

Ensembles of Restricted Hoeffding Trees ¹

ALBERT BIFET, Department of Computer Science, University of Waikato

EIBE FRANK, Department of Computer Science, University of Waikato

GEOFF HOLMES, Department of Computer Science, University of Waikato

BERNHARD PFAHRINGER, Department of Computer Science, University of Waikato

The success of simple methods for classification shows that it is often not necessary to model complex attribute interactions to obtain good classification accuracy on practical problems. In this paper, we propose to exploit this phenomenon in the data stream context by building an ensemble of Hoeffding trees that are each limited to a small subset of attributes. In this way, each tree is restricted to model interactions between attributes in its corresponding subset. Because it is not known *a priori* which attribute subsets are relevant for prediction, we build exhaustive ensembles that consider all possible attribute subsets of a given size. As the resulting Hoeffding trees are not all equally important, we weigh them in a suitable manner to obtain accurate classifications. This is done by combining the log-odds of their probability estimates using sigmoid perceptrons, with one perceptron per class. We propose a mechanism for setting the perceptrons' learning rate using the ADWIN change detection method for data streams, and also use ADWIN to reset ensemble members (i.e. Hoeffding trees) when they no longer perform well. Our experiments show that the resulting ensemble classifier outperforms bagging for data streams in terms of accuracy when both are used in conjunction with adaptive naive Bayes Hoeffding trees, at the expense of runtime and memory consumption. We also show that our stacking method can improve the performance of a bagged ensemble.

Categories and Subject Descriptors: H.2.8 [Database applications]: Data Mining

General Terms: Design, Algorithms, Performance

Additional Key Words and Phrases: Data Streams, Decision Trees, Ensemble Methods

ACM Reference Format:

Bifet, A., Eibe, F., Holmes, G., and Pfahringer, B. 2011. Ensembles of Restricted Hoeffding Trees. ACM Trans. Intell. Syst. Technol. V, N, Article A (January YYYY), 20 pages.

DOI = 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

1. INTRODUCTION

When applying boosting algorithms to build ensembles of decision trees using a standard tree induction algorithm, the tree inducer has access to all the attributes in the data. Thus, each tree can model arbitrarily complex interactions between attributes in principle. However, often the full modeling power of unrestricted decision tree learners is not necessary to yield good accuracy, and it may even be harmful because it can lead to overfitting. Hence, it is common to apply boosting in conjunction with depth-limited decision trees, where the growth of each tree is restricted to a certain level. In this way, only a limited set of attributes can be used in each tree, but the risk of overfitting is

¹This is an extended version of a paper that appeared in the 2nd Asian Conference on Machine Learning.

Author's addresses: Albert Bifet, Eibe Frank, Geoff Holmes and Bernhard Pfahringer, Computer Science Department, University of Waikato, Hamilton, New Zealand.

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© YYYY ACM 0000-0003/YYYY/01-ARTA \$10.00

DOI 10.1145/0000000.0000000 <http://doi.acm.org/10.1145/0000000.0000000>

reduced. In the extreme case, only a single split is allowed and the decision trees degenerate to decision stumps. Despite the fact that each restricted tree can only model interactions between a limited set of attributes, the overall ensemble is often highly accurate.

The method presented in this paper is based on this observation. We present an algorithm that produces a classification model based on an ensemble of restricted decision trees, where each tree is built from a distinct subset of the attributes. The overall model is formed by combining the log-odds of the predicted class probabilities of these trees using sigmoid perceptrons, with one perceptron per class. In contrast to boosting, which forms an ensemble classifier in a greedy fashion, building each tree in sequence and assigning corresponding weights as a by-product, our method generates each tree in parallel and combines them using perceptron classifiers by adopting the stacking approach [Wolpert 1992]. Because we are working in a data stream scenario, Hoeffding trees [Domingos and Hulten 2000] are used as the ensemble members. They, as well as the perceptrons, can be trained incrementally, and we also show how ADWIN-based change detection [Bifet and Gavaldà 2007] can be used to apply the method to evolving data streams.

There is existing work on boosting for data streams [Oza and Russell 2001b], but the algorithm has been found to yield inferior accuracy compared to bagging [Bifet et al. 2009; Bifet et al. 2009]. Moreover, it is unclear how a boosting algorithm can be adapted to data streams that evolve over time: because bagging generates models independently, a model can be replaced when it is no longer accurate, but the sequential nature of boosting prohibits this simple and elegant solution. Because our method generates models independently, we can apply this simple strategy, and we show that it yields more accurate classifications than Online Bagging [Oza and Russell 2001b] on the real-world and artificial data streams that we consider in our experiments.

