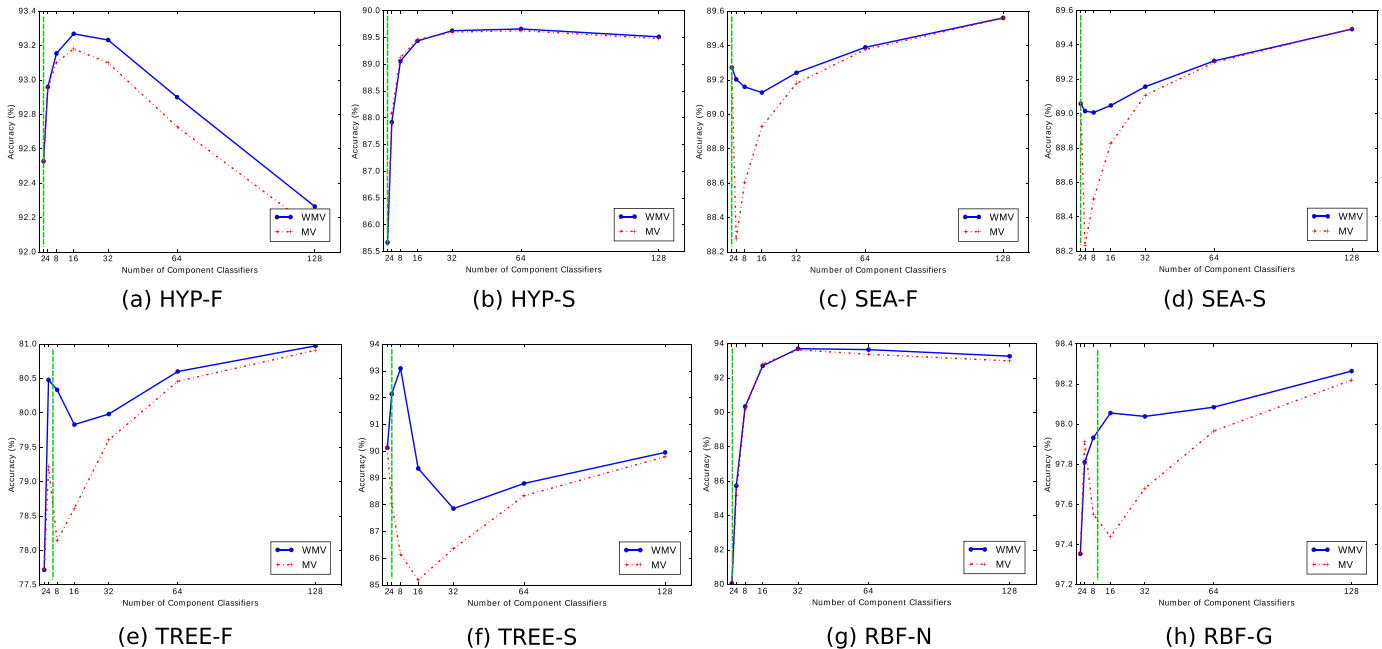


Fig. 4. Prediction behavior of WMV and MV aggregation rules, in terms of accuracy, for miscellaneous synthetic data sets, with increasing both the number of component classifiers, m , and class labels, p . The equality case, $m = p$, is marked on each plot using a green dashed vertical line.

premise conditions: generating independent scores and aggregating with optimum weights.

B. Impact of Data Streams

1) *Setup*: There are many factors when the complexity of classification problems are considered—concept drift, the number of features, and so on. To this end, we investigate the number of component classifiers for WMV and MV on a wide range of data sets. We use an experimental setup similar to the previous experiments on different synthetic and real-world data sets. We aim to investigate some general patterns in more realistic problems.

2) *Data Set*: We select eight synthetic and eight real-world benchmark data sets used for stream classification problems in the literature. A summary of our data sets is given in Table II. For this selection, we aim to have a mixture of different concept drift types, number of features, number of class labels, and noise percentages. Synthetic data sets are similar to the ones used for the GOOWE evaluation [36]. For real-world data sets, Sensor, PowerSupply, and HyperPlane data sets are taken from.² The remainder of real-world data sets are taken from.³ See [1], [36] for a detailed explanations of the data sets.

