

GOOWE: Geometrically Optimum and Online-Weighted Ensemble Classifier for Evolving Data Streams

Hamed R. Bonab, Bilkent University
Fazli Can, Bilkent University

Designing adaptive classifiers for an evolving data stream is a challenging task due to the data size and its dynamically changing nature. Combining individual classifiers in an online setting, the ensemble approach, is a well-known solution. It is possible that a subset of classifiers in the ensemble outperforms others in a time-varying fashion. However, optimum weight assignment for component classifiers is a problem which is not yet fully addressed in online evolving environments. We propose a novel data stream ensemble classifier, called Geometrically Optimum and Online-Weighted Ensemble (GOOWE), which assigns optimum weights to the component classifiers using a sliding window containing the most recent data instances. We map vote scores of individual classifiers and true class labels into a spatial environment. Based on the Euclidean distance between vote scores and ideal-points, and using the linear least squares (LSQ) solution, we present a novel, dynamic, and online weighting approach. While LSQ is used for batch mode ensemble classifiers, it is the first time that we adapt and use it for online environments by providing a spatial modeling of online ensembles. In order to show the robustness of the proposed algorithm, we use real-world datasets and synthetic data generators using the MOA libraries. First, we analyze the impact of our weighting system on prediction accuracy through two scenarios. Second, we compare GOOWE with 8 state-of-the-art ensemble classifiers in a comprehensive experimental environment. Our experiments show that GOOWE provides improved reactions to different types of concept drift compared to our baselines. The statistical tests indicate a significant improvement in accuracy, with conservative time and memory requirements.

CCS Concepts: •Information systems → Data stream mining; •Theory of computation → Online learning theory;

Additional Key Words and Phrases: Ensemble classifier, concept drift, evolving data stream, dynamic weighting, geometry of voting, least squares, spatial modeling for online ensembles

ACM Reference Format:

Hamed R. Bonab and Fazli Can, 2017. GOOWE: Geometrically Optimum and Online-Weighted Ensemble Classifier for Evolving Data Streams. *ACM Trans. Knowl. Discov. Data.* 9, 4, Article 39 (August 2017), 33 pages.

DOI: 0000001.0000001

1. INTRODUCTION

The automation of several processes in daily life has dramatically increased the number of data stream generators. Mining the data generated in real-world applications; like traffic management data, click streams in web exploration, detailed call logs, stock market and business transactions, social and computer network logs, and many other such examples; introduced several challenges to the domain. These challenges are

The paper is accepted for publication in the ACM Transactions on Knowledge Discovery from Data (TKDD) in August 2017.

The authors are with the Bilkent Information Retrieval Group, Computer Engineering Department, Bilkent University, 06800, Ankara, Turkey.

The current address of Hamed R. Bonab: College of Information and Computer Sciences, University of Massachusetts, Amherst, MA 01003, USA.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM. 1556-4681/2017/08-ART39 \$15.00

DOI: 0000001.0000001

mostly due to the size and time-evolving nature of these data streams. The cost and effort of storing and retrieving this type of data made the on-the-fly real-time analysis of incoming data crucial [Gama 2010].

In such dynamically evolving and non-stationary environments, data distribution can change over time, this is referred to as concept drift [Gama et al. 2014]. However, some of these changes are not real concept drifts, and they do not need to be reacted to by adaptive classifiers. Real concept drift is referred to as change in the conditional distribution of the output, given the input features, while the distribution of the input may stay unchanged [Gama 2010; Gama et al. 2014]. An example of evolving environments is filtering spam emails, in which the definition of the spam class label may change with time. Since users specify these class labels, and their preferences may also change with time, the conditional distribution of labels for incoming emails can change [Kuncheva 2004].

Designing a classifier for time-evolving data streams has some considerations to be addressed, compared to traditional classifiers. Since data arrives continuously, any proposed algorithm needs to process it under strict time constraints. Handling large volumes of data in main memory is impractical, so the proposed algorithm must use limited memory. Patterns of change in target concepts are categorized into sudden/abrupt, incremental, gradual, and reoccurring drifts [Bifet et al. 2009; Gama et al. 2014; Kuncheva 2008; Gomes et al. 2017; Krawczyk et al. 2017]. Effective classifiers should be able to handle these concept drifts.

More recently, many drift-aware adaptive learning algorithms have been developed. Among these algorithms, ensemble methods are naturally more consistent with the needs of the problem, and they are proven to outperform single algorithms statistically and computationally [Bifet et al. 2009; Brzezinski and Stefanowski 2014b; Kolter and Maloof 2005; Kuncheva 2004; Wang et al. 2003; Gomes et al. 2017; Krawczyk et al. 2017]. It is possible that a subset of classifiers in the ensemble outperforms others in a time-varying fashion. However, optimum weight assignment for component classifiers is a problem which is not yet fully addressed in online evolving environments [Zhu et al. 2010]. We propose a novel data stream ensemble classifier which assigns optimum weights to the component classifiers using a sliding window containing the most recent data instances. Since ensemble methods use individual classifiers inside their models, this does not decrease the importance of designing more adaptive individual classifiers for evolving data streams. Improving the performance of individual classifiers in terms of accuracy and resource usage can also increase the performance of an ensemble.

In this article, we concentrate on designing a geometric framework for dynamic weighting of component classifiers for ensemble methods. We model our ensemble in a spatial environment and use the Euclidean distance as our measure of closeness. We try to find an optimum weighting function based on LSQ, leading to a system of linear equations which describes the ensemble more precisely. Based on this system of linear equations, we design our algorithm called Geometrically Optimum and Online-Weighted Ensemble (GOOWE)—pronounced gooe (/ˈɡü-ē/). It is inspired from the geometry of voting, which is a well-known domain in the political and social sciences, and economics. The geometric analysis of individual votes for aggregation is proven to outperform existing solutions in these fields. In aggregation, various rules may have conflicting votes, i.e., “the paradox of voting.” Finding classes of profiles, uncovering paradoxes, and determining the likelihood of disagreements are among the problems addressed by the geometry of voting [Saari 2008].

For evaluating the performance of an algorithm in a time-evolving data stream domain, it is necessary to use tens of millions of examples [Bifet et al. 2009]. However, gathering this much real-world data, especially with substantial concept drifts, is not

Table I. Symbol Notation

Notation	Definition
S	Data stream
$I = \{I_1, I_2, \dots, I_n\}$	Instance window, $I_i; (1 \leq i \leq n)$
$I_t = x_t \in S$	Incoming data instance in time t
y_t / \hat{y}_t	Vector of true/predicted class label
$C = \{C_1, C_2, \dots, C_p\}$	Set of p class labels, $C_k; (1 \leq k \leq p)$
$\xi = \{CS_1, CS_2, \dots, CS_m\}$	Ensemble of m individual classifiers, $CS_j; (1 \leq j \leq m)$
$s_{ij} = \langle S_{ij}^1, S_{ij}^2, \dots, S_{ij}^p \rangle$	Score vector for I_i and CS_j , $S_{ij}^k; (1 \leq k \leq p)$
$o_i = \langle O_i^1, O_i^2, \dots, O_i^p \rangle$	Ideal-point for I_i , $O_i^k; (1 \leq k \leq p)$
$w = \langle W_1, W_2, \dots, W_m \rangle$	Weight vector for ξ , $W_j; (1 \leq j \leq m)$

feasible. There is a shortage in trusted evolving real-world publicly available datasets for testing stream classifiers [Krawczyk et al. 2017]. Moreover, we cannot verify concept drift phases in the course of time for real-world data streams. Some popular real-world data streams, used in the literature, questionably represent sufficiently real concept drifts (e.g. discussions on electricity data [Zliobaite 2013]). Because of these problems, like earlier studies in the literature, we use a combination of real-world and synthetic data streams in our experiments.

We experimentally evaluate our algorithm using several real-world and synthetic datasets representing gradual, incremental, sudden/abrupt, and reoccurring concept drifts. We use the most popular real-world datasets, and for generating synthetic data streams, we use the MOA libraries [Bifet et al. 2009]. For the sake of comparison, we use 8 state-of-the-art ensemble methods as baselines in our experiments. We follow the tradition and use classification accuracy, processing time, and memory costs as our comparison measurements. For classification accuracy measurement, we use the Interleaved Test-Then-Train approach [Bifet et al. 2009].

Contributions of our study. The summary of main contributions of this study are the following. We

- Provide a spatial modeling for online ensembles and use the linear least squares (LSQ) solution [Hansen et al. 2013] for optimizing the weights of components of an ensemble classifier for evolving environments. While LSQ is used for batch mode component weighting [Chan 1999; Friedman 2002], for the first time in the literature, we adapt and use it for online environments, as a stacking algorithm,
- Introduce an ensemble algorithm, called GOOWE. We use data chunks for training, and a sliding instance window containing the latest available data for testing; such an approach provides more robust behavior, as shown by our experiments,
- Analyze the impact of GOOWE's weighting system on component weighting strategy and ensemble model management strategy,
- Conduct an extensive experimental evaluation on 16 synthetic and 4 real-world data streams for comparing GOOWE with 8 state-of-the-art ensemble classifiers, and
- Carry out comprehensive statistical tests to show that GOOWE provides a statistically significant improvement in terms of accuracy while using conservative resources.

We present a brief chronological survey of the related work in Section 2; GOOWE in Section 3; our experimental evaluation setup in Section 4; experimental analysis in Section 5; comparative evaluation in Section 6; and statistical tests in Section 7. Section 8 offers a conclusion and directions for future research. Table I presents the notation of symbols that we use in the succeeding sections.

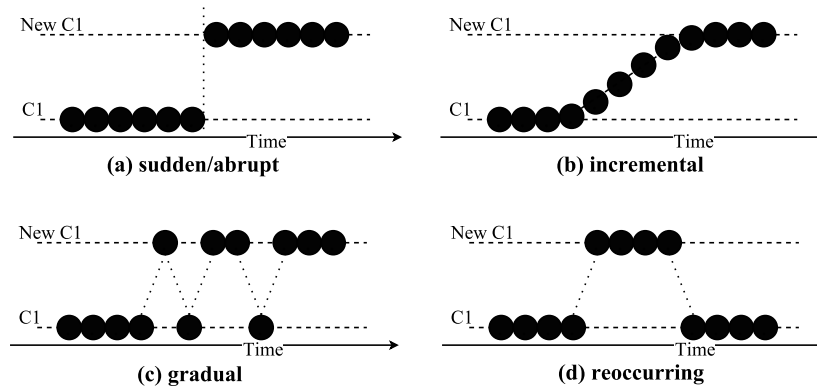


Fig. 1. Four patterns of real concept drift over time (revised from [Gama et al. 2014]).

2. BACKGROUND AND RELATED WORK

In this section, we explain our assumptions and specifications for time-evolving data streams. We distinguish different types of concept drifts based on the literature. We discuss different approaches of adapting concept drifts in evolving environments, focusing on ensemble methods, since they are naturally more capable of handling concept drift and they proved to outperform individual classifiers [Bifet et al. 2009; Gama et al. 2014; Gomes et al. 2017; Krawczyk et al. 2017].

2.1. Basic Concepts and Notations

The traditional supervised classification problem aims to map a vector of attributes, x , into a vector of class labels, y' , i.e. $x \mapsto y'$. The domain of attribute values in x , can be either numerical or nominal. However, for the domain of class labels in y' , we assume binary values for each label indicating selection or not-selection of that specific class label. We compare mapped class label vectors, y' , with true class label vectors, y . Instances from our data stream, $I_t = x_t \in S$, appear sequentially in temporal order, and we must process the data in an online fashion. We map x_t into y'_t , and when the true class labels, y_t , are available, we can evaluate our predictions. Due to the size of stream data, we are only able to store a limited number of instances in a window to process, and we need to discard old instances. Based on the availability of true class labels (data constraints) and our resources (solution/resource constraints), we can determine the length of the window. Classifiers are supposed to use limited memory and limited processing time per instance [Bifet et al. 2009; Gama et al. 2014; Kuncheva 2004].

In dynamically evolving environments, the conditional distribution of the output (i.e. true class labels) given the input vector, may change with time, i.e. $P(y_{t+1}|x_{t+1}) \neq P(y_t|x_t)$, while the distribution of the input vector itself, $P(x_t)$, may remain the same [Gama et al. 2014]. This is referred to as real concept drift and has raised several challenges for detecting and reacting to these changes.

Zhang et al. [Zhang et al. 2008] categorized real concept drifts into two scenarios; *Loose Concept Drift (LCD)* where only a change in $P(y_t|x_t)$ causes the concept drift, and *Rigorous Concept Drift (RCD)*, where change in both $P(y_t|x_t)$ and $P(x_t)$ cause the concept drift. The general assumption in the concept drift setting is that the change happens unexpectedly and is unpredictable. We do not consider the situation for some real-world problems where the change is predictable. We do not address concept-evolution, the arrival of a novel class label, and time-constrained classification [Farid et al. 2013;

Han et al. 2015; Masud et al. 2011; Sun et al. 2016; Wang et al. 2015; Zamani et al. 2016]. The reader is referred to [Gama et al. 2014] for various settings of the problem. We assume the most general setting of the evolving data stream classification problem.