Our method is also related to work on building ensembles using random subspace models [Ho 1998]. In the random subspace approach, each model in an ensemble is trained based on a randomly chosen subset of attributes. The models' predictions are then combined in an unweighted fashion. Because of this latter property, the number of attributes available to each model must necessarily be quite large in general, so that each individual model is powerful enough to yield accurate classifications. In contrast, in our method, we exhaustively consider all subsets of a given, small, size, build a tree for each subset of attributes, and then combine their predictions using stacking. In random forests [Breiman 2001], the random subspace approach is applied locally at each node of a decision tree in a bagged ensemble. Random forests have been applied to data streams, but did not yield substantial improvements on bagging in this scenario [Bifet et al. 2010]. Another related weighted ensemble method is Dynamic Weighted Majority [Kolter and Maloof 2007].

The paper is structured as follows. The next section presents the new ensemble learner we propose. Section 3 details the set-up for our experiments. Section 4 presents the experimental comparison to bagging. In Section 5 we consider ensemble pruning and in Section 6 we evaluate the effect of using stacking rather than simple averaging to combine the trees' predictions, and also consider the effect of different ADWIN-based classifier replacement strategies. Section 7 concludes the paper.

2. COMBINING RESTRICTED Hoeffding TREES USING STACKING

The basic method we present in this paper is very simple: enumerate all attribute subsets of a given user-specified size k , learn a Hoeffding tree from each subset based on the incoming data stream, gather the trees' predictions for each incoming instance, and use these predictions to train simple perceptron classifiers. The question that needs to

be addressed is how exactly to prepare the “meta” level data for the perceptrons and what shape it should take.

Rather than using discrete classifications to build the meta level model in stacking, it is common to use class probability estimates instead because they provide more information due to the fact that they represent the degree of confidence that each model has in its predictions. Hence, we also adopt this approach in our method: the meta-level data is formed by collecting the class probability estimates for an incoming new instance, obtained from the Hoeffding trees built from the data observed previously.

The meta level combiner we use is based on simple perceptrons with sigmoid activation functions, trained using stochastic gradient descent to minimize the squared loss with respect to the actual observed class labels in the data stream. We train one perceptron per class value and use the Hoeffding trees’ class probability estimates for the corresponding class value to form the input data for each perceptron.

There is one caveat. Because the sigmoid activation function is used in the perceptrons, we do not use the raw class probability estimates as the input values for training them. Rather, we use the log-odds of these probability estimates instead. Let $\hat{p}(c_{ij}|\vec{x})$ be the probability estimate for class i and instance \vec{x} obtained from Hoeffding tree j in the ensemble. Then we use

$$a_{ij} = \log(\hat{p}(c_{ij}|\vec{x})/(1 - \hat{p}(c_{ij}|\vec{x})))$$

as the value of input attribute j for the perceptron associated with class value i .

Let \vec{a}_i be the vector of log-odds for class i . The output of the sigmoid perceptron for class i is then $f(\vec{a}_i) = 1/(1 + e^{-(\vec{w}_i\vec{a}_i+b_i)})$, based on per-class coefficients \vec{w}_i and b_i . We use the log-odds as the inputs for the perceptron because application of the sigmoid function presupposes a linear relationship between $\log(f(\vec{a}_i)/(1 - f(\vec{a}_i)))$ and \vec{a}_i .

To avoid the zero-frequency problem, we slightly modify the probability estimates obtained from a Hoeffding tree by adding a small constant ϵ to the probability for each class, and then renormalize. In our experiments, we use $\epsilon = 0.001$, but smaller values lead to very similar results.

The perceptrons are trained using stochastic gradient descent: the weight vector is updated each time a new training instance is obtained from the data stream. Once the class probability estimates for that instance have been obtained from the ensemble of Hoeffding trees, the input data for the perceptrons can be formed, and the gradient descent update rule can be used to perform the update. The weight values are initialized to the reciprocal of the size of the ensemble, so that the perceptrons give equal weight to each ensemble member initially.