3) *Results*: Figs. 4 and 5 present the prediction accuracy difference for WMV and MV for increasing component classifier counts, m , on each data set. For marking the equality of m and p , we use a green dashed vertical line, similar to the previous experiments. As we can see, given more broad types of data sets, each with completely different complexities, it is difficult to conclude strict patterns. We have the following

TABLE II
SUMMARY OF DATA SET CHARACTERISTICS

Dataset	#Instance	#Att	#CL	%N	Drift Spec.
HYP-F	1×10^6	10	2	5	Incrm., DS=0.1
HYP-S	1×10^6	10	2	5	Incrm., DS=0.001
SEA-F	1×10^6	3	2	10	Abrupt, #D=9
SEA-S	1×10^6	3	2	10	Abrupt, #D=3
TREE-F	1×10^6	10	6	0	Reoc., #D=15
TREE-S	1×10^6	10	4	0	Reoc., #D=4
RBF-G	1×10^6	20	10	0	Gr., DS=0.01
RBF-N	1×10^6	20	2	0	No Drift
Airlines	539,383	7	2	-	Unknown
ClickPrediction	399,482	11	2	-	Unknown
Electricity	45,311	8	2	-	Unknown
HyperPlane	1×10^5	10	5	-	Unknown
CoverType	581,012	54	7	-	Unknown
PokerHand	1×10^7	10	10	-	Unknown
PowerSupply	29,925	2	24	-	Unknown
Sensor	2,219,802	5	58	-	Unknown

Note: #CL: No. of Class Labels, %N: Percentage of Noise, DS: Drift Speed, #D: No. of Drifts, Gr.: Gradual.

interesting observations: 1) for almost all the data sets, WMV, with optimum weights, outperforms MV; 2) we can see the same results as the previous experiments: there is a peak point in the accuracy value and it is dependent on the number of class labels; 3) the theoretical vertical line, i.e., the equality case $m = p$, seems to precede the peak point on each plot; and 4) optimum weighting seems to be more resilient in the evolving environments, i.e., data streams with concept drift, regardless of the type of concept drift.

The observations we have with the real-world data streams provide strong evidence that supports our claim which indicates that the number of class labels has an important influence on the ideal number of component classifiers and prediction