There are several forms of change patterns over time for real concept drift, as shown in Fig. 1. If we consider a non-changing conditional distribution of the output given the input as one concept, a drift may happen *suddenly/abruptly* by replacing one concept with another (e.g. C1 with New C1 in Fig. 1-(a)) at a moment in time t . Drift may happen *incrementally* between the first and last concepts (e.g. C1 and New C1 in Fig. 1-(b), respectively), where there are many intermediate concepts which smoothly connect the dots. *Gradual drift* happens when there are no intermediate concepts and both of the first and last concepts are occurring for a period of time, Fig. 1-(c). Drifts may introduce new concepts that were not seen before, or previously seen concepts may *reoccur* after some time, Fig. 1-(d). Once-off random anomalies or blips are called *outlier/noise* and there should not be any reaction, as we do not consider them to be concept drift. Since most of the real-world problems are complex mixtures of these concept drifts, we expect any classifier to react and adapt reasonably to different types of concept drifts and remain robust to outliers, predicting with acceptable resource requirements [Gama et al. 2014].

2.2. Ensemble Classifiers for Evolving Online Environments

A recently published survey on concept drift adaption [Gama et al. 2014], presents a new taxonomy of adaptive classifiers using four existing modules of various learning methods in time-evolving environments. They are *memory management*, *change detection*, *learning property*, and *loss estimation*. In this study, we concentrate on model management strategies, as a learning property, to present state-of-the-art ensemble methods in chronological order. Model management strategies are techniques used in maintaining ensemble components as new data become available in the course of time. In addition, since we provide a novel stacking algorithm for online ensemble classifiers, we cover vote combination techniques of these ensembles. The remaining modules, other than learning property, are out of the scope of this paper.

Two more recently published surveys on ensemble learning for data stream analysis [Gomes et al. 2017; Krawczyk et al. 2017] show the importance of ensemble learning methods, especially on changing environments, and present ongoing research challenges. Gomes et al. cover existing data stream ensemble learning methods, propose a consistent taxonomy among them, and compare them based on some important aspects like vote aggregation, diversity measurement, and dynamic updates [Gomes et al. 2017]. Krawczyk et al. discuss more advanced topics such as imbalanced data streams, novelty detection, active and semi-supervised learning, complex data representations, and structured outputs with a focus on ensemble learning [Krawczyk et al. 2017].

Based on the model management categories of Kuncheva [Kuncheva 2004], there are five possible strategies for adaptive online classifiers:

- (1) *Horse Racing*: The dynamic combination ensemble strategy that aims to have the most proper combination rule of existing individual components in an ensemble;
- (2) *Updated Data Feeding*: Feeding individual classifiers with the most recent available data;
- (3) *Scheduled Feeding of Ensemble Members*: Scheduling the update of individual classifiers, either by retraining in a batch mode, or incrementally in an online mode with newly available data;
- (4) *Add/Drop Classifiers*: Adding fresh classifiers to the ensemble or pruning the deteriorating classifiers; and

- (5) *Feature Regulation*: Regulating the importance of features along with the life of an ensemble.

Practically any combination of these strategies can be used together and they do not need to be necessarily mutually exclusive.

Elwell and Polikar [Elwell and Polikar 2011] explain *active* versus *passive* approaches. Active approaches benefit from a drift detection mechanism, reacting only when drift is detected. On the other hand, passive approaches continuously update the model with each incoming data. Since training identical hypotheses with the same data produces identical classifiers, we need some mechanisms to increase their diversity. This is accomplished mostly by Kuncheva's third and fourth strategies. In addition, there are some works to measure and maintain the diversity of component classifiers [Minku et al. 2010; Minku and Yao 2012].

The WINNOW [Littlestone 1987], Weighted Majority (WM) [Littlestone and Warmuth 1994], and Hedge(β) [Freund and Schapire 1997] algorithms are the initial adaptive ensemble methods for large-scale changing environments. They mainly use the horse racing strategy for developing better combination rules in an off-line setting. They begin by creating a set of classifiers with an initial weight (usually 1). They adapt the ensemble's behavior using a reward-punishment system to keep track of the most trustworthy expert in each time slot. In particular, WINNOW uses $\alpha > 1$ (usually $\alpha = 2$) for its promotion ($w_i \leftarrow w_i \times \alpha$) and demotion ($w_i \leftarrow w_i \div \alpha$) steps. WM excludes the promotion step, and if an expert incorrectly classifies the instance, the algorithm decreases its weight by a multiplicative constant, $\beta \in [0, 1]$. The Hedge(β) algorithm operates in the same way, but instead of taking the weighted majority vote, chooses one classifier's decision as the ensemble decision. They provide a general framework for weighting component classifiers. However, they do not suggest any mechanism for dynamically adding or removing new components.

The Streaming Ensemble Algorithm (SEA) [Street and Kim 2001] provides a block-based and fixed-size ensemble of classifiers, each trained on the incoming chunk of instances—addressing Kuncheva's fourth model management strategy. If the ensemble has space, SEA adds the new classifier to the ensemble; otherwise, it puts the new classifier in the place of a weaker classifier. SEA uses a majority vote for predictions in an off-line setting. Due to batch mode component classifiers stopping learning after being formed, replacing the worst performing classifier in an unweighted ensemble, the learner is unable to properly track concept drifts in the stream data.

Oza [Oza 2001; Oza and Russell 2001] uses Kuncheva's second and third model management strategies together with the traditional bagging and boosting algorithms in online settings for designing OzaBagging and OzaBoosting. For stream data environments, as the number of training examples and component classifiers tend to go to infinity, Oza uses the Poisson distribution with $\lambda = 1$ for approximating the binomial distribution of sampling. A similar idea is used for the OzaBoosting algorithm. It employs incremental values of λ , starting from 1, for the training and sampling of classifiers.

The Dynamic Weighted Majority (DWM) [Kolter and Maloof 2003; 2007] introduced an ensemble of incremental learning algorithms, each with an associated weight in an online setting. Models are generated by the same learning algorithm on different batches of data. DWM uses the WM approach for assigning weights and makes predictions using a weighted-majority vote of the components where weights are dynamically changing. Pruning components with weights less than a threshold helps to avoid creating an excessive number of components. An extension to DWM, additive expert ensemble (AddExp) [Kolter and Maloof 2005], provides a general theoretical expert analysis to prove mistakes and loss bounds for a discrete and a continuous ensemble.

Table II. Summary of Related Ensemble Classifiers for Evolving Online Environments

Ensemble	Spec.		Kuncheva's Strategies				
	Study	Type	St. 1	St. 2	St. 3	St. 4	St. 5
WINNOW	[Littlestone 1987]	Passive	✓	×	✓	×	×
WM	[Littlestone and Warmuth 1994]	Passive	✓	×	✓	×	×
Hedge(β)	[Freund and Schapire 1997]	Passive	✓	×	✓	×	×
SEA	[Street and Kim 2001]	Passive	×	×	✓	✓	×
OzaBag/OzaBoost	[Oza 2001; Oza and Russell 2001]	Passive	×	✓	✓	×	×
DWM	[Kolter and Maloof 2003; 2007]	Passive	✓	×	✓	✓	×
AWE	[Wang et al. 2003]	Passive	✓	×	✓	✓	×
ACE	[Nishida et al. 2005]	Active	✓	✓	✓	×	×
LevBag	[Bifet et al. 2010a]	Active	✓	✓	✓	✓	×
Learn++.NSE	[Elwell and Polikar 2011]	Passive	✓	×	✓	✓	×
AUE2	[Brzezinski and Stefanowski 2014b]	Passive	✓	×	✓	✓	×
OAUE	[Brzezinski and Stefanowski 2014a]	Passive	✓	✓	✓	✓	×
GOOWE	Current work	Passive	✓	×	✓	✓	×

The Accuracy Weighted Ensemble (AWE) [Wang et al. 2003] alternatively suggests a general framework for mining changing data streams using weighted ensemble classifiers by re-evaluating ensemble components with incoming data chunks. Inspired by the framework of SEA, a new static learning algorithm is trained and the previous components of ensemble are evaluated on each incoming data chunk. However, these evaluations are done with a special version of Mean Square Error (MSE) allowing the algorithm to select the k best classifiers to create a new ensemble ($MSE_i = \frac{1}{|D|} \sum_{x \in D} (1 - M_c^i(x))^2$; where D is the latest data chunk and $M_c^i(x)$ is the probability score that x belongs to its true class label c , generated by a specific classifier system indexed i). Briefly, it assigns weights to component classifiers based on their expected classification accuracy—according to Bayes error optimization [Tumer and Ghosh 1996]. Moreover, the structure of the ensemble is pruned if errors of individual classifiers are worse than the MSE of a random classifier ($MSE_r = \sum_c P(c) \times (1 - P(c))^2$; where $P(c)$ is the probability of observing class label c). All in all, the weight of classifier i is determined by a linear function ($w_i = MSE_r - MSE_i$).

Since larger data chunks can provide a better distribution of data, they are more capable of building more accurate classifiers but may contain more than one change. Smaller chunks can separate drifting places better, but usually lead to poorer classifiers. In particular, ensembles built on large data chunks may react too slowly to sudden drifts occurring inside the chunk [Bifet et al. 2009; Brzezinski and Stefanowski 2014b]. To overcome this problem, Adaptive Classifier Ensemble (ACE) [Nishida et al. 2005], proposed an algorithm which uses a hybrid of one online classifier and a collection of batch classifiers (a mixture of active and passive approaches) along with a drift detection mechanism. ACE does not benefit from pruning strategies, and the possible use of a drift detector leads to poor reactions for gradual drifts.

Bifet [Bifet et al. 2010a] introduced Leverage Bagging (LevBag) as an extended version of OzaBagging, using the first four strategies of Kuncheva. It aims to increase the resampling rate using a larger value of λ in the Poisson distribution. Additionally, it adapts output detection codes [Dietterich and Bakiri 1995] for handling multi-class problems using only binary classifiers and the ADWIN [Bifet and Gavalda 2007] change detector for dealing better with concept drifts in stream data.

Learn++.NSE (NSE) [Elwell and Polikar 2011] is a batch learning ensemble that uses weighted majority voting. It updates weights dynamically with respect to the

time-adjusted errors of the classifiers on current and past environments. Similar to the AWE model management approach, evaluation of classifiers is considered by giving more credit to the ones capable of identifying previously unknown instances. On the other hand, classifiers that misclassify previously known instances are penalized. Moreover, NSE does not discard any component from the ensemble when its knowledge is not relevant to the current chunk of data. Although temporarily forgetting model management is particularly useful in cyclical environments, it causes some resource overuse. Ditzler and Polikar extended NSE for class imbalanced data stream [Ditzler and Polikar 2013].

Brzezinski et al. [Brzezinski and Stefanowski 2014b] proposed the Accuracy Updated Ensemble (AUE2), for combining the chunk-based algorithms with incremental learning components. Its model management strategy is based on AWE, and suggests a non-linear weighting function using the same MSE functions ($w_{ij} = \frac{1}{(MSE_r + MSE_{ij} + \epsilon)}$). The online version of AUE2 [Brzezinski and Stefanowski 2014a], called Online Accuracy Updated Ensemble (OAUE), uses a sliding window for the last n instances of the data stream.

A summary of these online ensemble classifiers is provided in Table II. Our ensemble, GOOWE, that we present in the next section, is also included in the table for comparison. As we can see, GOOWE's model management strategies are the same as AWE and AUE2.

Ensemble size. It is also called ensemble cardinality in some studies. Determining the number of component classifiers for an ensemble, discussed briefly in [Gomes et al. 2017; Krawczyk et al. 2017], is an important problem since it has high impact on the prediction ability of an ensemble, and resource consumptions, in terms of time and memory. Our study [Bonab and Can 2016] shows that the intuition of adding more classifiers will result in greater accuracy, is incorrect in practice. In the context of data stream classification, the ensemble size can either be defined fixed, or dynamic, prior to the execution. While there is a lack of studies for determining the size of an online ensemble, most of the existing studies for batch ensembles use statistical tests for determining the proper number of components [Latinne et al. 2001; Oshiro et al. 2012; Hernández-Lobato et al. 2013]. Our geometric framework used for the weighting of components of GOOWE, is also used for determining the ideal number of classifiers for online ensembles, in a theoretical perspective. Increasing or decreasing the number of classifiers from this ideal point deteriorates predictions. We called it “the law of diminishing returns in ensemble construction.” Our theoretical study shows that using the same number of independent component classifiers as class labels gives the highest accuracy [Bonab and Can 2016].

3. GOOWE: GEOMETRICALLY OPTIMUM AND ONLINE-WEIGHTED ENSEMBLE

Concepts and Motivation. Unlike traditional batch learning, the assumption of independent and identical distribution (i.i.d) of the whole stream data is not true for evolving online environments [Gama et al. 2013]. The possibilities of changes are; “feature changes”, or evolving of $p(x)$ with time stamp t , “conditional change”, or the changes of class label y assignment to feature vector x , and “dual changes”, which includes both [Gao et al. 2007]. Four recognized patterns of conditional change are given in Fig. 1. The same patterns of change are possible for feature changes. As mentioned in Section 2.1, Zhang et al. [Zhang et al. 2008] categorized these change into LCD and RCD scenarios. An effective classification algorithm should be able to handle these continuous changes.

The data stream is sliced into chunks, each representing a single distribution. Almost all state-of-the-art stream classifiers divide the data into fixed chunk sizes, as




Fig. 2. Data Chunk (DC) vs. Instance Window (I)—stream data is sliced into equal chunks with size of h and sliding instance window takes the latest n instances with available labels; filled circles are instances with available labels and unfilled circles are yet-to-come instances.

h [Mustafa et al. 2014]. There is a recent study for dynamic determination of chunk size according to concept drift speed [Mustafa et al. 2014]. This problem is beyond the scope of our study.

Depending on when the labeled training data becomes available, Gao et al. [Gao et al. 2007] categorized stream classifiers into two groups: The first group updates the training distribution as soon as the labeled instance becomes available, and the second group receives labeled data in chunks and updates the model. Since updating classifiers is a costly operation, the second group of classifiers can be more time efficient. However, these methods perform well when the up-to-date data chunk has identical or similar distributions to the yet-to-come data chunk, which is called a stationary assumption in the data stream. This assumption ignores the instable nature of evolving data streams when concept drift occurs frequently.