A crucial aspect of stochastic gradient descent is an appropriately chosen learning rate, which determines the magnitude of the update. If it is chosen too large, there is a risk that the learning process will not converge. A common strategy is to decrease it as the amount of training data increases. Let n be the number of training instances seen so far in the data stream, and let m be the number of attributes. We set the learning rate α based on the following equation:

$$\alpha = \frac{2}{2 + m + n}.$$

However, there is a problem with this approach for setting the learning rate in the context we consider here: it assumes that the training data is identically distributed. This is not actually the case in our scenario because the training data for the perceptrons is derived from the probability estimates obtained from the Hoeffding trees, and these change over time, generally becoming more accurate. Setting the learning rate based on the above equation means that the perceptrons will adapt too slowly once the initial data in the data stream has been processed.

There is a solution to this problem: the stream of predictions from the Hoeffding trees can be viewed as an *evolving* data stream (regardless of whether the underlying data stream forming the training data for the Hoeffding trees is actually evolving) and we can use an existing change detection method for evolving data streams to detect when the learning rate should be reset to a larger value. We do this very simply by setting the value of n to zero when change has been detected. To detect change, we use the ADWIN change detector [Bifet and Gavaldà 2007], which is discussed in more detail in the next section. It detects when the accuracy of a classifier increases or decreases significantly as a data stream is processed. We apply it to monitor the accuracy of each Hoeffding tree in the ensemble. When accuracy changes significantly for one of the trees, the learning rate is reset by setting n to zero. The value of n is then incremented for each new instance in the stream until a new change is detected. This has the effect that the learning rate will be kept relatively large while the learning curve for the Hoeffding trees has a significant upward trend.

2.1. ADWIN-based Change Detection

The ADWIN change detector comes with theoretical guarantees on the ratio of false positives and the scale of change. In addition to using it to reset the learning rate when necessary, we also use it to make our ensemble classifier applicable to an evolving data stream, where the original data stream, used to train the Hoeffding trees, changes over time.

ADWIN maintains a window of observations and automatically detects and adapts to the current rate of change. Its only parameter is a confidence bound δ , indicating the desired confidence in the algorithm's output, inherent to all algorithms dealing with random processes. ADWIN does not maintain the window explicitly, but compresses it using a variant of the exponential histogram technique. Consequently it keeps a window of length W using only $O(\log W)$ memory and $O(\log W)$ processing time per item.

The strategy we use to cope with evolving data streams using ADWIN is based on the approach that has been used to make Online Bagging applicable to an evolving data stream in [Bifet et al. 2009]. The idea is to replace ensemble members when they start to perform poorly. To implement this, we use ADWIN to detect when the accuracy of one of the Hoeffding trees in the ensemble has dropped significantly. To do this, we can use the same ADWIN change detectors that are also applied to detect when the learning rate needs to be reset. When one of the change detectors associated with a particular tree reports a significant drop in accuracy, the tree is reset and the coefficients in the perceptrons that are associated with this tree (one per class value) are set to zero. A new tree is then generated from new data in the data stream, so that the ensemble can adapt to changes.

Note that all trees for which a significant drop is detected are replaced and that a change detection event automatically triggers a reset of the learning rate, which is important because the perceptrons need to adapt to the changed ensemble of trees.

2.2. A Note on Computational Complexity

Our approach is based on generating trees for all possible attribute subsets of size k . If there are m attributes in total, there are $\binom{m}{k}$ of these subsets. Clearly, only moderate values of k , or values of k that are very close to m , are feasible. When k is one, there is no penalty with respect to computational complexity compared to building a single Hoeffding tree, because then there is only one tree per attribute, and an unrestricted tree also scales linearly in the number of attributes. If k is two, there is an extra factor of m in the computational complexity compared to building a single unrestricted Hoeffding tree, i.e. overall effort becomes quadratic in the number of attributes.

$k = 2$ is very practical even for datasets with a relatively large number of attributes, although certainly not for very high-dimensional data (for which linear classifiers are usually sufficient anyway). Larger values of k are only practical for small numbers of attributes, unless k is very close to m (e.g. $k = m - 1$). We have used $k = 4$ for datasets with 10 attributes in our experiments, with very acceptable runtimes. It is important to keep in mind that many practical classification problems appear to exhibit only very low-dimensional interactions, which means small values of k are sufficient to yield high accuracy.

3. EXPERIMENT SETUP

We have performed several experiments to investigate the performance and resources needed by our new method, and to compare to bagged Hoeffding trees. To measure use of resources, we use time and memory, and also RAM-Hours, a new evaluation measure introduced in [Bifet et al. 2010], where every RAM-Hour equals a GB of RAM deployed for 1 hour.