²Access URL: <http://www.cse.fau.edu/~xqzhu/stream.html>

³Access URL: <https://www.openml.org/>

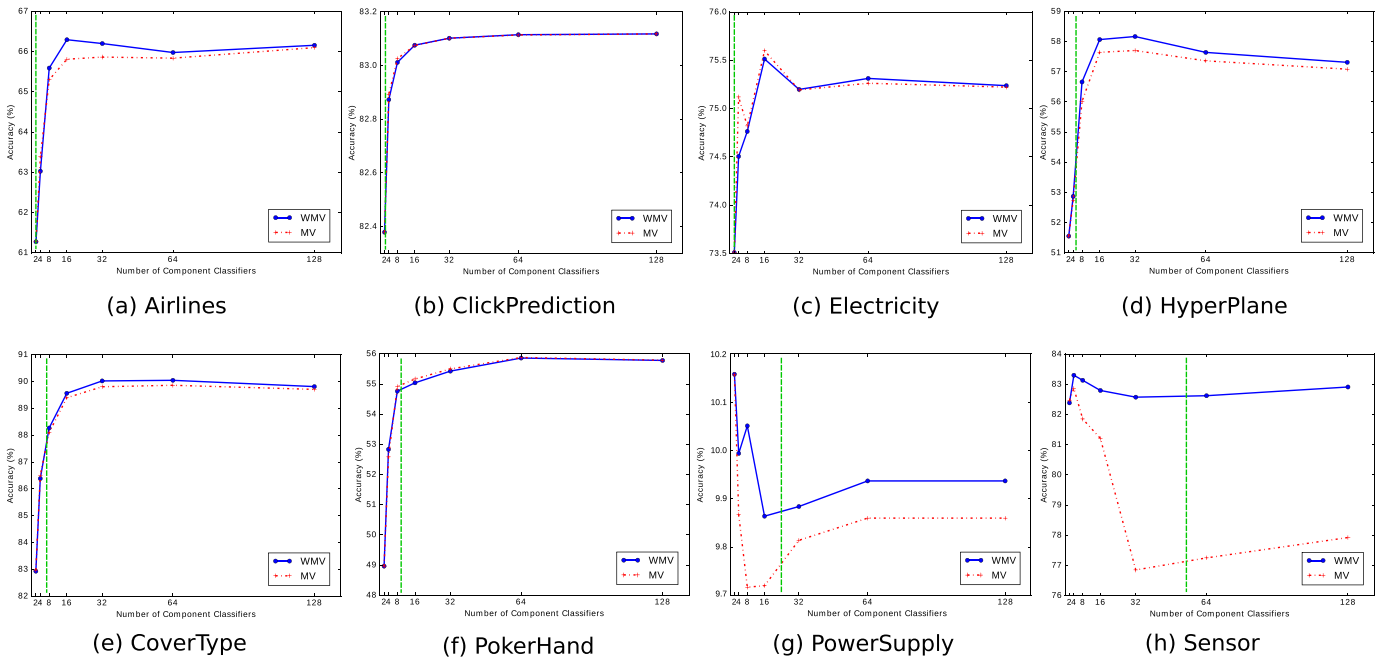


Fig. 5. Prediction behavior of WMV and MV aggregation rules, in terms of accuracy, for miscellaneous real-world data sets with increasing both the number of component classifiers, m , and class labels, p . The equality case, $m = p$, is marked on each plot using a green dashed vertical line.

performance. In Fig. 5, we observe that the peak performances, with one exception, are not observed with the maximum ensemble size. In other words, as we increase the number of component classifiers and move away from the green line and employ an ensemble of size 128, in all cases, prediction performance becomes lower than that of a smaller size ensemble. The only exception is observed with ClickPrediction; even with that one, no noticeable improvement is provided with the largest ensemble size. Furthermore, in all data streams, except ClickPrediction, the peak performances are closer to the green line rather than being closer to the largest ensemble size.

C. Impact of Diversity

1) *Setup*: In order to study the impact of diversity in ensemble classifiers and show the practical value of our theorems, we design two different scenarios for the binary classification problem. We select a binary classification for the purpose of this experiment since the difference between WMV and MV are almost always insignificant for binary classification, compared with multiclass problems. In addition, multiclass problems can potentially be modeled as several binary classification problems [48].

To this end, we recruit a well-known and state-of-the-art online ensemble classifier, called leverage bagging (LevBag) [49] as the base ensemble method for our comparisons. It is based on the OzaBagging ensemble [50], [51] and is proven to react well in online environments. It exploits resampling with replacement (i.e., input randomization), using a $Poisson(\lambda)$ distribution to train diversified component classifiers.

We use LevBag in our experiments since it initializes a fixed number of component classifiers—i.e., unlike GOOWE, where

component classifiers are dynamically added and removed during training in the course of incoming stream data [36], for LevBag the number of component classifiers are fixed from the initialization and the ensemble does not alter them. In addition, LevBag uses error-correcting output codes for handling multiclass problems and transforms them into several binary classification problems [49]. MV is used for vote aggregation, as the baseline of our experiments.

2) *Design*: For our analysis, we train different LevBag ensembles with 2, 4, ..., 64, and 128 components of classifiers—named LevBag-2, LevBag-4, ..., respectively. The HT [46] is used as the component classifier.

We design two experimental scenarios and compare them with LevBag ensembles as baselines. Each scenario is designed to show the practical value of our theorems with different perspectives. Here is a brief description.

- 1) *Scenario 1*: We select the two most diverse components out of a LevBag-10 ensemble's pool of component classifiers, called Sel2Div ensemble, and aggregate their votes. For pairwise diversity measurements among the components of the ensemble, Yule's Q-statistic [19] is used. Minku *et al.* [52] used it for pairwise diversity measurements of online ensemble learning. Q-statistic is measured between all the pairs, and the highest diverse pair is chosen. For two classifiers CS_r and CS_s , the Q-statistic is defined in the following. N^{ab} is the number of instances in the instance window that CS_r predicts a and CS_s predicts b

$$Q_{r,s} = \frac{N^{11}N^{00} - N^{01}N^{10}}{N^{11}N^{00} + N^{01}N^{10}}.$$

- 2) *Scenario 2*: We train a hybrid of two different algorithms as component classifiers, a potentially diverse ensemble.