To make our ensemble more efficient, we update component classifiers when a new chunk of labeled data is received. Although we do not address concept drift adaption directly, our extensive experiments show that using a proper component weighting system based on very recent instances would adapt existing component classifiers for recent concept changes. Consequently, having an optimum weighting function would be extremely beneficial for handling concept drift. For this purpose, we exploit a sliding instance window with the latest n labeled instances. The size of the instance window can vary with chunk size, $h \neq n$. The instance window size can be determined by the performance and accuracy requirements of the problem. Fig. 2 shows this combination usage of data chunk and instance window.

Inspired from the geometry of voting [Saari 2008] and using the least squares problem (LSQ) [Hansen et al. 2013], we designed a geometrically optimum and online-weighted ensemble method for evolving environments, called GOOWE. While LSQ is used for component weighting of ensemble classifiers in batch mode [Chan 1999; Friedman 2002], it is the first time that we provide a spatial modeling for online environments as a stacking algorithm.

The motivation of this study is to design an ensemble that assigns optimum weights to component classifiers, in an on-line setting with different types of concept drifts. For combining votes, as a stacking algorithm, we model scores of the ensemble's individual classifiers in a spatial environment as vectors, and try to establish a clear relationship between a geometric feature of vectors, and their effectiveness. Its novelty is based on a dynamically changing component optimum weight assignment approach for online ensembles in evolving data streams.

Design. GOOWE's model management approach is similar to AWE and AUE2, with a passive approach for handling concept drift. Basically, a new incremental learning algorithm is trained on each incoming data chunk, and the previous components of the




Fig. 3. General schema of GOOWE; each $I_t \in S$ delivered to $CS_j (1 \leq j \leq m)$, produces relevance score vector, s_{tj} . GOOWE maps these score vectors, as a score-polytope, and the true class label, as an ideal-point, to a p-dimensional space. It assigns weights, W_j , using the linear least squares (LSQ) solution. The predicted class label, y'_t , is obtained using the weighted majority voting [Bonab and Can 2016].

ensemble are re-evaluated on the same data chunk. However, these evaluations are done with a special function of mean square error (MSE), allowing the algorithm to assign the weights of component classifiers dynamically, relative to each other, and in an on-line setting.

In the training scenario, we use data chunks according to Fig. 2, as they become available. When a new data chunk is received, we train a new component classifier using these instances and we add it to the ensemble. If there is no space for the new classifier, we substitute it with the worst-performing component. For testing the ensemble and classifying a new instance, we use our LSQ-based stacking algorithm based on the sliding instance window for getting the most updated weights for adapting existing components. Briefly, GOOWE uses a combination of data chunks and instance windows, as shown in Fig 2. A *data chunk (DC)* has h instances of a equally divided data stream; an *instance window (I)* has the latest n instances of a data stream, with available true class labels. In our implementation, we build the instance window with the length of $\max(n, h)$, and simply add a counter with the maximum value of h into the instance window for providing the data chunk. If the length of the instance window is less than the length of data chunk (i.e. $n < h$), we set the length of instance window to h and use the latest n instances.

In our geometric framework, we use the Euclidean norm as the system's loss function for optimization purposes. There are clear statistical, mathematical, and computational advantages of using the Euclidean norm [Hansen et al. 2013]. We calculate weights based on the latest n instances in our window, and to make a prediction we use a weighted majority voting approach.

As shown in Fig. 3, we have an ensemble of component classifiers $\xi = \{CS_1, CS_2, \dots, CS_m\}$. Each component classifier, $CS_j (1 \leq j \leq m)$, processes instance I_t of an evolving data stream, S , and produces relevance scores, $s_j = \langle S_j^1, S_j^2, \dots, S_j^p \rangle$, for each of the class labels, $C = \{C_1, C_2, \dots, C_p\}$. Since each classifier produces relevance scores in different ranges, we use Eq. 1 for normalizing the scores into the range of $[0, 1]$.

$$S_j^k \leftarrow \frac{S_j^k}{\sum_{a=1}^p S_j^a} \quad (1 \leq k \leq p) \quad (1)$$

Taking each class label as one dimension, enables us to map each component's score ($s_j; 1 \leq j \leq m$) into a point in a p-dimensional Euclidean space. Mapping all score points of I_t in the same way, builds a polytope in a p-dimensional Euclidean space, which we call the *score-polytope* of I_t . We define *score-vector* by using the origin point as the starting point and score point as the terminal point in our spatial environment. Using the vector of the true class label for I_t as y_t , we can assume an *ideal-point* in the p-dimensional space as $o = \langle O^1, O^2, \dots, O^p \rangle$. For example, if the number of class labels is 4, and the true class label of I_t is C_2 , then the ideal-point would be $o = \langle 0, 1, 0, 0 \rangle$.

Optimum Weight Assignment. For making predictions, we use n latest instances $I = \{I_1, I_2, \dots, I_n\}$, as an instance window, where I_n is the latest instance and all true class labels are available. For each instance $I_i (1 \leq i \leq n)$, each component classifier $CS_j (1 \leq j \leq m)$ has a score-vector as $s_{ij} = \langle S_{ij}^1, S_{ij}^2, \dots, S_{ij}^p \rangle$. 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. We aim to find the optimum weight vector $w = \langle W_1, W_2, \dots, W_m \rangle$ to minimize the distance between score-polytope and ideal-point. Using the squared Euclidean norm as our measure of closeness for the linear least squares problem (LSQ) results in

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

The corresponding residual vector is $r = o - Sw$, where for each instance I_i , $S \in \mathbb{R}^{m \times p}$ is the matrix with relevance scores s_{ij} in each row, w is the vector of weights to be determined, and o is the vector of ideal-point [Hansen et al. 2013]. Since we have n instances in our window, we use the following function for our optimization solution.

$$f(W_1, W_2, \dots, W_m) = \sum_{i=1}^n \sum_{k=1}^p \left(\sum_{j=1}^m (W_j S_{ij}^k) - O_i^k \right)^2 \quad (3)$$

Taking a partial derivation over $W_q (1 \leq q \leq m)$ and finding optimum points will give us our weight vector. The gradient equations become

$$\frac{\partial f}{\partial W_q} = \sum_{i=1}^n \sum_{k=1}^p 2 \left(\sum_{j=1}^m (W_j S_{ij}^k) - O_i^k \right) S_{iq}^k, \quad (1 \leq q \leq m) \quad (4)$$

Setting the gradient to zero, $\nabla f = 0$

$$\sum_{j=1}^m W_j \left(\sum_{i=1}^n \sum_{k=1}^p S_{iq}^k S_{ij}^k \right) = \sum_{i=1}^n \sum_{k=1}^p O_i^k S_{iq}^k, \quad (1 \leq q \leq m) \quad (5)$$

and assuming below summations as a_{qj} and d_q

$$a_{qj} = \sum_{i=1}^n \sum_{k=1}^p S_{iq}^k S_{ij}^k, \quad (1 \leq q, j \leq m) \quad (6)$$

$$d_q = \sum_{i=1}^n \sum_{k=1}^p O_i^k S_{iq}^k, \quad (1 \leq q \leq m) \quad (7)$$

lead to m linear equations with m variables (weights). The proper weights in the following matrix equation are our intended optimum weight vector. We present the weight

assignment equation in matrix representation to make the later example easier to follow.

$$\begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1m} \\ a_{21} & a_{22} & \cdots & a_{2m} \\ \vdots & \vdots & \ddots & \vdots \\ a_{m1} & a_{m2} & \cdots & a_{mm} \end{bmatrix} \times \begin{bmatrix} W_1 \\ W_2 \\ \vdots \\ W_m \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \\ \vdots \\ d_m \end{bmatrix} \quad (8)$$

Briefly, $Aw = d$, where A is the coefficients matrix and d is the remainders vector. According to Eq. 6, A is a symmetric square matrix. In the sense of the least squares solution [Hansen et al. 2013], since it is probable that A is rank-deficient, we may not have a unique solution and we denote the minimizer by w^* . According to Theorem 9 of [Hansen et al. 2013], the *normal equations* for w^* can be written as

$$A^T A w = A^T d \quad (9)$$

In this equation, $A^T A$, is also a symmetric square matrix. In addition, if A has full rank, $A^T A$ is positive definite and our problem has a unique solution. In the rank-deficient case, it is a non-negative definite, and we have a set of possible weight vectors. The QR factorization suggests less expensive solutions for both full rank and rank-deficient cases [Hansen et al. 2013]. In such cases, the weights are nearly optimal.

Since we predict scores for each incoming instance separately, we define A_i and d_i ($1 \leq i \leq n$) according to Eq. 6 and Eq. 7. Matrix A and vector d can be calculated simply by adding all A_i and d_i for all instances of a given window, respectively.

$$a_{qj}^i = \sum_{k=1}^p S_{iq}^k S_{ij}^k, \quad (1 \leq i \leq n) \quad (10)$$

$$d_q^i = \sum_{k=1}^p O_i^k S_{iq}^k, \quad (1 \leq i \leq n) \quad (11)$$

Using the weighted majority vote approach gives the aggregated score vector. Since we calculate scores in a spatial environment, it is possible that these score values become negative. Using the following normalization before Eq. 1 gives the proper aggregated score vector.

$$S_k \leftarrow \frac{S_k - \min(S_k)}{\max(S_k) - \min(S_k)}, \quad (1 \leq k \leq p) \quad (12)$$

Example - Assigning Optimal Weights for Component Classifiers. Suppose that we have 2 classifiers and 2 class labels, as shown in Fig. 4. Our instance window has 2 instances as I_1 and I_2 . We want to find the optimum weight vector for aggregating scores for a newly arrived instance as I_t .




Fig. 4. An example of GOOWE component classifiers weighting.




Fig. 5. Score vectors for instance window of example.

We have a 2-dimensional Euclidean space, as shown in Fig. 5. Score vectors and their intended projections are illustrated with black and red lines, respectively. Putting the values into Equation 6 and 7, gives the following matrix equation.

$$\begin{bmatrix} 1.37 & 1.11 \\ 1.11 & 1.05 \end{bmatrix} \times \begin{bmatrix} W_1 \\ W_2 \end{bmatrix} = \begin{bmatrix} 1.61 \\ 1.18 \end{bmatrix}$$

Solving this equation gives the intended weight vector, $w = \langle 1.88, -0.87 \rangle$. Multiplying these weights with the score vectors of the components, results in the aggregated score vector, $s = \langle 0.86, 0.14 \rangle$. We have a much stronger vote compared to each individual classifier.

Pseudocode of GOOWE Algorithm. It is given in Algorithm 1. In the training scenario (lines 9-23), having the proper number of instances from each class label, as our training data, is crucial for more accurate individual classifiers. On the other hand, for the testing scenario (lines 3-8), static weighting component classifiers can result in relatively poor aggregated predictions, especially with the existence of frequent concept drifts in data stream. Using a combination of data chunk and instance window enables us to think about training and testing of our algorithm separately. These two values can be adjusted according to the drift rate of the data stream.

When the number of instances in the data chunk, DC, reaches its maximum value (line 9), GOOWE trains a new incremental classifier (line 10). If the ensemble has its maximum number of classifiers, m , then GOOWE calculates the weights of classifiers using Eq. 8 and the instances in the data chunk (lines 12-16). The more the obtained weight value is close to zero, the more we want to cancel its effectiveness in our aggregated score vector. As a result, we take the absolute value of weight values and omit the classifier with the least weight (line 17). We first incrementally update all the existing classifiers with DC (lines 19-21), and then add a fresh classifier into the ensemble (line 22). Most of the incrementally updated classifiers need to be pruned after some updates. Since we have memory constraints in our problem, we prune these classifiers when the consumed memory exceeds the memory limit (lines 24-26). For example, in our experiments we use the Hoeffding tree [Domingos and Hulten 2000], and prune the least active leaves of the tree to satisfy the user-specified memory constraint.

For making the class label prediction for each incoming instance, GOOWE calculates the weights of classifiers using Eq. 8 and the instances in the instance window (lines 3-7). It multiplies the resulting weights with score vectors, and using the weighted majority voting approach, calculates the aggregated score vector. Adjusting the length of the instance window and data chunk depends on the data stream and types of con-

ALGORITHM 1: GOOWE (Geometrically Optimum and Online-Weighted Ensemble)

Require: S : data stream, I : window of n latest instances, DC : latest data chunk with length of h , m : maximum number of classifiers, CS : single classifier system, p : number of class labels, L : memory limit.

Ensure: ξ : set of weighted classifiers, s_T : aggregated score vector.

```

1:  $\xi \leftarrow \emptyset$ ;
2: while  $S$  has more instances do
3:   for all instances  $I_i \in I$  do
4:      $A \leftarrow A + A_i$ ; {using Eq. 10}
5:      $d \leftarrow d + d_i$ ; {using Eq. 11}
6:   end for
7:    $w \leftarrow \text{solve}(Aw = d)$ ; {see Eq. 8}
8:    $s_T \leftarrow \sum_{j=1}^m (W_j s_j)$ ; {weighted majority vote}
9:   if  $DC$  has  $h$  instances then
10:     $CS' \leftarrow$  new single classifier built on  $DC$ ;
11:    if  $\xi$  has  $m$  classifiers then
12:      for all instances  $I_i \in DC$  do
13:         $A' \leftarrow A' + A'_i$ ; {using Eq. 10}
14:         $d' \leftarrow d' + d'_i$ ; {using Eq. 11}
15:      end for
16:       $w' \leftarrow \text{solve}(A'w' = d')$ ; {see Eq. 8}
17:       $\xi \leftarrow \xi \setminus \{\text{classifier with } \min(|W'_j|); 1 \leq j \leq m\}$ ;
18:    end if
19:    for all  $CS_j \in \xi$  do
20:      train  $CS_j$  with  $DC$ ;
21:    end for
22:     $\xi \leftarrow \xi \cup \{CS'\}$ ;
23:  end if
24:  if  $\text{memory\_usage}(\xi) \geq L$  then
25:    prune all component classifiers;
26:  end if
27: end while

```

cept drift. There is no general solution to this problem. However, setting relatively small values to the instance window, and relatively large values to the data chunk, according to available resources, can result in better accuracy.