We performed our experiments using the **Massive Online Analysis (MOA)** framework [Bifet et al. 2010], a data mining software environment for implementing algorithms and running experiments for evolving data streams. We implemented all algorithms in Java within the MOA framework. We used the default parameters of the classifiers in MOA. The experiments were performed on 2.66 GHz Core 2 Duo E6750 machines with 4 GB of memory.

In the rest of this section we briefly describe the datasets and methods used (excluding our new method, which was described in the previous section).

3.1. Real-World Data

We consider two types of data streams, synthetic ones and ones corresponding to real-world problems. To evaluate our method on real-world data, we use three different datasets: two large datasets from the UCI repository of machine learning databases [Frank and Asuncion 2010], namely Forest Covertype and Poker-Hand, and one other large dataset that has previously been used to study data stream methods, namely the Electricity dataset [Harries 1999; Gama et al. 2004]. In our experiments with these datasets all numeric attributes have been normalized to the $[0, 1]$ range.

Forest Covertype. This data contains the forest cover type for 30 x 30 meter cells obtained from US Forest Service (USFS) Region 2 Resource Information System (RIS) data. The data has 581,012 instances and 54 attributes. There are 7 classes. The data is often used to evaluate methods for data stream classification, e.g. in [Gama et al. 2003; Oza and Russell 2001a].

Poker-Hand. This data contains 1,000,000 instances. Each instance in this dataset is an example of a hand consisting of five playing cards drawn from a standard deck of 52 cards. There are 10 predictor attributes in this data because each card is described using two attributes (suit and rank). The class attribute has 10 possible values and describes the poker hand, e.g. full house. The cards are not ordered in the data and a hand can be represented by any permutation, which makes it very difficult to learn for propositional learners such as the ones that are commonly used for data stream classification. Hence, we use a modified version, in which cards are sorted by rank and suit and duplicates have been removed [Bifet et al. 2009]. The resulting dataset contains 829,201 instances.

Electricity. The electricity data is described in [Harries 1999] and has been used for data stream classification in [Gama et al. 2004]. It originates from the New

South Wales Electricity Market in Australia. In this market, prices are not fixed but set every five minutes, as determined by supply and demand. The data comprises 45,312 instances and 8 predictor attributes. The class label is determined based on the change of price relative to a moving average of the last 24 hours, yielding 2 class values.

In the data stream scenario, it is important to consider evaluation under concept drift. We can use the method from [Bifet et al. 2009] in conjunction with the above three datasets to create a data stream exhibiting concept drift. To this end, we assume the underlying distributions are fixed and then combine these "pure" distributions by modeling concept drift events as weighted combinations of two distributions. As originally suggested in [Bifet et al. 2009], the sigmoid function is used to define the probability that an instance is sampled from the "old" stream—the one in place before the concept drift event—or the "new" stream. In this way, two distributions are mixed using a soft threshold function. The exact specification of the operator used to combine two data streams is given in the following definition from [Bifet et al. 2009].

Definition 3.1. Given two data streams a, b , we define $c = a \oplus_{t_0}^W b$ as the data stream built by joining the two data streams a and b , where t_0 is the point of change, W is the length of change, $\Pr[c(t) = b(t)] = 1/(1 + e^{-4(t-t_0)/W})$ and $\Pr[c(t) = a(t)] = 1 - \Pr[c(t) = b(t)]$.

This operator can be applied repeatedly to introduce multiple concept change events, thus making it possible to combine multiple data streams into a single evolving data stream. Different parameters can be used for each change event. For example, the following expression corresponds to the combination of four data streams a, b, c , and d , joined with three parametrized change events: $((a \oplus_{t_0}^{W_0} b) \oplus_{t_1}^{W_1} c) \oplus_{t_2}^{W_2} d) \dots$

We use this operator to join the above three real-world datasets into a single data stream that exhibits concept drift. More specifically, we define a new data stream

$$\text{CovPokElec} = (\text{CoverType} \oplus_{581,012}^{5,000} \text{Poker-Hand}) \oplus_{829,201}^{5,000} \text{ELECTRICITY}$$

Note that, to perform the operation, we concatenate the lists of attributes from the three datasets, and the number of classes is set to the maximum number of classes amongst the three datasets.

3.2. Synthetic Data Streams

There is a shortage of publicly available large real-world datasets that are suitable for the evaluation of data stream methods. Thus we also consider synthetic data in our experiments. To avoid bias, we use synthetic data generators that are commonly found in the literature. Most of these data generators exhibit a mechanism that implements concept drift. For each generator, we use 1 million instances in our experiments.