TABLE III
CLASSIFICATION ACCURACY IN PERCENTAGE (%)—THE HIGHEST ACCURACY FOR EACH DATA SET IS BOLD

Dataset	LevBag							Select 2 most diverse	Hybrid of HT and NB
	2	4	8	16	32	64	128	2	2
Airlines	85.955	86.747	87.455	88.136	88.527	89.516	90.638	88.430	86.392
ClickPrediction	95.395	95.516	95.515	95.531	95.532	95.524	95.525	95.520	95.524
Electricity	83.481	84.299	84.172	84.842	84.332	84.797	84.906	85.406	82.538
RBF	78.391	79.210	79.750	79.789	80.065	80.321	80.339	78.575	77.483
SEA	86.011	86.354	86.070	86.246	85.387	84.976	84.872	86.166	84.650
HYP	87.620	87.957	88.839	88.388	88.292	88.176	88.313	87.957	88.283

TABLE IV
MULTIPLE COMPARISONS FOR FRIEDMAN STATISTICAL TEST RESULTS.
MINIMUM REQUIRED DIFFERENCE OF MEAN RANK IS **2.635**. HIGHER
MEAN RANK MEANS BETTER PERFORMANCE

Idx	Ensemble	Mean Rank	Different From ($P < 0.05$)
(1)	LevBag-2	2.000	(3)(4)(5)(6)(7)(8)
(2)	LevBag-4	4.250	(4)(7)
(3)	LevBag-8	4.833	(1)
(4)	LevBag-16	7.000	(1)(2)(9)
(5)	LevBag-32	6.333	(1)(9)
(6)	LevBag-64	5.750	(1)(9)
(7)	LevBag-128	7.000	(1)(2)(9)
(8)	Sel2Div	5.250	(1)(9)
(9)	Hyb-HTNB	2.583	(4)(5)(6)(7)(8)

For this, one instance of the HT and the naive Bayes (NB) [47] algorithms are exploited; both are trained on the same instances of data stream—without input randomization. We call this the Hyb-HTNB ensemble.

The ensemble sizes for both of these scenarios are two. For each instance, vote aggregation in both scenarios is done using our geometric weighting framework. An instance window of 100 latest incoming instances are kept, and using (8) and (9) weights are calculated— $\Lambda w = \gamma$.

3) *Data Set*: We examine our experiments using three real-world and three synthetic data streams, all with two-class labels. For real-world data sets, we use the exact same real-world data sets with two-class labels as with the previous experiments, see Table II. For synthetic data sets, we generate 500 000 instances of RBF, streaming ensemble algorithm (SEA), and HyperPlane generator (HYP) stream generator from the MOA framework [45]. For the settings of these generators, the default values are used, except for the number of class labels, which are two.

4) *Results*: Table III shows the prediction accuracy of different ensemble sizes and experimental scenarios for examined data sets. The highest accuracy for each data set is bold. We can see that the ensemble size and component selection have a crucial impact on the accuracy of prediction.

To differentiate the significance of differences in accuracy values, we exploited the nonparametric Friedman statistical test, with $\alpha = 0.05$ and $F(8, 40)$. The *null-hypothesis* for this statistical test claims that there is no statistically significant difference among all examined ensembles, in terms of accuracy. The resulting two-tailed probability value, $P = 0.002$, rejects the null-hypothesis and shows that the differences are significant.

The Friedman multiple pairwise comparisons are conducted and presented in Table IV. We observe that there is no significant difference among LevBag-8, LevBag-16, LevBag-32, LevBag-64, LevBag-128, and Sel2Div ensembles. Given that all are trained using the same component classifier, the impact of this result is important; only two base classifiers can be comparably good with 128 of them when they trained in a diverse enough fashion and weighted optimally.

On the other hand, the Hyb-HTNB ensemble performs equivalently as good as LevBag-2, LevBag-4, and LevBag-8, according to statistical significance tests. Hyb-HTNB is a naturally diverse ensemble; we included this in our experiment to show the impact of diversity on prediction accuracy. Since NB is a weak classifier compared with HT, it is reasonable that Hyb-HTNB is not performing as good as the Sel2Div ensemble.

V. CONCLUSION

In this paper, we studied the impact of ensemble size using a geometric framework. The entire decision-making process through voting is adapted to a spatial environment and weighting combination rules, including MV, are considered for providing better insight. The main focus of the study is online ensembles; however, nothing prevents us from using the proposed model on batch ensembles.

The ensemble size is crucial for online environments, due to the dimensionality growth of data. We discussed the effect of ensemble size with MV and optimal weighted voting aggregation rules. The highly important outcome is that we do not need to train a near-infinite number of components to have a good ensemble.