Experimental evaluations, presented in the following sections, illustrate that GOOWE can react statistically significantly better compared to its state-of-the-art rivals.

4. EXPERIMENTAL EVALUATION SETUP

The main concerns for evolving data stream classifiers are more accurate predictions with less memory consumption, and less processing time. In addition, any proposed method for an evolving data stream needs to be careful with concept drift, and react accordingly. In the following sections, we present our experimental evaluation for different simulation scenarios conducted to evaluate our proposed ensemble.

In summary, our experimental evaluation is presented as follows. We

- First*, describe synthetic and real-world evolving data streams used in our experiments. We explain each with the type of concept drift, the number of class labels, and the number of instances used. While there is a shortage in trusted evolving real-world streams [Krawczyk et al. 2017], we try to include all possible known/unknown categories of concept drift in our experiments. We also specify our experimental

- framework setup, implementation details, used libraries, and more for reproducibility purpose (current section).
- *Second*, provide an analysis conducted for examining two major differentiating elements of GOOWE, component weighting strategy and ensemble model management strategy (Section 5). Our optimum and online weighting system shows its effectiveness for both vote aggregation and ensemble maintenance.
- *Last but not least*, present our extensive and comparative experiments. We compare GOOWE with state-of-the-art rival ensembles and extensively discuss the superiority conditions. For the sake of comparison, we include 8 state-of-the-art adaptive ensemble methods proposed for evolving data streams (Section 6).

4.1. Datasets as Data Streams

Selecting proper time-evolving data streams is one of the vital steps for comparing different algorithms. There are two types of data stream sets—synthetic and real-world datasets. We generate the whole dataset before the experiment, and use the terms dataset and data stream equivalently. Similarly to other domains of prediction algorithms, real-world datasets are the best. However, their problem is that we do not know when drift occurs, or if there is any drift at all. Some studies use real-world datasets with artificial concept drifts, called real-world data with forced/synthetic concept drift [Gama et al. 2014]. These datasets cannot be considered as real examples of drifts. Synthetic data has several benefits like being easy to reproduce, having a low cost of storage and transmission, but most importantly, it provides an advantage of knowing where exactly drift has happened [Bifet et al. 2009; Gama et al. 2014].

A proposed algorithm should be capable of handling large data streams—with potentially an infinite number of instances [Bifet et al. 2009]. As a result, for the comparison of several algorithms, we need to have large datasets in the order of tens of millions of instances. Similar to common approaches [Bifet et al. 2009; Brzezinski and Stefanowski 2014b; 2014a; Street and Kim 2001], in order to cover all patterns of changes over time; sudden/abrupt, incremental, gradual, and reoccurring as concept drifts including blips or noise; we use synthetic data stream generators, implemented in the MOA framework. Using these generators, we prepared 16 synthetic datasets. In addition, we have 4 widely used real-world data streams.

Following are a brief description of each dataset including their generation and preparation. Table III summarizes the specifications of each dataset. We report the average of accuracy, processing time, and maximum memory consumption for each dataset in Table VI, VII, and VIII, respectively.

4.1.1. Synthetic Datasets . According to the concept drift scenarios of Zhang et al. [Zhang et al. 2008], we have 8 Rigorous Concept Drifting (RCD) and 8 Loose Concept Drifting (LCD) synthetic datasets. Bifet et al. [Bifet et al. 2009] specified Random RBF generator as the RCD data stream, and the rest of synthetic data stream generators as the LCD data stream.

Random RBF. It assigns a fixed number of random positioned centroids, with a random standard deviation value, class label, and weight. For generating new instances, we randomly select a center, considering weights, so that centroids with higher weights are more likely to be chosen. A random direction is chosen for displacement, using a Gaussian distribution, and drift is defined by moving the centroids, with constant speed. Attributes are all numerical values. Using this generator we prepared 8 different datasets, each containing 1 million instances, with 20 attributes, and 0 percent noise. Here are 3 important alternate factors we changed among these 8 datasets. We reflect these, respectively, in the naming of RBF datasets in Table III.

Table III. Summary of Dataset Characteristics

Dataset	#Instance	#Att	#CL	%N	Drift Spec.
RBF-G-4-S	1×10^6	20	4	0	Gr., Bp., DS=0.0001
RBF-G-4-F	1×10^6	20	4	0	Gr., Bp., DS=0.01
RBF-G-10-S	1×10^6	20	10	0	Gr., Bp., DS=0.0001
RBF-G-10-F	1×10^6	20	10	0	Gr., Bp., DS=0.01
RBF-A-4-S	1×10^6	20	4	0	Abrupt, #D=10
RBF-A-4-F	1×10^6	20	4	0	Abrupt, #D=100
RBF-A-10-S	1×10^6	20	10	0	Abrupt, #D=10
RBF-A-10-F	1×10^6	20	10	0	Abrupt, #D=100
SEA-S	1×10^6	3	2	10	Abrupt, #D=3
SEA-F	2×10^6	3	2	10	Abrupt, #D=9
HYP-S	1×10^6	10	2	5	Incrm., DS=0.001
HYP-F	1×10^6	10	2	5	Incrm., DS=0.1
TREE-S	1×10^6	10	4	0	Reoc., #D=4
TREE-F	1×10^5	10	6	0	Reoc., #D=15
LED-M	1×10^6	24	10	10	Mixed, #D=3
LED-ND	1×10^7	24	10	20	No drift
CoverType	581,012	54	7	-	Unknown
PokerHand	1×10^7	10	10	-	Unknown
CovPokElec	1,455,525	72	10	-	Unknown
Airlines	539,383	7	2	-	Unknown

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

- *Concept Drift Type (Gradual: G and Abrupt: A)*. The way the generator moves centroids make the data stream gradually changing. We add some outliers during generations of gradual changing datasets in order to have blips. We generate abruptly changing data streams using the sigmoid join operator ($c = a \oplus_{t_0}^W b$; t_0 : point of change, W : length of change) [Bifet et al. 2009].
- *Number of Classes (Four: 4 and Ten: 10)*. The ability to generate an arbitrary number of classes is useful for evaluating an algorithm. We generate our datasets with either four or ten class labels.
- *Drift Frequency (Slow: S and Fast: F)*. For gradually changing datasets, we generate instances with 0.01 (fast) and 0.0001 (slow) concept change speed (defined as moving centroids in a random direction for a predefined distance of 0.01 or 0.001, within each 500 instances). For abruptly changing datasets, we switch to a new random stream generator that generates data stream with zero concept changing speed, 10 (slow) or 100 (fast), times evenly distributed over 1 million instances.

SEA Concepts. It involves 3 numerical attributes varying between 0 and 10 [Street and Kim 2001]. In our experiment, we use this generator in 2 different settings, both with 10 percent noise. First, 1 million instances, with drifts occurring every 250,000 examples (slow: SEA-S), and second, 2 million instances with drifts occurring every 200,000 examples (fast: SEA-F) are generated.

Rotating Hyperplane. It assigns points in a multi-dimensional hyperplane and classifies them positively and negatively. Concept drift is defined by changing the orientation and position of the hyperplane [Hulten et al. 2001]. We set the hyperplane generator to create 2 datasets, each with 1 million instances described by 10 numerical features. We add 5 percent class noise to both of them. The modification weight of slowly changing dataset (HYP-S) is set to $w_i = 0.001$, and for the rapidly changing one (HYP-F) to $w_i = 0.1$.

Random Tree. It produces nominal and numerical attributes using a randomly constructed tree. Drift is defined by abruptly changing the tree after a given number of examples [Bifet et al. 2010b]. For both slow and fast tree datasets, we set the generator to have 5 nominal and 5 numerical attributes. The slowly changing dataset (TREE-S) consists of 1 million instances, with 4 evenly distributed reoccurring drifts. The rapidly changing dataset (TREE-F) contains 100,000 instances with 15 sudden drifts; it is the fastest changing dataset in our experiments.

LED. It tries to predict the digit displayed on a seven-segment LED display. Each instance has 24 binary attributes and each has a possibility of being inverted, which is defined as noise. We have 2 LED datasets. The first dataset, LED-M, has 1 million instances with 2 gradually drifting concepts abruptly switching after 0.5 million instances, and 10 percent noise. The second, LED-ND, has 10 million instances without any drift and 20 percent noise, making it the noisiest and largest dataset [Brzezinski and Stefanowski 2014b].

4.1.2. Real-World Datasets. The noise values, number of drifts, and drift speeds are unknown for these datasets. Access URL links are given in the footnote.

CoverType.¹ It contains the forest cover type from the US Forest Service (USFS), comprised of 581,012 instances and 54 attributes.

PokerHand.² It consists of 1 million instances and 10 attributes. Each record is a hand of 5 playing cards—with 2 attributes as suit and rank.

CovPokElec.³ It combines the normalized CoverType, normalized PokerHand, and Electricity datasets using the sigmoid join operator. The Electricity dataset comes from the Australian New South Wales Electricity Market. CovPokElec is obtained by merging all attributes, and assuming that each dataset corresponds to a different concept [Bifet et al. 2009].

Airlines.⁴ It consists of 539,383 examples described by 7 attributes. The task is to predict whether or not a given flight will be delayed, given the information of the scheduled departure.

4.2. Experimental Framework: Detailed Design

Implementation details. In this paper, we use the Massive Online Analysis (MOA)⁵ framework [Bifet et al. 2010b]. MOA is an open-source software package to run data streaming experiments and, to the best of our knowledge, is the most popular framework for data stream mining. We use the Java MAtrix (JAMA)⁶ package, a basic linear algebra library, for matrix operations and to find least squares solutions in our implementation of GOOWE. We extended MOA for GOOWE implementation using the Java programming language. Some of the other ensemble algorithms, we used as baselines; they are implemented as part of the MOA framework. We used the MOA extensions library for DWM and NSE. In addition, our implementation of GOOWE, and some detailed information about experimental evaluation, such as standard deviations, and dataset generations are available on our GitHub webpage⁷.

Experimental Analysis. We first study the impact of the proposed weighting system on vote aggregation and ensemble maintenance using two scenarios. In both of these scenarios, we use a fixed block-based ensemble, while different weighting sys-

¹Access link: <http://archive.ics.uci.edu/ml/datasets/Covertype>

²Access link: <http://archive.ics.uci.edu/ml/datasets/Poker+Hand>

³Access link: <http://www.openml.org/d/149>

⁴Access link: <http://moa.cms.waikato.ac.nz/datasets/>

⁵MOA webpage: <http://moa.cms.waikato.ac.nz/>

⁶JAMA webpage: <http://math.nist.gov/javanumerics/jama/>

⁷GOOWE webpage: <https://hamedrab.github.io/GOOWE/>

tems are implemented in parallel to the original weighting system. In this way, we may study a single impact factor, and cancel all other impact factors. Through this analysis, GOOWE's weighting system is compared to most similar block-based ensembles, i.e. AUE2, AWE, and DWM, and some other baselines based on GOOWE's weighting system.

Comparative Study. For our comparative study, we evaluate GOOWE by comparing it with 8 well-known ensemble classifiers for non-stationary environments using the online block-based, bagging, and boosting methods as baselines. We select AWE, AUE2, DWM, and NSE ensemble methods from block-based approaches. In addition to these, we include OAUE, OzaBag, OzaBoost, and LevBag ensemble methods as popular online ensembles proven to have reasonable performance in evolving environments.

Ensemble Size. As discussed in Section 2, ensemble size has an important impact on performance of different algorithms. We suggest in [Bonab and Can 2016] to have the same number of component classifiers as class labels. For our experimental analyses, we use the same number of classifiers as the number of class labels for each data stream. However, in order to ease the comparisons of time and memory consumption values, and to follow the convention in the literature of using a fixed maximum number of classifiers, we fixed ensemble size for our comparative study. We set the maximum number of classifiers to 10. Studies based on a fixed number of classifiers are acceptable, since in such cases all ensemble methods can be equally disadvantaged [Bonab and Can 2016].

Base Classifier. We use the Hoeffding tree [Domingos and Hulten 2000] as the base classifier component for all examined ensemble methods. We use the Hoeffding tree enhanced with adaptive Naive Bayes leaf predictions, with a grace period $n_{min} = 100$, split confidence $\delta = 0.01$, and tie-threshold $\tau = 0.05$ similar to experiments in [Brzezinski and Stefanowski 2014b; 2014a; Domingos and Hulten 2000].

Chunk and Instance Window Size. In our experiments, according to the chunk size analysis of [Wang et al. 2003] and similar to the experimental evaluations of [Brzezinski and Stefanowski 2014b], the chunk size for block-based ensembles (namely DWM, NSE, AWE, AUE2, and GOOWE) is set to 500 instances. OAUE and GOOWE use a sliding window of recent data instances. To ensure a fair comparison, similar to block-based ensembles, we set the instance window length to 500 instances. Although this length can be smaller for most of the ensembles, to perform an equivalent comparison, we choose this value based on the suggested minimum chunk length of AWE [Wang et al. 2003]. The data chunk size and instance window size analysis is possible as a future work.