Rotating Hyperplane. This data generator, introduced in [Hulten et al. 2001], provides an elegant way of generating time-changing data for the evaluation of data stream classification. The data is generated based on hyperplanes whose orientation and position is varied smoothly over time. Classes are assigned to data points based on which side of a hyperplane they are located. The generator has a parameter that controls the speed of change. We use 10 predictor attributes for this data and 2 class values.

Random RBF Generator. This data generator, from [Kirkby 2007], is also able to produce time-changing data. Like the hyperplane generator, it produces data that is difficult to model accurately using decision tree learners. It first generates a

fixed number of radial basis functions with randomly chosen centroids and assigns weights and class labels to them. Data is then generated based on these basis functions, taking their relative weights into account. Concept drift is introduced by moving the centroids with constant speed, as controlled by a parameter. We use several variants of this data, varying the numbers of centroids and the speed of change, but keeping the number of predictor attributes constant at 10 and the number of class values fixed at 5.

LED Generator. This generator is one of the oldest data generators that can be found in the literature. It was introduced in [Breiman et al. 1984] and a C implementation can be found in the UCI repository of machine learning databases [Frank and Asuncion 2010]. The task is to classify the 10 digits based on boolean attributes that correspond to the LEDs of a seven-segment LED display. Each attribute value has a 10% chance of being inverted, so the data includes noise. The variant of the data generator that we use additionally also generates 17 completely irrelevant attributes. Concept drift is controlled by a drift parameter that determines the number of attributes with drift.

Waveform Generator. This source of synthetic data was also introduced in the CART book [Breiman et al. 1984] and a C implementation is available from the UCI repository. The data has three classes that correspond to three different waveforms. Each waveform is based on a combination of two or three base waves. As in the LED generator, there is a drift parameter determining the number of attributes with drift. The data has 21 attributes.

Random Tree Generator. This generator, which is the only generator we use that does not implement concept drift, first constructs a decision tree by randomly choosing attributes to define splits for the internal nodes of the tree. It then assigns randomly chosen class labels to the leaf nodes of this tree. To generate data, it routes uniformly distributed synthetic instances down the appropriate paths in the tree and assigns class labels based on the corresponding leaf nodes. This generator was introduced in [Domingos and Hulten 2000] and favours decision tree learners. We use it with 10 attributes and 5 class values.

3.3. Hoeffding Naive Bayes Trees and ADWIN Bagging

Hoeffding trees [Domingos and Hulten 2000] are state-of-the-art tree inducers in classification for data streams. They exploit the fact that a small sample can often be enough to choose an optimal splitting attribute. This idea is supported mathematically by the Hoeffding bound, which quantifies the number of observations (in our case, examples) needed to estimate some statistics within a prescribed precision (in our case, the goodness of an attribute). Using the Hoeffding bound one can show that the inducer's output is asymptotically nearly identical to that of a non-incremental learner using infinitely many examples.

Hoeffding trees perform prediction by choosing the majority class at each leaf. Their predictive accuracy can be increased by adding naive Bayes models at the leaves of the trees. However, [Holmes et al. 2005] identified situations where the naive Bayes method outperforms the standard Hoeffding tree initially but is eventually overtaken. They proposed a *Hoeffding Naive Bayes Tree (hnbt)*, a hybrid adaptive method that generally outperforms the two original prediction methods for both simple and complex concepts. This method works by performing a naive Bayes prediction per training instance, comparing its prediction with the majority class. Counts are stored to measure how many times the naive Bayes prediction gets the true class correct as com-

Table I. Performance for a Hoeffding Naive Bayes Tree (**hnbt**). Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB.

	hnbt		
	Time	Acc.	Mem.
COVTYPE	25.74	80.70	2.65
ELECTRICITY	0.86	79.20	0.09
POKER-HAND	10.10	77.10	0.59
COVPOKELEC	68.25	78.72	7.54
RBF(50,0)	17.16	83.46 ± 0.74	0.55
RBF(50,0.001)	18.06	32.31 ± 0.09	0.42
RBF(10,0.001)	14.68	76.44 ± 0.21	0.52
RBF(50,0.0001)	19.15	45.14 ± 0.59	0.61
RBF(10,0.0001)	15.93	79.17 ± 0.51	0.61
RT	11.64	89.24 ± 0.25	1.68
LED(50000)	18.51	68.53 ± 0.26	1.95
Waveform(50)	19.27	78.24 ± 6.52	1.00
HYP(10,0.001)	10.33	83.22 ± 2.27	1.45
HYP(10,0.0001)	9.50	89.00 ± 0.41	1.26
74.32 Acc.			
0.00 RAM-Hours			