We delivered a framework which heightens the understanding of the diversity and explains why diversity contributes to the accuracy of predictions.

Our experimental evaluations showed the practical value of our theorems and highlighted existing challenges. Practical imperfections across different algorithms and different learning complexities on our various data sets prevent us to clearly show that $m = p$ is the ideal ensemble size and diversity are the core decisions to be used in the ensemble design. The experimental results show that the number of class labels has an important effect on the ensemble size. For example, in seven out of eight real-world data sets, the peak performances are closer to the ideal $m = p$ point rather than being closer to the largest ensemble size.

As a future work, we aim to define some diversity measures based on this framework, while also studying the coefficient matrix specifications. We also plan to study the ideal number of components for multilabel classification [37], and the use of an ensemble of ensembles in multistream environments.

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers A. Can and A. Büyükcakır for their valuable comments. Any opinions, findings, and conclusions expressed in this paper are those of the authors and do not necessarily reflect those of the sponsors.

REFERENCES

- [1] H. R. Bonab and F. Can, "A theoretical framework on the ideal number of classifiers for online ensembles in data streams," in *Proc. 25th Conf. Inf. Knowl. Manage. (CIKM)*, 2016, pp. 2053–2056.
- [2] F. Can, "Incremental clustering for dynamic information processing," *ACM Trans. Inf. Syst.*, vol. 11, no. 2, pp. 143–164, 1993.
- [3] Y. Freund and R. E. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *J. Comput. Syst. Sci.*, vol. 55, no. 1, pp. 119–139, Aug. 1997.
- [4] T. G. Dietterich, "Ensemble methods in machine learning," in *Multiple Classifier Systems*. Cham, Switzerland: Springer, 2000, pp. 1–15.
- [5] J. Gama, *Knowledge Discovery from Data Streams*. Boca Raton, FL, USA: CRC Press, 2010.
- [6] S. Wang, L. Minku, and X. Yao, "Resampling-based ensemble methods for online class imbalance learning," *IEEE Trans. Knowl. Data Eng.*, vol. 27, no. 5, pp. 1356–1368, May 2015.
- [7] B. Krawczyk, L. L. Minku, J. Gama, J. Stefanowski, and M. Woźniak, "Ensemble learning for data stream analysis: A survey," *Inf. Fusion*, vol. 37, pp. 132–156, Sep. 2017.
- [8] L. Rokach, "Ensemble-based classifiers," *Artif. Intell. Rev.*, vol. 33, nos. 1–2, pp. 1–39, 2010.
- [9] L.-W. Chan, "Weighted least square ensemble networks," in *Proc. IEEE Int. Joint Conf. Neural Netw. (IJCNN)*, vol. 2, Jul. 1999, pp. 1393–1396.
- [10] P. Latinne, O. Debeir, and C. Decaestecker, "Limiting the number of trees in random forests," in *Multiple Classifier System*. Berlin, Germany: Springer-Verlag, 2001, pp. 178–187.
- [11] X. Hu, "Using rough sets theory and database operations to construct a good ensemble of classifiers for data mining applications," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, Nov./Dec. 2001, pp. 233–240.
- [12] T. M. Oshiro, P. S. Perez, and J. A. Baranauskas, "How many trees in a random forest?" in *Machine Learning and Data Mining in Pattern Recognition*. Springer, 2012, pp. 154–168.
- [13] D. Hernández-Lobato, G. Martínez-Muñoz, and A. Suárez, "How large should ensembles of classifiers be?" *Pattern Recognit.*, vol. 46, no. 5, pp. 1323–1336, 2013.
- [14] H. M. Gomes, J. P. Barddal, F. Enembreck, and A. Bifet, "A survey on ensemble learning for data stream classification," *ACM Comput. Surv.*, vol. 50, no. 2, pp. 23:1–23:36, Jun. 2017.
- [15] G. Tsoumakas, I. Partalas, and I. Vlahavas, "A taxonomy and short review of ensemble selection," in *Proc. Workshop Supervised Unsupervised Ensemble Methods Appl.*, 2008, pp. 41–46.
- [16] J. B. Gomes, M. M. Gaber, P. A. C. Sousa, and E. Menasalvas, "Mining recurring concepts in a dynamic feature space," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 25, no. 1, pp. 95–110, Jan. 2014.
- [17] K. Jackowski, "New diversity measure for data stream classification ensembles," *Eng. Appl. Artif. Intell.*, vol. 74, pp. 23–34, Sep. 2018.
- [18] L. I. Kuncheva, *Combining Pattern Classifiers: Methods and Algorithms*. Hoboken, NJ, USA: Wiley, 2004.
- [19] 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.
- [20] H. Liu, A. Mandvikar, and J. Mody, "An empirical study of building compact ensembles," in *Advances in Web-Age Information Management*. Berlin, Germany: Springer-Verlag, 2004, pp. 622–627.
- [21] G. Brown, J. Wyatt, R. Harris, and X. Yao, "Diversity creation methods: A survey and categorisation," *Inf. Fusion*, vol. 6, no. 1, pp. 5–20, 2005.