Measurements. By considering the main requirements of data stream environments [Bifet et al. 2009; Brzezinski and Stefanowski 2014b; Street and Kim 2001] in our experimental setup, we chose the interleaved Test-Then-Train procedure for measuring prediction accuracy values. For time and memory measurements, we use CentiSecond (CS) and MegaByte (MB), respectively. Our initial experiments showed that for synthetic datasets with the exact same settings of data stream generators, accuracy, time, and memory measurements showed variations. In order to have confident conclusions, for each synthetic data stream, we generate 10 time-seeded random datasets. For example, when we say that RBF-G-4-F dataset has 1 million of instances, we examine 10 such datasets (i.e. a total of 10 millions of instances) and report the mean value among these 10.

Machine Specification. The experiments were performed on a machine equipped with an Intel Xeon E3-1200 v3 @ 3.40 GHz processor and 32 GB of ECC RAM.

Table IV. Classification Accuracy in Percentage (%) for Vote Aggregation Analysis on Data Streams with Concept Drift—*Base1* ensemble method with different weighting systems used for aggregating votes

Dataset	MV	DWM ($\beta = 0.5$)	DWM ($\beta = 0.2$)	AWE	AUE2	GOOWE	GOOWE-Min	GOOWE-Max
RBF-G-4-S	31.854	31.983	31.989	30.834	31.214	33.853	29.692	33.627
RBF-G-4-F	91.746	85.888	85.666	90.868	91.668	91.626	72.478	87.066
RBF-G-10-S	15.444	14.857	14.867	14.674	15.036	17.395	13.444	15.733
RBF-G-10-F	80.794	80.956	80.929	80.939	80.864	84.062	77.054	78.378
RBF-A-4-S	93.040	90.476	89.851	90.099	93.037	92.983	71.727	90.768
RBF-A-4-F	93.737	90.110	90.498	93.794	93.699	93.627	70.295	91.008
RBF-A-10-S	90.460	90.017	90.011	90.980	90.675	93.869	80.094	85.469
RBF-A-10-F	89.402	88.474	87.950	86.842	89.572	92.622	80.934	84.916
SEA-S	85.636	86.927	86.921	85.289	85.847	84.510	82.364	83.670
SEA-F	89.433	89.302	89.288	89.520	89.230	89.409	86.289	89.310
HYP-S	83.140	87.461	87.462	86.811	84.587	82.429	80.351	82.134
HYP-F	90.742	90.955	90.953	90.981	90.382	91.189	88.434	91.007
TREE-S	94.632	94.452	94.452	94.750	94.599	94.796	58.737	94.470
TREE-F	82.280	81.993	81.963	81.445	82.199	82.560	54.021	82.255
LED-M	73.649	73.266	73.256	72.646	73.645	73.599	69.046	73.575
CoverType	86.516	87.841	87.881	85.655	83.306	88.139	75.597	87.609
PokerHand	66.851	68.707	68.441	66.451	66.826	71.823	60.204	63.281
CovPokElec	74.845	75.911	75.831	75.584	74.909	79.519	68.341	71.845
Airlines	62.136	62.663	62.689	62.041	62.293	62.368	60.957	62.116

5. EXPERIMENTAL ANALYSES: THE IMPACT OF WEIGHTING AND MODEL MANAGEMENT STRATEGIES OF GOOWE

In this section, we mainly focus on answering the question: why should GOOWE work better in terms of prediction accuracy, or to put it in other words, when/where in the learning process does GOOWE get its advantage? To answer this question, we need to study the impact of GOOWE's weighting system on vote aggregation and ensemble maintenance in evolving environments as two major features of GOOWE. These two features differentiate GOOWE from other block-based ensembles, and we show the superiority of GOOWE compared to other ensembles based on these two key features.

We designed two scenarios for studying the impact of the weighting system of GOOWE on vote aggregation and ensemble maintenance. Detailed information regarding each of these scenarios are given in the following. The main idea in both analyses is that by isolating the examined feature the impact can be studied. We choose a basic and comparably good ensemble method, and fix all settings for training and testing, except for the studied one (vote aggregation or ensemble maintenance). Here for our analyses, we exploit AUE2 implementation from the MOA framework as the base ensemble, since the weighting system of other block-based ensembles can be applied easily; it is also one of the leading ensembles. For the following scenarios of analyses, we created two versions as *Base1* and *Base2*. *Base1* includes every detail of the AUE2 ensemble, except its vote aggregation. *Base2* includes every detail of the AUE2 ensemble, except decisions on add/drop components. Further explanations are provided for each of these in the following. Using these analyses, we can verify GOOWE's weighting system superiority without benefiting from other specifications of each ensemble.

5.1. Analysis of Vote Aggregation

For evaluating the impact of the weighting strategy proposed for GOOWE on vote aggregation, as previously described, we use the AUE2 implementation from the MOA framework, except for its vote aggregation, as the base ensemble method, called *Base1*. We implement GOOWE's weighting system for the *Base1* ensemble classifier. As a result, the only variant to this new ensemble, compared to the original AUE2 version,




Fig. 6. The Friedman statistical test average ranks for different vote aggregation rules. Higher average rank means better prediction accuracy. The minimum required difference of average ranks for considering a statistically significant difference is 1.197.

is our weighting system for vote aggregation. In this way, we are able to study the impact of any weighting function in vote aggregation on the accuracy of predictions.

In order to have different vote aggregation rules as our baselines, we also implement Majority Voting (MV), DWM with punishment constant values of 0.5 and 0.2, and also AWE's weighting systems for *Base1* ensemble. In addition, we include the prediction accuracy of the component corresponding to the least/highest weight obtained from GOOWE weighting system (in Table IV illustrated as GOOWE-Min and GOOWE-Max). GOOWE-Min presents the worst-performing component and GOOWE-Max presents the best performing component, according to GOOWE's weights. We conduct our analysis using these as state-of-the-art baselines of weighting systems.

Table IV presents the accuracy values obtained from the mentioned vote aggregation rules. Note that in all of these scenarios, the data stream has concept drift. In order to compare these aggregation rules we conduct the non-parametric Friedman statistical test with pairwise comparisons. The null-hypothesis states that all aggregation rules are equal [Demšar 2006; Conover 1999]. Since we have 8 vote aggregation rule and 19 datasets in our experiment, F_F is distributed according to the F distribution with $8 - 1 = 7$ and $(8 - 1) \times (19 - 1) = 126$ degrees of freedom. We run the statistical test at the significance level of $\alpha = 0.05$ and reject the null-hypothesis with a *p-value* of < 0.00001 .

The multiple comparisons average ranks are plotted in Figure 6. The Critical Distances (CD) for $F(8, 152) = 14.802$ is 1.197, meaning average ranks of aggregation rules need to have at least this amount of difference to be considered statistically significantly different, in a pairwise comparison. The weighting system of GOOWE is statistically significantly better compared to all other baseline aggregation rules, as shown in Figure 6. While MV acts very well among remaining aggregation rules in evolving environments, we are not able to claim a statistically significant difference among them—excluding GOOWE-Max and GOOWE-Min.

Based on our preliminary tests, GOOWE's weighting system shows its superiority in evolving environments. For this purpose, we tested our analysis scenario on RBF and LED data streams without any concept drift; there was no meaningful difference between MV and GOOWE's weighting systems. This is because when concept drift happens, GOOWE reacts much faster. The same can be concluded when we compare rapidly changing data streams with slowly changing ones, in Table IV. We will show this with more details through our comparative experiments in the next section.

Table V. Classification Accuracy in Percentage (%) for the Model Management Analysis on Data Streams with Concept Drift—*Base2* ensemble method with different weighting systems used for the decision of components add/drop

Dataset	DWM	AWE	AUE2	GOOWE
RBF-G-4-S	34.077	30.846	31.854	32.110
RBF-G-4-F	89.424	89.99	91.746	92.176
RBF-G-10-S	17.404	14.104	15.444	15.241
RBF-G-10-F	89.992	80.161	80.794	90.995
RBF-A-4-S	93.651	87.399	93.040	93.610
RBF-A-4-F	94.310	87.265	93.737	94.389
RBF-A-10-S	95.015	85.528	90.460	95.437
RBF-A-10-F	95.395	86.230	89.402	95.259
SEA-S	89.353	87.75	85.636	89.196
SEA-F	89.565	89.010	89.433	89.469
HYP-S	86.779	83.726	83.140	83.009
HYP-F	88.035	90.797	90.742	91.189
TREE-S	94.813	94.632	94.632	84.529
TREE-F	82.280	82.280	82.280	78.830
LED-M	73.644	73.619	73.649	73.596
CoverType	88.204	87.344	86.516	88.004
PokerHand	85.716	67.637	66.851	81.702
CovPokElec	88.935	74.818	74.845	81.849
Airlines	64.570	63.084	62.136	62.146

5.2. Analysis of Model Management Strategy

For examining the superiority of our model management strategy, similar to the previous analysis, we use the implementation of AUE2 from the MOA framework, as *Base2*. In this analysis, we use GOOWE and other baseline weights in the process of making decisions for add/drop components. We implement these baselines for the *Base2* ensemble. Note that, for aggregating the votes of components in this analysis, we use majority voting to equally disadvantage all the ensembles. We use DWM with $\beta = 0.5$ and AWE as baselines of this analysis. DWM with $\beta = 0.2$ gave the exact same results as DWM with $\beta = 0.5$. We construct and maintain the *Base2* ensemble using these weighting algorithms for each data stream. Table V presents the resulting accuracies.

In Table V we observe a similar superiority of the GOOWE weighting system for rapidly changing data streams, compared to slowly changing data streams. The same scenario is valid here; GOOWE gets its advantage when more concept drifts happen, while reacting similarly in non-changing environments.

Similarly to previous analysis, we conduct the non-parametric Friedman statistical test with pairwise comparisons. The null-hypothesis states that all model management strategies are equal. Since we have 4 algorithms and 19 datasets in our experiment, F_F is distributed according to the F distribution with $4 - 1 = 3$ and $(4 - 1) \times (19 - 1) = 54$ degrees of freedom. We run the statistical test at the significance level of $\alpha = 0.05$, and get $F(3, 54) = 8.3937$. We are able to reject the null-hypothesis with a *p-value* of 0.0001. Moreover, pairwise multiple comparisons indicate no statistically significant superiority for GOOWE, in ensemble maintenance, compared to DWM and its superiority compared to AUE2 and AWE.

Conclusion of the Experimental Analyses. Our first analysis shows the superiority of GOOWE vote aggregation in evolving environments. The second analysis shows GOOWE's conservative behavior in ensemble maintenance. We can conclude that GOOWE gets its advantage with vote aggregation, while reacting similarly as the best block-based ensembles for model management.

Table VI. Average Classification Accuracy in Percentage (%) —for each synthetic dataset a one-way ANOVA using Scheffe multiple comparisons statistical test is conducted, and the top tier algorithms are underlined; for real-world datasets the most accurate algorithms are underlined

Dataset	DWM	NSE	AWE	AUE2	GOOWE	OAUE	OzaBag	LevBag	OzaBoost
RBF-G-4-S	75.157	72.355	75.329	<u>91.174</u>	<u>92.014</u>	<u>91.817</u>	87.084	85.779	88.353
RBF-G-4-F	74.102	72.041	73.837	<u>94.250</u>	<u>94.590</u>	<u>93.322</u>	87.213	85.947	87.995
RBF-G-10-S	79.549	77.365	81.326	83.102	<u>92.298</u>	83.059	80.901	80.671	76.951
RBF-G-10-F	79.669	78.455	80.875	83.055	<u>92.189</u>	82.726	80.748	80.256	76.459
RBF-A-4-S	76.628	73.308	78.046	<u>96.543</u>	<u>96.901</u>	<u>96.267</u>	95.618	95.676	97.367
RBF-A-4-F	75.452	72.519	77.591	<u>96.779</u>	<u>97.019</u>	<u>95.867</u>	<u>95.461</u>	<u>95.988</u>	<u>96.258</u>
RBF-A-10-S	81.297	79.446	84.832	91.943	<u>96.477</u>	85.771	<u>95.017</u>	<u>94.901</u>	<u>95.136</u>
RBF-A-10-F	82.338	80.471	85.657	92.592	<u>96.730</u>	88.473	<u>95.504</u>	<u>95.480</u>	<u>95.923</u>
SEA-S	86.030	86.847	87.897	<u>89.718</u>	89.031	<u>89.749</u>	89.628	89.633	89.360
SEA-F	86.084	86.849	87.923	<u>89.812</u>	89.637	<u>89.831</u>	89.739	89.742	89.551
HYP-S	86.819	87.175	<u>90.483</u>	88.486	<u>88.891</u>	<u>89.044</u>	83.467	<u>89.222</u>	86.306
HYP-F	90.734	88.714	90.994	<u>92.564</u>	<u>92.567</u>	<u>92.748</u>	82.032	92.148	89.495
TREE-S	32.921	24.638	33.926	35.222	35.932	36.286	37.135	<u>37.124</u>	24.639
TREE-F	28.870	09.512	30.154	31.858	<u>33.827</u>	32.024	32.217	<u>32.217</u>	9.518
LED-M	65.880	70.354	<u>74.002</u>	<u>73.975</u>	73.989	<u>73.973</u>	73.984	74.008	73.778
LED-ND	43.336	46.849	<u>51.210</u>	<u>51.196</u>	<u>51.191</u>	<u>51.208</u>	<u>51.194</u>	<u>51.212</u>	51.034
CoverType	86.266	79.437	80.600	84.951	89.779	88.205	84.264	84.080	90.570
PokerHand	47.591	49.550	49.470	50.217	<u>54.604</u>	50.031	52.995	52.995	53.681
CovPokElec	85.377	65.249	66.330	73.668	<u>88.718</u>	86.390	82.265	77.647	87.154
Airlines	61.206	60.797	60.650	61.395	61.834	<u>62.516</u>	61.404	62.164	61.015

6. COMPARATIVE EVALUATION

In this section we examine GOOWE as an ensemble algorithm, as described in Algorithm 1, and compare it with the 8 most state-of-the-art ensemble methods. We measure class label prediction accuracy (in percentage), maximum memory usage (in MegaByte), and total processing time of every one thousand instances (in CentiSecond) for each of the ensemble algorithms—average values for synthetic datasets and exact values for real-world datasets reported in Table VI, VII, and VIII, respectively. For each synthetic dataset, a one-way analysis of variance (ANOVA) using Scheffe multiple comparisons [Scheffe 1959] are conducted, and the best-performing algorithms are underlined. It is not possible to conduct the Scheffe statistical test for real-world datasets, since they only have a single value. For each of them, we underline the most accurate and least resource consuming algorithm.