Table II. Comparison using sets of 1 or 2 attributes. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	<i>k</i> =1			<i>k</i> =2		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	0.22	88.26	309.53	10.57	90.64	11397.01
ELECTRICITY	0.05	85.50	2.06	0.14	89.30	6.36
POKER-HAND	0.04	76.17	38.31	0.28	79.74	185.95
COVPOKELEC	0.26	80.84	1255.49		<i>DNC</i>	
RBF(0,0)	33.79	52.35 ± 0.23	134.35	5.63	80.57 ± 0.08	492.00
RBF(50,0.001)	18.37	32.18 ± 0.56	100.98	0.14	43.14 ± 0.08	429.24
RBF(10,0.001)	32.78	49.85 ± 0.20	131.64	4.89	74.18 ± 0.07	484.87
RBF(50,0.0001)	0.75	49.20 ± 0.12	62.00	0.21	65.72 ± 0.12	452.44
RBF(10,0.0001)	20.78	50.67 ± 0.16	111.49	2.23	77.27 ± 0.08	476.51
RT	17.95	48.20 ± 0.28	98.34	124.33	66.84 ± 0.93	558.00
LED(50000)	0.50	71.16 ± 0.27	222.10	6.29	72.65 ± 0.09	3620.69
Waveform(50)	51.54	82.00 ± 0.36	243.66	56.64	85.89 ± 0.21	1732.22
HYP(10,0.001)	5.79	75.91 ± 0.39	61.27	2.33	88.94 ± 0.38	228.39
HYP(10,0.0001)	10.98	75.68 ± 1.32	79.08	5.22	90.34 ± 0.14	257.16
65.57 Acc.			77.32 Acc.			
0.15 RAM-Hours			1.21 RAM-Hours			

pared to the majority class. When performing a prediction on a test instance, the leaf will only return a naive Bayes prediction if it has been more accurate overall than the majority class, otherwise it resorts to a majority class prediction.

Bagging using ADWIN [Bifet et al. 2009] is based on the online bagging method of [Oza and Russell 2001b] with the addition of the ADWIN algorithm to detect changes in accuracy for each ensemble member. It yields excellent predictive performance for evolving data streams. When a change is detected, the classifier of the ensemble with the lowest accuracy, as predicted by its ADWIN detector, is removed and a new classifier is added to the ensemble. We use ADWIN Bagging to compare to our new ensemble learner, in both cases using **hnbt** as the base learner for the ensemble members.

Table III. Comparison using sets of 3 or 4 attributes. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	$k=3$			$k=4$		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	240.83	91.46	241144.76		<i>DNC</i>	
ELECTRICITY	0.27	90.33	13.78	0.43	90.52	19.68
POKER-HAND	0.72	80.80	588.20	1.17	80.67	1284.38
COVPOKELEC		<i>DNC</i>			<i>DNC</i>	
RBF(0,0)	21.84	85.78 ± 0.07	1834.72	50.27	87.67 ± 0.04	4041.01
RBF(50,0.001)	0.38	48.59 ± 0.10	1563.64	0.73	51.19 ± 0.06	3484.85
RBF(10,0.001)	17.89	79.13 ± 0.08	1799.88	37.92	81.30 ± 0.08	4105.68
RBF(50,0.0001)	0.68	74.28 ± 0.13	1609.09	1.46	77.44 ± 0.18	3677.13
RBF(10,0.0001)	5.26	83.58 ± 0.03	1780.45	8.05	85.78 ± 0.04	4102.96
RT	249.04	85.06 ± 0.11	1467.27	141.63	90.90 ± 0.34	2459.67
LED(50000)	68.97	72.66 ± 0.08	47380.04	659.22	72.79 ± 0.01	254752.81
Waveform(50)	199.21	85.84 ± 0.29	18774.08		<i>DNC</i>	
HYP(10,0.001)	7.70	90.40 ± 0.36	788.07	17.21	90.93 ± 0.28	1770.61
HYP(10,0.0001)	17.70	91.47 ± 0.15	865.02	37.23	91.68 ± 0.14	1950.51
	81.49 Acc.			81.90 Acc.		
	72.00 RAM-Hours			72.99 RAM-Hours		