- [22] P. C. Hansen, V. Pereyra, and G. Scherer, *Least Squares Data Fitting With Applications*. Baltimore, MD, USA: Johns Hopkins Univ. Press, 2013.
- [23] 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.
- [24] A. Ulaş, M. Semerci, O. T. Yıldız, and E. Alpaydın, "Incremental construction of classifier and discriminant ensembles," *Inf. Sci.*, vol. 179, no. 9, pp. 1298–1318, 2009.
- [25] L. Rokach, "Collective-agreement-based pruning of ensembles," *Comput. Statist. Data Anal.*, vol. 53, no. 4, pp. 1015–1026, 2009.
- [26] D. D. Margineantu and T. G. Dietterich, "Pruning adaptive boosting," in *Proc. Int. Conf. Mach. Learn. (ICML)*, vol. 97, 1997, pp. 211–218.
- [27] C. Toraman and F. Can, "Squeezing the ensemble pruning: Faster and more accurate categorization for news portals," in *Proc. Eur. Conf. Inf. Retr. (ECIR)*. Berlin, Germany: Springer-Verlag, 2012, pp. 508–511.
- [28] R. Elwell and R. Polikar, "Incremental learning of concept drift in nonstationary environments," *IEEE Trans. Neural Netw.*, vol. 22, no. 10, pp. 1517–1531, Oct. 2011.
- [29] T. Windeatt and C. Zor, "Ensemble pruning using spectral coefficients," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 24, no. 4, pp. 673–678, Apr. 2013.
- [30] K. Tumer and J. Ghosh, "Analysis of decision boundaries in linearly combined neural classifiers," *Pattern Recognit.*, vol. 29, no. 2, pp. 341–348, 1996.
- [31] E. Bauer and R. Kohavi, "An empirical comparison of voting classification algorithms: Bagging, boosting, and variants," *Mach. Learn.*, vol. 36, nos. 1–2, pp. 105–139, 1999.
- [32] G. Fumera, F. Roli, and A. Serrau, "A theoretical analysis of bagging as a linear combination of classifiers," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 7, pp. 1293–1299, Jul. 2008.
- [33] G. Fumera and F. Roli, "A theoretical and experimental analysis of linear combiners for multiple classifier systems," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 6, pp. 942–956, Jun. 2005.
- [34] L. Pietruczuk, L. Rutkowski, M. Jaworski, and P. Duda, "A method for automatic adjustment of ensemble size in stream data mining," in *Proc. IEEE Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2016, pp. 9–15.
- [35] L. Pietruczuk, L. Rutkowski, M. Jaworski, and P. Duda, "How to adjust an ensemble size in stream data mining?" *Inf. Sci.*, vol. 381, pp. 46–54, Mar. 2017.
- [36] H. R. Bonab and F. Can, "GOOWE: Geometrically optimum and online-weighted ensemble classifier for evolving data streams," *ACM Trans. Knowl. Discovery Data*, vol. 12, no. 2, pp. 25:1–25:33, Jan. 2018.
- [37] A. Büyükcakır, H. Bonab, and F. Can, "A novel online stacked ensemble for multi-label stream classification," in *Proc. 27th Int. Conf. Inf. Knowl. Manage. (CIKM)*, 2018, pp. 1063–1072.
- [38] J. Z. Kolter and M. A. Maloof, "Dynamic weighted majority: An ensemble method for drifting concepts," *J. Mach. Learn. Res.*, vol. 8, pp. 2755–2790, Dec. 2007.
- [39] J. Z. Kolter and M. A. Maloof, "Dynamic weighted majority: A new ensemble method for tracking concept drift," in *Proc. IEEE Int. Conf. Data Mining (ICDM)*, Nov. 2003, pp. 123–130.
- [40] D. J. Miller and L. Yan, "Critic-driven ensemble classification," *IEEE Trans. Signal Process.*, vol. 47, no. 10, pp. 2833–2844, Oct. 1999.
- [41] S. Wu and F. Crestani, "A geometric framework for data fusion in information retrieval," *Inf. Syst.*, vol. 50, pp. 20–35, Jun. 2015.
- [42] M. Abramowitz and I. A. Stegun, *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*, vol. 55. North Chelmsford, MA, USA: Courier Corp., 1964.
- [43] K. Tumer and J. Ghosh, "Error correlation and error reduction in ensemble classifiers," *Connection Sci.*, vol. 8, no. 3, pp. 385–404, 1996.
- [44] H. Wang, W. Fan, P. S. Yu, and J. Han, "Mining concept-drifting data streams using ensemble classifiers," in *Proc. Int. Conf. Knowl. Discovery Data Mining (SIGKDD)*, 2003, pp. 226–235.
- [45] A. Bifet, G. Holmes, R. Kirkby, and B. Pfahringer, "MOA: Massive online analysis," *J. Mach. Learn. Res.*, vol. 11, pp. 1601–1604, May 2010.
- [46] P. Domingos and G. Hulten, "Mining high-speed data streams," in *Proc. Int. Conf. Knowl. Discovery Data Mining (SIGKDD)*, 2000, pp. 71–80.
- [47] A. Bifet, G. Holmes, B. Pfahringer, R. Kirkby, and R. Gavaldà, "New ensemble methods for evolving data streams," in *Proc. Int. Conf. Knowl. Discovery Data Mining (SIGKDD)*, 2009, pp. 139–148.
- [48] R. Rifkin and A. Klautau, "In defense of one-vs-all classification," *J. Mach. Learn. Res.*, vol. 5, pp. 101–141, Dec. 2004.
- [49] A. Bifet, G. Holmes, and B. Pfahringer, "Leveraging bagging for evolving data streams," in *Proc. Int. Conf. Mach. Learn. Knowl. Discovery Databases (ECML-PKD)*, 2010, pp. 135–150.
- [50] N. C. Oza, "Online ensemble learning," Ph.D. dissertation, Comput. Sci. Division, Univ. California, Berkeley, CA, USA, Sep. 2001.