We draw scatter diagrams of the algorithms on the arrival of new chunks of data streams, as in [Bifet et al. 2009; Elwell and Polikar 2011; Brzezinski and Stefanowski 2014b]. We provide one plot of accuracy and memory behavior for each category of RCD, LCD, and real-world datasets. For better understanding the behavior of ensembles in these situations, we present accuracy and memory plots for gradual changing RCD, and abrupt changing RCD datasets, separately. We provide these plots in Fig. 7, 8, 9, and 10—note that the plots are in different scales.

6.1. RCD Data Streams with Gradual/Abrupt Drift Patterns

Table VI for Random RBF data streams (the first 8 rows) shows the superiority of GOOWE over other algorithms, in terms of accuracy. Its superiority is more significant for the gradually changing data streams, with respect to the abruptly changing data streams. Comparing the number of class labels suggests that GOOWE performs better for RCD datasets with 10 class-labels, rather than 4. The preliminary experiments show that this relationship changes with the number of component classifiers of the

Table VII. Average Processing Time in CentiSecond (CS), for processing every one thousand instances—for each synthetic dataset a one-way ANOVA using Scheffe multiple comparisons statistical test is conducted, and the top tier algorithms are underlined; for real-world datasets the least time-consuming algorithms are underlined

Dataset	DWM	NSE	AWE	AUE2	GOOWE	OAUE	OzaBag	LevBag	OzaBoost
RBF-G-4-S	6.718	439.071	18.458	20.563	14.887	17.482	8.299	17.825	7.231
RBF-G-4-F	<u>6.628</u>	444.261	21.930	21.039	14.400	17.986	11.234	18.801	14.403
RBF-G-10-S	31.854	1085.484	41.223	45.602	31.569	37.233	<u>17.146</u>	38.402	<u>14.871</u>
RBF-G-10-F	30.501	1124.648	49.549	47.105	31.294	37.599	<u>23.036</u>	38.789	28.333
RBF-A-4-S	<u>6.598</u>	443.676	21.900	17.218	13.784	15.367	9.901	16.425	12.124
RBF-A-4-F	<u>6.621</u>	445.520	21.913	17.037	13.806	15.376	10.007	16.381	11.921
RBF-A-10-S	30.027	1125.125	49.025	44.556	31.202	36.867	<u>21.545</u>	36.214	27.433
RBF-A-10-F	29.693	1119.683	48.678	43.537	30.718	36.231	<u>20.482</u>	35.397	26.59
SEA-S	<u>0.615</u>	43.714	2.752	2.772	2.428	2.760	1.509	2.840	1.808
SEA-F	<u>0.616</u>	94.267	2.740	3.065	2.452	2.958	1.684	3.102	1.963
HYP-S	<u>1.504</u>	115.595	7.390	6.617	6.321	6.547	3.955	5.103	4.241
HYP-F	<u>1.443</u>	79.532	7.534	5.680	5.707	5.877	3.580	4.290	3.873
TREE-S	<u>6.337</u>	85.346	11.468	9.479	9.333	9.528	7.217	8.777	7.009
TREE-F	<u>8.193</u>	14.276	13.639	11.026	10.518	11.036	<u>7.641</u>	9.478	<u>7.295</u>
LED-M	35.260	1288.899	57.355	54.824	45.966	40.135	<u>13.874</u>	25.615	<u>14.764</u>
LED-ND	38.603	1280.011	58.002	54.971	45.615	42.924	<u>16.068</u>	23.699	<u>17.102</u>
CoverType	<u>11.526</u>	335.569	28.269	24.409	21.858	21.952	11.721	18.868	14.359
PokerHand	<u>8.578</u>	68.910	<u>5.184</u>	7.123	6.842	8.816	5.634	8.012	6.148
CovPokElec	15.483	546.412	25.803	24.005	24.076	24.832	<u>13.786</u>	18.814	17.804
Airlines	<u>0.982</u>	11.304	2.692	2.710	3.338	3.146	2.086	2.673	2.676

Table VIII. Maximum Memory Usage in MegaByte (MB) —for each synthetic dataset a one-way ANOVA using Scheffe multiple comparisons statistical test is conducted, and the top tier algorithms are underlined; for real-world datasets the least memory-consuming algorithms are underlined

Dataset	DWM	NSE	AWE	AUE2	GOOWE	OAUE	OzaBag	LevBag	OzaBoost
RBF-G-4-S	<u>0.261</u>	116.193	0.320	0.483	1.632	0.467	13.379	15.221	13.232
RBF-G-4-F	<u>0.294</u>	116.193	<u>0.319</u>	2.885	9.750	2.711	35.258	39.568	24.517
RBF-G-10-S	<u>0.333</u>	116.201	<u>0.394</u>	3.470	11.066	3.358	18.063	20.293	8.085
RBF-G-10-F	<u>0.691</u>	116.200	<u>0.394</u>	5.451	11.837	5.408	16.877	17.382	13.343
RBF-A-4-S	<u>0.931</u>	116.192	<u>0.319</u>	4.535	15.657	4.687	27.618	42.854	30.531
RBF-A-4-F	<u>0.553</u>	116.192	<u>0.319</u>	5.004	12.107	5.026	26.646	40.845	28.919
RBF-A-10-S	<u>0.569</u>	116.200	<u>0.395</u>	4.453	11.320	4.194	19.184	35.068	16.292
RBF-A-10-F	<u>0.842</u>	116.200	<u>0.394</u>	4.682	13.784	5.334	17.168	34.660	15.368
SEA-S	<u>0.181</u>	116.109	<u>0.229</u>	18.974	3.975	21.146	7.046	9.009	6.861
SEA-F	<u>0.204</u>	464.351	<u>0.230</u>	36.584	6.332	42.023	13.826	15.858	13.629
HYP-S	<u>0.384</u>	116.145	<u>0.614</u>	2.376	4.849	2.636	22.918	25.218	21.19
HYP-F	<u>0.505</u>	116.147	<u>0.643</u>	1.063	2.999	1.863	20.912	22.882	21.097
TREE-S	<u>0.172</u>	116.137	<u>0.210</u>	1.426	5.931	7.343	6.288	9.265	6.343
TREE-F	<u>0.181</u>	1.251	<u>0.224</u>	0.681	3.118	0.762	0.538	0.564	0.606
LED-M	<u>0.788</u>	116.222	<u>0.453</u>	<u>0.453</u>	<u>1.052</u>	<u>0.415</u>	13.707	18.715	13.927
LED-ND	<u>0.581</u>	116.221	<u>0.453</u>	<u>0.453</u>	<u>2.871</u>	<u>0.454</u>	12.891	20.912	12.863
CoverType	1.252	39.376	<u>0.437</u>	0.736	1.128	1.079	55.631	71.475	62.185
PokerHand	<u>0.186</u>	116.146	0.224	0.254	0.924	0.229	3.460	4.390	3.524
CovPokElec	4.534	246.144	<u>0.576</u>	11.543	54.435	36.529	132.399	162.592	153.337
Airlines	<u>0.131</u>	33.825	0.157	0.300	2.507	0.535	7.592	10.322	7.773

ensemble. For example, having 4 component classifiers can benefit more from a data stream with 4 class labels.

As shown in Table VI, in most of the cases, GOOWE has higher average accuracy for the fast-changing datasets, compared to the slow changing ones. We can intuitively un-

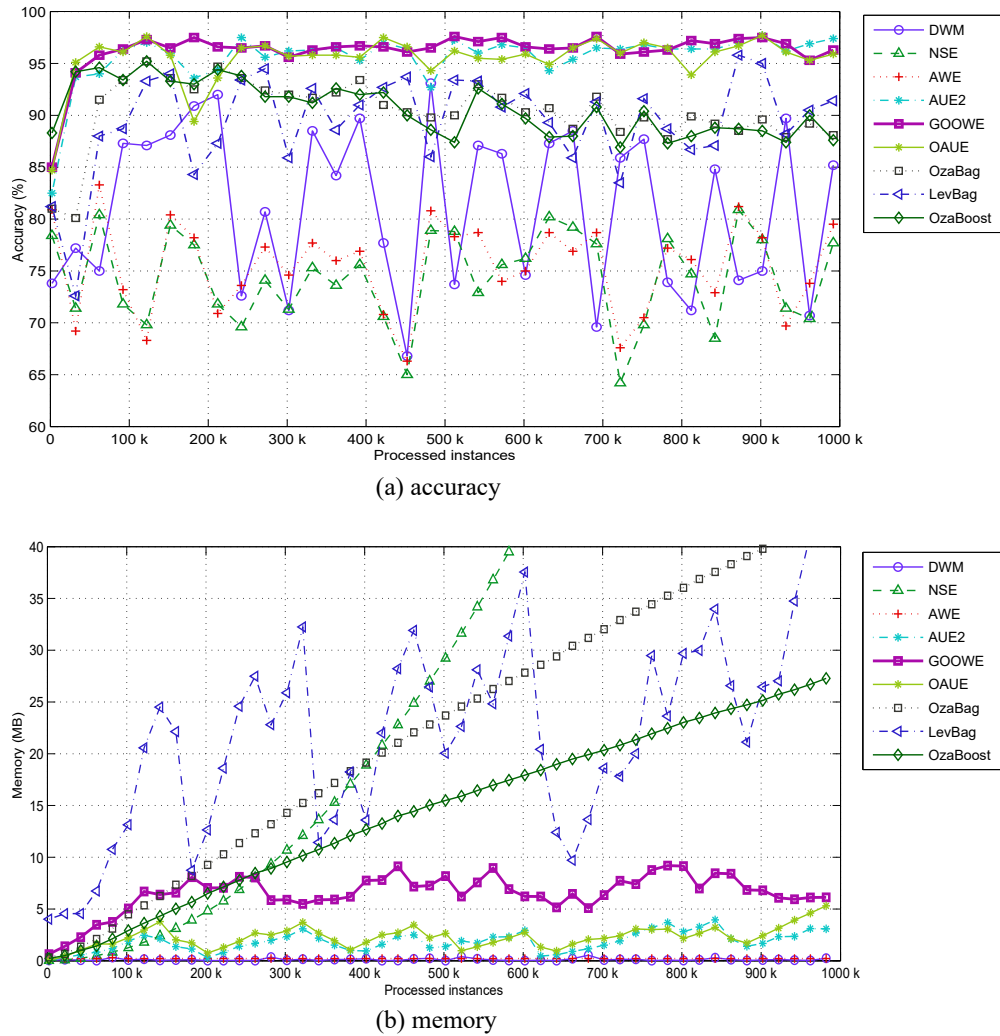


Fig. 7. RCD example with a gradually changing data stream: Classification accuracy and memory consumption for RBF-G-4-F dataset.

derstand the reason in accuracy plots of Fig. 7-(a) and 8-(a). They present behaviors of different ensemble methods with the arrival of new data chunks of gradually/abruptly changing RBF data streams. The place of abrupt drifts is clear in the classification accuracy plots, consistent with what we know from the generation step of these synthetic datasets. In most abruptly changing points, it is obvious that GOOWE has significantly faster adaptive reactions than the others. While OzaBoost, LevBag and OzaBag perform similarly to GOOWE in stationary phases of the data stream, they react slowly in changing phases. As a result, when more changes exist in stream data, GOOWE provides better performance. DWM, NSE and AWE are among the poorly performing algorithms.

Table VII and VIII for the RBF datasets (the first 8 rows) show the conservative resource consumption of GOOWE, in terms of time and memory. We present memory usage behavior for the algorithms in RBF-G-4-F and RBF-A-10-S datasets in Fig.




Fig. 8. RCD example with abruptly changing data stream: Classification accuracy and memory consumption for RBF-A-10-S dataset.

7-(b) and 8-(b). They show that most ensemble methods drop one of the most memory-hungry component classifiers with drift occurrence. Among these algorithms, the memory consumption of GOOWE is less than those of NSE, LevBag, OzaBag, and OzaBoost. Although it uses more memory than DWM, AWE, AUE2, and OAUE, it does not grow exponentially. As Brzezinski explained in [Brzezinski and Stefanowski 2014b], no pruning was used to limit the number of components for NSE, and it requires much more time and memory than the other algorithms. As a result, memory usage of NSE does not react to concept drifts and grows exponentially with the arrival of new instances.