Table IV. Comparison using ADWIN Bagging of 10 and 100 trees. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	ADWIN Bagging 10 Trees			ADWIN Bagging 100 Trees		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	0.83	84.75	201.59	13.06	84.76	2662.80
ELECTRICITY	0.17	84.11	3.62	3.98	84.59	30.48
POKER-HAND	0.09	74.04	65.27	1.40	71.13	886.78
COVPOKELEC	0.99	78.36	425.65	16.90	76.55	6692.23
RBF(0,0)	6.11	87.67 ± 0.52	191.84	60.52	88.48 ± 0.28	2143.47
RBF(50,0.001)	0.09	51.15 ± 0.11	220.50	0.92	44.65 ± 0.04	2332.09
RBF(10,0.001)	5.10	80.35 ± 0.25	200.24	56.51	81.31 ± 0.12	2126.37
RBF(50,0.0001)	0.19	59.94 ± 0.28	225.65	1.90	60.67 ± 0.32	2401.50
RBF(10,0.0001)	5.76	83.70 ± 0.24	202.33	67.14	84.27 ± 0.29	2175.74
RT	20.32	91.16 ± 0.15	156.06	202.99	91.27 ± 0.13	1686.98
LED(50000)	1.35	73.12 ± 0.07	203.19	24.44	73.07 ± 0.06	3044.94
Waveform(50)	1.68	83.63 ± 0.52	277.35	52.33	85.33 ± 0.34	2759.82
HYP(10,0.001)	2.59	89.68 ± 0.34	103.86	42.64	89.01 ± 0.46	1192.34
HYP(10,0.0001)	6.80	91.04 ± 0.24	117.97	118.73	90.75 ± 0.48	1367.94
	79.48 Acc.			78.99 Acc.		
	0.04 RAM-Hours			5.67 RAM-Hours		

4. COMPARISON WITH BAGGING

Our first experimental evaluation compares our new stacking method using restricted Hoeffding trees with bagging, and we use the datasets discussed in the previous section for this evaluation. The evaluation methodology used was Interleaved Test-Then-Train or Prequential evaluation (based on 10 runs for the artificial data): every example was used for testing the model before using it for training [Gama et al. 2009]. This interleaved test followed by train procedure was carried out on the full training set in each case. The parameters of the artificial streams are as follows:

- RBF(x,v): RandomRBF data stream of x centroids moving at speed v .
- HYP(x,v): Hyperplane data stream with x attributes changing at speed v .
- RT: Random Tree data stream
- WAVEFORM(v): Waveform dataset, with v as the value of the drift parameter.
- LED(v): LED dataset, with v as the value of the drift parameter.

We test our method using restricted Hoeffding trees containing 1, 2, 3, and 4 attributes respectively, against Bagging using ADWIN. We do not show results for Online Bagging [Oza and Russell 2001b] as in [Bifet et al. 2009] it was shown that Bagging using ADWIN outperforms Online Bagging. Note that we use the average class probability estimates of the ensemble members in bagging, because we found that this yields better results than simply majority voting.

Results for a single Hoeffding Naive Bayes Tree (**hnbt**) are shown in Table I. Bagging using ADWIN (10 trees) in Table IV is at least two percent more accurate than **hnbt** on all but two of the datasets we used, POKER-HAND and COVPOKELEC. The former dataset is the only one on which **hnbt** achieves substantially higher accuracy (77.1%) than Bagging using ADWIN.

Tables II, III and IV report the final accuracy, speed, and memory consumption of the ensemble classification models induced on the synthetic data and the real-world datasets: FOREST COVERTYPE, POKER-HAND, ELECTRICITY and COVPOKELEC (the mixture of the three). Accuracy is measured as the final percentage of examples correctly classified over the test/train interleaved evaluation. Time is measured in seconds, and memory in MB. Some methods did not complete (DNC in the tables) due to the lack of memory required by the new methods that generate a large number of ensemble members on some datasets. Because of this, averages are computed across those datasets that finished.

The results of our method for different values of k , and the comparison to bagging, show that excellent predictive performance can be obtained on the practical datasets for small values of k : even with $k = 1$, the proposed method outperforms bagging on the real-world datasets. Further substantial improvements can be obtained by moving from $k = 1$ to $k = 2$. These results indicate that modeling low-dimensional interactions is sufficient for obtaining good performance on the practical datasets we considered.

Note that, in a separate experiment, we have also investigated the effect of disabling the ADWIN-based change detection mechanism for resetting the learning rate. Disabling the mechanism yields a reduction in average accuracy of about 1.6% for $n=2$ and 0.9% for $n=3$. This shows that change detection for resetting the learning rate is indeed beneficial.