- [51] N. C. Oza and S. Russell, "Experimental comparisons of online and batch versions of bagging and boosting," in *Proc. Int. Conf. Knowl. Discovery Data Mining (SIGKDD)*, 2001, pp. 359–364.
- [52] L. L. Minku, A. P. White, and X. Yao, "The impact of diversity on online ensemble learning in the presence of concept drift," *IEEE Trans. Knowl. Data Eng.*, vol. 22, no. 5, pp. 730–742, May 2010.



Hamed Bonab received the B.S. degree in computer engineering from the Iran University of Science and Technology, Tehran, Iran, and the M.S. degree in computer engineering from Bilkent University, Ankara, Turkey. He is currently pursuing the Ph.D. degree with the College of Information and Computer Sciences, University of Massachusetts Amherst, Amherst, MA, USA.

His current research interests include stream processing, data mining, machine learning, and information retrieval.



Fazli Can received the B.S. and M.S. degrees in electrical and electronics and computer engineering and the Ph.D. degree in computer engineering from Middle East Technical University, Ankara, Turkey, in 1976, 1979, and 1985, respectively. He conducted his Ph.D. research under the supervision of Prof. E. Ozkarahan; at Arizona State University, Tempe, AZ, USA, and Intel, Chandler, AZ, USA; as a part of the RAP Database Machine Project.

He is currently a Faculty Member at Bilkent University, Ankara. Before joining Bilkent, he was a tenured Full Professor at Miami University, Oxford, OH, USA. He co-edited ACM SIGIR Forum from 1995 to 2002 and is a Co-Founder of the Bilkent Information Retrieval Group, Bilkent University. His interest in dynamic information processing dates back to his 1993 incremental clustering paper in *ACM Transactions on Information Systems* and some other earlier work with Prof. E. Ozkarahan on dynamic cluster maintenance. His current research interests include information retrieval and data mining.