6.2. LCD Data Streams with Miscellaneous Drift Patterns

Table VI for the LCD generators (the second 8 rows) shows that GOOWE is among the top-tier algorithms in the Rotating Hyperplane, TREE-F, and LED datasets, in terms

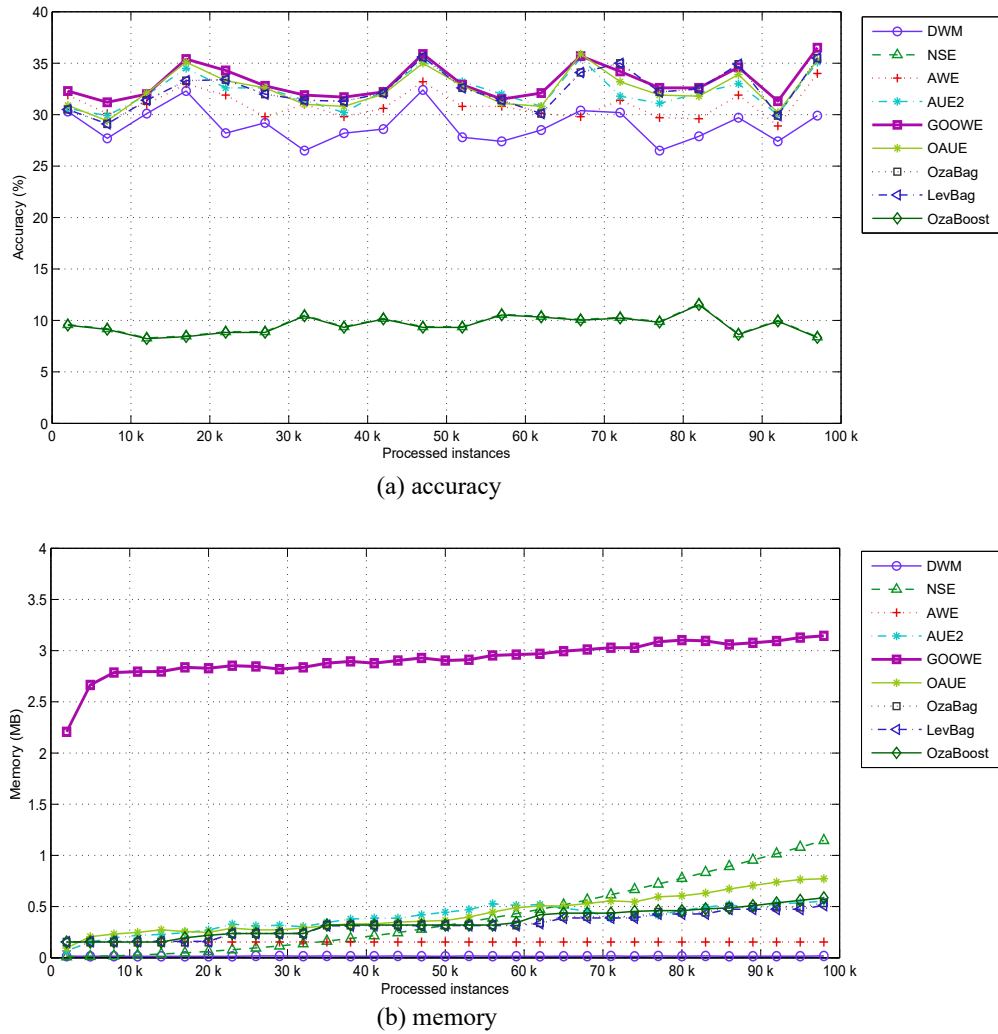


Fig. 9. LCD example with reoccurring data stream: Classification accuracy and memory consumption for TREE-F dataset.

of accuracy. Similarly, to the gradually changing RBF datasets, the Rotating Hyperplane dataset has incremental drifts. TREE-F is the smallest and fastest changing dataset with reoccurring drift patterns. These characteristics show the superiority of GOOWE. Generally, for the LCD datasets, we can say that GOOWE acts better for the fast-changing datasets, compared to the slow ones. For the LED datasets there is no significant difference among various algorithms, and most of them reacting similarly, since there are no clear concept drifts. For the SEA datasets, although GOOWE is not among the top tier algorithms, the differences among accuracy values are small. Moreover, Table VII and VIII for LCD data streams (second 8 rows) show comparable resource consumption of GOOWE, in terms of time and memory, similar to the RCD data streams.

We can see the behavior of the algorithms for the TREE-F dataset, in terms of accuracy (Fig. 9-(a)) and memory usage (Fig. 9-(b)). Similarly to other accuracy plots,




Fig. 10. Real-world example data stream: Classification accuracy and memory consumption for CovPokElec dataset.

GOOWE reacts robustly to concept drifts. The memory usage plot suggests that GOOWE is the worst memory consumer algorithm. However, its memory growth rate is slow and its maximum memory usage is under 4 MB. In other words, it uses a limited amount of memory. In contrast to other plots, in Fig. 9-(b), NSE shows a small consumption of memory. In general, for NSE on small datasets, when only a few components are created, memory usage is reasonable.

6.3. Real-World Data Streams with Unknown Drift Patterns

Table VI, for real-world datasets (the last 4 rows), shows the superiority of GOOWE over other algorithms in PokerHand and CovPokElec datasets, in terms of accuracy. For CoverType and Airlines datasets, although GOOWE is not the best performing algorithm, still the difference with the best performing algorithms are less than 1

percent. In addition, Table VII and VIII for real-world datasets (the last 4 rows) show reasonable resource consumption of GOOWE, in terms of time and memory.

Fig. 10 shows classification accuracy and memory usage behaviors for the CovPok-Elec real-world dataset with the arrival of new data chunks. The accuracy plot (Fig. 10-(a)) shows that GOOWE, OzaBoost, OAUE and DWM are among the best performing ensemble methods. By tracking the behavior of different algorithms, in different situations, it is obvious that there are diverse types of concept drift in the CovPok-Elec dataset. For example, comparing the behavior of the algorithms around 350k and 400k demonstrates that, although all of the algorithms prove the existence of concept drift, for the first evolving point OzaBoost reacts best, in contrast to the second point, where DWM shows the best reaction. By looking at 750k or 1050k points, we can say that, in some situations, different algorithms are not synchronous; while some of them (DWM, OzaBag, LevBag) show a decrease in performance, the others (NSE, AWE, AUE2, GOOWE, OAUE, OzaBoost) show an increase. Considering the first 500k of instances, belonging to the normalized CoverType dataset, DWM and OAUE outperform the others and react faster to unknown drift types. For the second 500k of instances, belonging to the normalized PokerHand dataset, we see more robust behavior from GOOWE and OAUE. However, for the last 500k of instances, belonging to the Electricity dataset, except one evolving point around 1250k, the best performing algorithm is GOOWE.

The memory plot (Fig. 10-(b)) suggests that while AWE, DWM and AUE2 are the least memory consumers, OAUE and GOOWE algorithms are far better than NSE, OzaBag, LevBag and OzaBoost which grow exponentially.

7. STATISTICAL ANALYSIS AND FURTHER DISCUSSION

In order to assure the significant difference of average values for classification accuracies, processing time, and memory usage, we carried out statistical tests. First, a one-way analysis of variance (ANOVA) test using Scheffe multiple comparisons [Scheffe 1959] were conducted on the results of different algorithms for each dataset. The null-hypothesis for each dataset when considered individually is: There is no significant difference between the algorithms.

We conducted the Scheffe test at the significance level of $\alpha = 0.05$. The most accurate, least consumer of processing time, and least consumer of memory algorithms are underlined for each row of Table VI, VII, and VIII. We underlined the top tier group of the Scheffe's comparison results for each synthetic dataset [Scheffe 1959]. As we mentioned earlier, it is not possible to conduct the Scheffe statistical test for real-world datasets, since they only have a single value. For each of them, we underline the most accurate and least resource consuming algorithm. As shown in Table VI, for 15 out of 20 datasets, GOOWE is consistently among the most accurate algorithms. OAUE is placed in the second rank of the most accurate algorithms, with 11 out of 20 datasets. For the cases of time and memory usage, Table VII and VIII, we can see that GOOWE is among the conservative consumers of resources. Despite this fact, when we compare resource usage with the worst ones, we can see that the costs are much less, and are affordable. In addition, comparing the resource usage of OAUE and LevBag with GOOWE shows that our algorithm is in the same range of memory and time consumption.

To extend the analysis, we carried out the non-parametric Friedman statistical test for comparing multiple classifiers over multiple datasets [Demšar 2006; Conover 1999]. The null-hypothesis for this test states that all the algorithms are equivalent on all datasets when considered together. Since we have 9 algorithms and 20 datasets in our experiment, F_F is distributed according to the F distribution with $9 - 1 = 8$ and $(9 - 1) \times (20 - 1) = 152$ degrees of freedom. We run the statistical tests at the significance

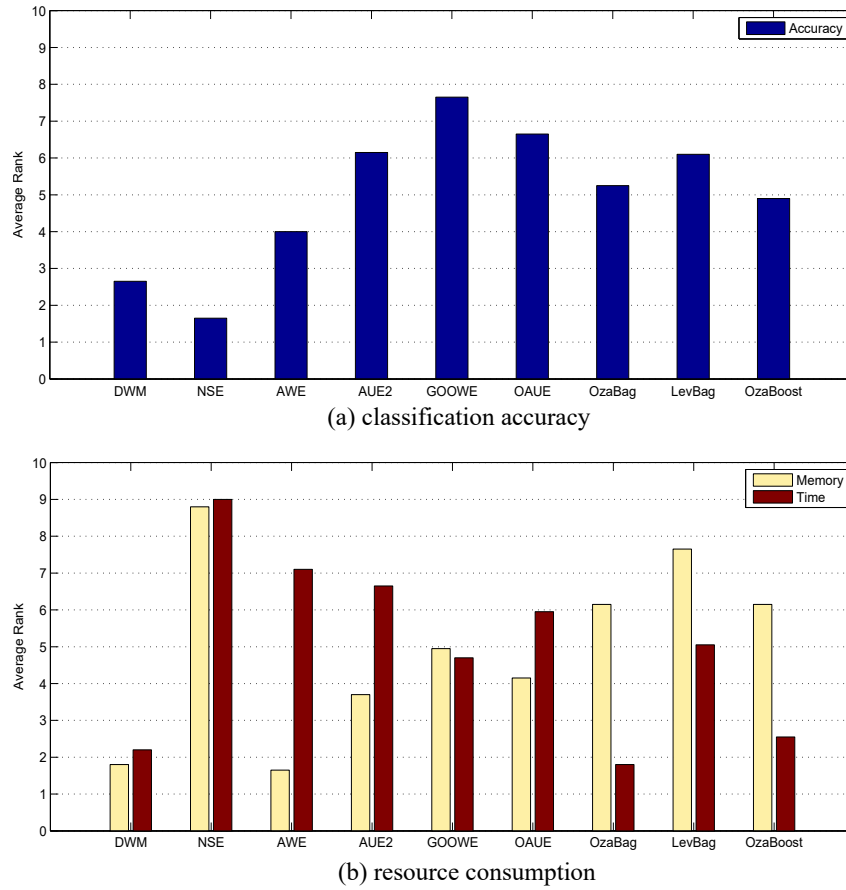


Fig. 11. The Friedman statistical test average rank plots; for classification accuracy plot (a) higher average rank means better prediction, and for resource consumption plot (b) lower average ranks mean better performance.

level of $\alpha = 0.05$; and the Critical Distances (CD) for $F(8, 152)$, and average ranks of algorithms are given in Table IX. If the Friedman test results in a p-value less than α , the null-hypothesis is rejected and we can conclude that at least 2 of the algorithms are significantly different from each other. The tests for accuracy, memory, and time results in p-values less than 0.00001, and the null-hypotheses are rejected for all cases. We plot these average ranks in Fig. 11. Note that for classification accuracy, Fig. 11-(a), higher average rank means better prediction; and for resource consumption, Fig. 11-(b), lower ranks mean better performance.

Table IX shows that, according to the Friedman test, GOOWE outperforms DWM, NSE, AWE, AUE2, OzaBag, LevBag, and OzaBoost, but not OAUE. The CD value is 1.238, and their rank difference is less than this value ($7.650 - 6.650 = 1.000$). Since the difference of average ranks between GOOWE and OAUE is close to the CD, we performed the Wilcoxon signed-rank test to further analyze this pair of algorithms [Demšar 2006]. It ranks the absolute values of the differences between paired samples, and calculates a statistic on the number of negative and positive differences. For our case, the positive differences are 13, and the negative differences are 7. The two-tailed

Table IX. Summary of the Friedman Statistical Test for Accuracy, Memory and Time; The underlined values are GOOWE and its rivals that are in the same range of rank with no significant difference

Test	Test Results		Average Algorithm Ranks								
	F_F	CD	DWM	NSE	AWE	AUE2	GOOWE	OAUE	OzaBag	LevBag	OzaBoost
Accuracy	19.153	1.238	2.650	1.650	4.000	6.150	<u>7.650</u>	<u>6.650</u>	5.250	6.100	4.900
Memory	76.198	0.784	1.800	8.800	1.650	3.700	<u>4.950</u>	<u>4.150</u>	6.150	7.650	6.150
Time	77.692	0.778	2.200	9.000	7.100	6.650	<u>4.700</u>	5.950	1.800	<u>5.050</u>	2.550

probability value, $P = 0.014$, is less than $\alpha = 0.05$; Therefore, it can be accepted that GOOWE is significantly better than OAUE, in terms of accuracy.

Similar to the accuracy test, we performed time and memory statistical comparisons using the Friedman test, summarized in Table IX. For the memory test, we reject the null-hypothesis. The average ranks show that GOOWE uses more memory compared to DWM, AWE, AUE2, and OAUE. On the other hand, it uses less memory compared to NSE, OzaBag, LevBag, and OzaBoost. The CD value shows that, GOOWE is in the middle of the baselines as a singleton, with a significant difference from the upper and lower range algorithms. For the processing time test, we again reject the null-hypothesis; the average ranks show that GOOWE is faster than NSE, AWE, AUE2, and OAUE. It is significantly slower than DWM, OzaBag, and OzaBoost, somehow equivalent to LevBag.

In summary, we can say that there is a trade-off between prediction accuracy and resource consumption. In this trade-off, GOOWE can predict statistically significantly better compared to the most accurate algorithms. Furthermore, it uses affordable resources compared to the most memory and time-efficient ensembles.

8. CONCLUSIONS AND FUTURE WORK

In this paper, we provide a geometrically optimum and online-weighted ensemble classifier, called GOOWE, for non-stationary environments. The main contribution of the proposed algorithm is providing a spatial modeling for using the linear least squares (LSQ) solution for dynamically optimizing the weights of components of an ensemble classifier for evolving environments. Our algorithm, different from the use of LSQ in batch mode ensembles, has dynamically changing optimum weight assignment to component classifiers and continuous training and testing. We use data chunks for training and a sliding instance window containing the latest available data for testing; such an approach provides more robust behavior as shown in our experiments. We use the Euclidean norm as the measure of closeness in LSQ. The LSQ proved to react well in noisy situations [Hansen et al. 2013], as it did in our algorithm for data streams with concept drifts.

First, we conduct an analysis for examining two major differentiating elements of GOOWE, the component weighting strategy and ensemble model management strategy. Our optimum and online weighting system shows its effectiveness for both vote aggregation and ensemble maintenance in evolving environments. Second, we experimentally compare GOOWE with 8 state-of-the-art ensemble classifiers, as our baselines, on 20 datasets as data streams with tens of millions of instances, where 16 of them are synthetic and 4 of them are real-world datasets. For the synthetic streams, we use two categories of data stream generators: Rigorous Concept Drift (RCD) and Loose Concept Drift (LCD), each with 8 datasets. We include all possible patterns of change (sudden/abrupt, incremental, gradual, and reoccurring as concept drifts, including blips and noise) in our datasets. The statistical tests prove the superiority and robustness of GOOWE in reacting to different types of concept drift, in terms of ac-

curacy. Furthermore, we show that it requires conservative resource consumption, in terms of memory and processing time.

In future work, the impact of the length of data chunk size on the performance in the presence of different concept drifts and its dynamic determination are possible studies. Effects of instance window size on performance can also be analyzed. In addition, GOOWE can be used in various sub-problem domains, such as semi-supervised and multi-label classification. It can be applied to unbalanced datasets and streams with concept-evolution.

ACKNOWLEDGMENTS

We would like to thank Manouchehr Takrimi from Bilkent University, Jon M. Patton from Miami University of OH, and Alper Can for their valuable comments and pointers on this paper.

REFERENCES

- Albert Bifet and Ricard Gavaldà. 2007. Learning from time-changing data with adaptive windowing. In *Proceedings of the Seventh SIAM International Conference on Data Mining, April 26-28, 2007, Minneapolis, Minnesota, USA*. 443–448.
- Albert Bifet, Geoff Holmes, Richard Kirkby, and Bernhard Pfahringer. 2010b. MOA: Massive Online Analysis. *Journal of Machine Learning Research* 11 (2010), 1601–1604.
- Albert Bifet, Geoffrey Holmes, and Bernhard Pfahringer. 2010a. Leveraging bagging for evolving data streams. In *Machine Learning and Knowledge Discovery in Databases, European Conference, ECML PKDD 2010, Barcelona, Spain, September 20-24, 2010, Proceedings, Part I*. 135–150.
- Albert Bifet, Geoffrey Holmes, Bernhard Pfahringer, Richard Kirkby, and Ricard Gavaldà. 2009. New ensemble methods for evolving data streams. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Paris, France, June 28 - July 1, 2009*. 139–148.
- Hamed R. Bonab and Fazli Can. 2016. A theoretical framework on the ideal number of classifiers for online ensembles in data streams. In *Proceedings of the 25th ACM International Conference on Information and Knowledge Management, CIKM 2016, Indianapolis, IN, USA, October 24-28, 2016*. 2053–2056.
- Dariusz Brzezinski and Jerzy Stefanowski. 2014a. Combining block-based and online methods in learning ensembles from concept drifting data streams. *Information Sciences* 265 (2014), 50–67.
- Dariusz Brzezinski and Jerzy Stefanowski. 2014b. Reacting to different types of concept drift: The accuracy updated ensemble algorithm. *IEEE Transactions on Neural Networks and Learning Systems* 25, 1 (2014), 81–94.
- Lai-Wan Chan. 1999. Weighted least square ensemble networks. In *International Joint Conference on Neural Networks (IJCNN 1999)*, Vol. 2. IEEE, 1393–1396.
- W. Conover. 1999. *Practical Nonparametric Statistics*. John Wiley & Sons, New York.
- Janez Demšar. 2006. Statistical comparisons of classifiers over multiple data sets. *Journal of Machine Learning Research* 7 (2006), 1–30.
- Thomas G. Dietterich and Ghulum Bakiri. 1995. Solving multiclass learning problems via error-correcting output codes. *Journal of Artificial Intelligence Research (JAIR)* 2 (1995), 263–286.
- Gregory Ditzler and Robi Polikar. 2013. Incremental learning of concept drift from streaming imbalanced data. *IEEE Transactions on Knowledge and Data Engineering* 25, 10 (2013), 2283–2301.
- Pedro M. Domingos and Geoff Hulten. 2000. Mining high-speed data streams. In *Sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Boston, MA, USA, August 20-23, 2000*. 71–80.
- Ryan Elwell and Robi Polikar. 2011. Incremental learning of concept drift in nonstationary environments. *IEEE Transactions on Neural Networks and Learning Systems* 22, 10 (2011), 1517–1531.
- Dewan Md Farid, Li Zhang, Alamgir Hossain, Chowdhury Mofizur Rahman, Rebecca Strachan, Graham Sexton, and Keshav Dahal. 2013. An adaptive ensemble classifier for mining concept drifting data streams. *Expert Systems with Applications* 40, 15 (2013), 5895–5906.
- Yoav Freund and Robert E. Schapire. 1997. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. System Sci.* 55, 1 (1997), 119–139.
- Jerome H Friedman. 2002. Stochastic gradient boosting. *Computational Statistics & Data Analysis* 38, 4 (2002), 367–378.
- Joao Gama. 2010. *Knowledge discovery from data streams*. CRC Press.

- João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. 2013. On evaluating stream learning algorithms. *Machine Learning* 90, 3 (2013), 317–346.
- João Gama, Indre Zliobaite, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. 2014. A survey on concept drift adaptation. *Comput. Surveys* 46, 4 (2014), 44:1–44:37.
- Jing Gao, Wei Fan, and Jiawei Han. 2007. On appropriate assumptions to mine data streams: Analysis and practice. In *Seventh IEEE International Conference on Data Mining (ICDM 2007)*. IEEE, 143–152.
- Heitor Murilo Gomes, Jean Paul Barddal, Fabrício Enembreck, and Albert Bifet. 2017. A survey on ensemble learning for data stream classification. *Comput. Surveys* 50, 2, Article 23 (March 2017), 36 pages.
- Dong-Hong Han, Xin Zhang, and Guo-Ren Wang. 2015. Classifying uncertain and evolving data streams with distributed extreme learning machine. *Journal of Computer Science and Technology* 30, 4 (2015), 874–887.
- Per Christian Hansen, Víctor Pereyra, and Godela Scherer. 2013. *Least squares data fitting with applications*. Johns Hopkins University Press, Baltimore, Md.
- Daniel Hernández-Lobato, Gonzalo MartíNez-Muñoz, and Alberto Suárez. 2013. How large should ensembles of classifiers be? *Pattern Recognition* 46, 5 (2013), 1323–1336.
- Geoff Hulten, Laurie Spencer, and Pedro M. Domingos. 2001. Mining time-changing data streams. In *Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 26-29, 2001*. 97–106.
- Jeremy Z. Kolter and Marcus A. Maloof. 2003. Dynamic weighted majority: A new ensemble method for tracking concept drift. In *Third IEEE International Conference on Data Mining (ICDM 2003)*. IEEE, 123–130.
- Jeremy Z. Kolter and Marcus A. Maloof. 2005. Using additive expert ensembles to cope with concept drift. In *Proceedings of 22nd International Conference on Machine Learning, (ICML 2005), Bonn, Germany, August 7-11, 2005*. 449–456.
- J Zico Kolter and Marcus A Maloof. 2007. Dynamic weighted majority: An ensemble method for drifting concepts. *Journal of Machine Learning Research* 8 (2007), 2755–2790.
- Bartosz Krawczyk, Leandro L Minku, João Gama, Jerzy Stefanowski, and Michał Woźniak. 2017. Ensemble learning for data stream analysis: A survey. *Information Fusion* 37 (2017), 132–156.
- Ludmila I. Kuncheva. 2004. Classifier ensembles for changing environments. In *Proceedings of the 5th International Workshop on Multiple Classifier Systems, (MCS 2004), Cagliari, Italy, June 9-11, 2004*. 1–15.
- Ludmila I Kuncheva. 2008. Classifier ensembles for detecting concept change in streaming data: Overview and perspectives. In *2nd Workshop SUEMA*, Vol. 2008. 5–10.
- Patrice Latinne, Olivier Debeir, and Christine Decaestecker. 2001. Limiting the number of trees in random forests. In *International Workshop on Multiple Classifier Systems*. Springer, 178–187.
- Nick Littlestone. 1987. Learning quickly when irrelevant attributes abound: A new linear-threshold algorithm. *Machine Learning* 2, 4 (1987), 285–318.
- Nick Littlestone and Manfred K. Warmuth. 1994. The weighted majority algorithm. *Information and Computation* 108, 2 (1994), 212–261.
- Mohammad M. Masud, Jing Gao, Latifur Khan, Jiawei Han, and Bhavani M. Thuraisingham. 2011. Classification and novel class detection in concept-drifting data streams under time constraints. *IEEE Transactions on Knowledge and Data Engineering* 23, 6 (2011), 859–874.
- Leandro L. Minku, Allan P. White, and Xin Yao. 2010. The impact of diversity on online ensemble learning in the presence of concept drift. *IEEE Transactions on Knowledge and Data Engineering* 22, 5 (2010), 730–742.
- Leandro L. Minku and Xin Yao. 2012. DDD: A new ensemble approach for dealing with concept drift. *IEEE Transactions on Knowledge and Data Engineering* 24, 4 (2012), 619–633.
- Ahmad Mustafa, Ahsanul Haque, Latifur Khan, Michael Baron, and Bhavani Thuraisingham. 2014. Evolving stream classification using change detection. In *International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom 2014)*. IEEE, 154–162.
- Kyosuke Nishida, Koichiro Yamauchi, and Takashi Omori. 2005. ACE: Adaptive classifiers-ensemble system for concept-drifting environments. In *Multiple Classifier Systems*. Springer, 176–185.
- Thais Mayumi Oshiro, Pedro Santoro Perez, and José Augusto Baranauskas. 2012. How many trees in a random forest?. In *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, 154–168.
- Nikunj C. Oza. 2001. *Online Ensemble Learning*. Ph.D. Dissertation. Computer Science Division, Univ. California, Berkeley, CA, USA.

- Nikunj C Oza and Stuart Russell. 2001. Experimental comparisons of online and batch versions of bagging and boosting. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 26-29, 2001*. ACM, 359–364.
- Donald G. Saari. 2008. Complexity and the geometry of voting. *Mathematical and Computer Modelling* 48, 9-10 (2008), 1335–1356.
- Henry Scheffe. 1959. *The Analysis of Variance*. John Wiley, New York.
- W. Nick Street and YongSeog Kim. 2001. A streaming ensemble algorithm (SEA) for large-scale classification. In *Proceedings of the Seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 26-29, 2001*. 377–382.
- Yu Sun, Ke Tang, Leandro L Minku, Shuo Wang, and Xin Yao. 2016. Online ensemble learning of data streams with gradually evolved classes. *IEEE Transactions on Knowledge and Data Engineering* 28, 6 (2016), 1532–1545.
- Kagan Tumer and Joydeep Ghosh. 1996. Analysis of decision boundaries in linearly combined neural classifiers. *Pattern Recognition* 29, 2 (1996), 341–348.
- Haixun Wang, Wei Fan, Philip S. Yu, and Jiawei Han. 2003. Mining concept-drifting data streams using ensemble classifiers. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Washington, DC, USA, August 24 - 27, 2003*. 226–235.
- Shuo Wang, Leandro L Minku, and Xin Yao. 2015. Resampling-based ensemble methods for online class imbalance learning. *IEEE Transactions on Knowledge and Data Engineering* 27, 5 (2015), 1356–1368.
- Mohammadzaman Zamani, Hamid Beigy, and Amirreza Shaban. 2016. Cascading randomized weighted majority: a new online ensemble learning algorithm. *Intelligent Data Analysis* 20, 4 (2016), 877–889.
- Peng Zhang, Xingquan Zhu, and Yong Shi. 2008. Categorizing and mining concept drifting data streams. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*. 812–820.
- Xingquan Zhu, Peng Zhang, Xiaodong Lin, and Yong Shi. 2010. Active learning from stream data using optimal weight classifier ensemble. *IEEE Transactions on Systems, Man, and Cybernetics, Part B* 40, 6 (2010), 1607–1621.
- Indre Zliobaite. 2013. How good is the electricity benchmark for evaluating concept drift adaptation. *CoRR* abs/1301.3524 (2013).

Received July 2016; revised May 2017; accepted August 2017