Also, we note that Bagging using ADWIN with 100 trees is no more accurate overall than using 10 trees. Hence, the good performance of our method is not due to the fact that it uses a larger number of trees. The combination of using attribute subsets and the perceptron weighting methodology is essential for the improvement in accuracy obtained.

As the RandomRBF, Random Tree, and Hyperplane data streams have 10 attributes, the number of attribute sets used by the restricted trees for these datasets are

$$\binom{10}{1} = 10 \quad \binom{10}{2} = 45 \quad \binom{10}{3} = 120 \quad \binom{10}{4} = 210 \quad \binom{10}{5} = 252$$

respectively. When the number of attributes is bigger than 15, the number of combinations is huge as combinations grow exponentially. For example, for the FOREST COVERTYPE data with 54 attributes we obtain the following figures:

$$\binom{54}{1} = 54 \quad \binom{54}{2} = 1,431 \quad \binom{54}{3} = 24,804 \quad \binom{54}{4} = 316,251$$

We observe that for $k = 2$ ($\binom{m}{2} = m \cdot (m - 1)/2$) the number of possible combinations is relatively small, and we can run experiments on all datasets. However, for $k = 4$,

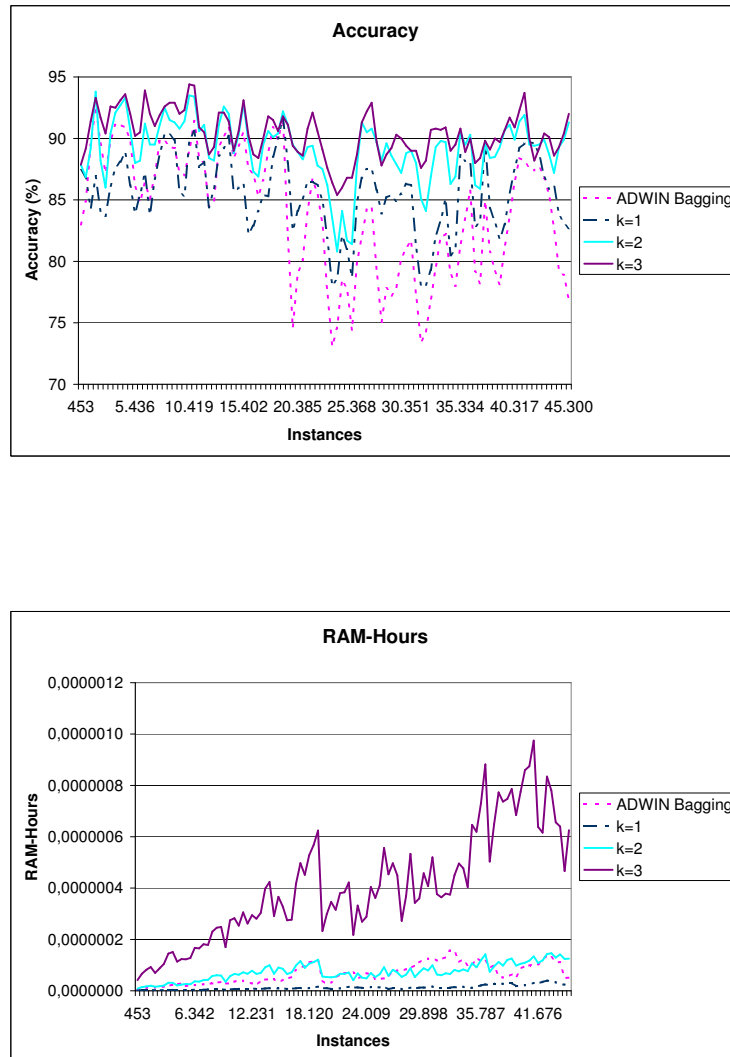


Fig. 1. Accuracy and RAM-Hours on the Electricity dataset.

the number of combinations is so large ($\binom{m}{4} = m \cdot (m-1) \cdot (m-2) \cdot (m-3)/24$) that for some datasets insufficient memory was available.

The learning curves and model growth curves for the ELECTRICITY dataset are plotted in Figures 1 and 2. We use the prequential method [Gama et al. 2009] with a sliding window of 1,000 examples to compute the learning curves. We observe that using $k = 1$ we obtain the fastest and most resource-efficient method, but also the worst in accuracy. On the other hand, with $k = 3$ we have the slowest method, the one

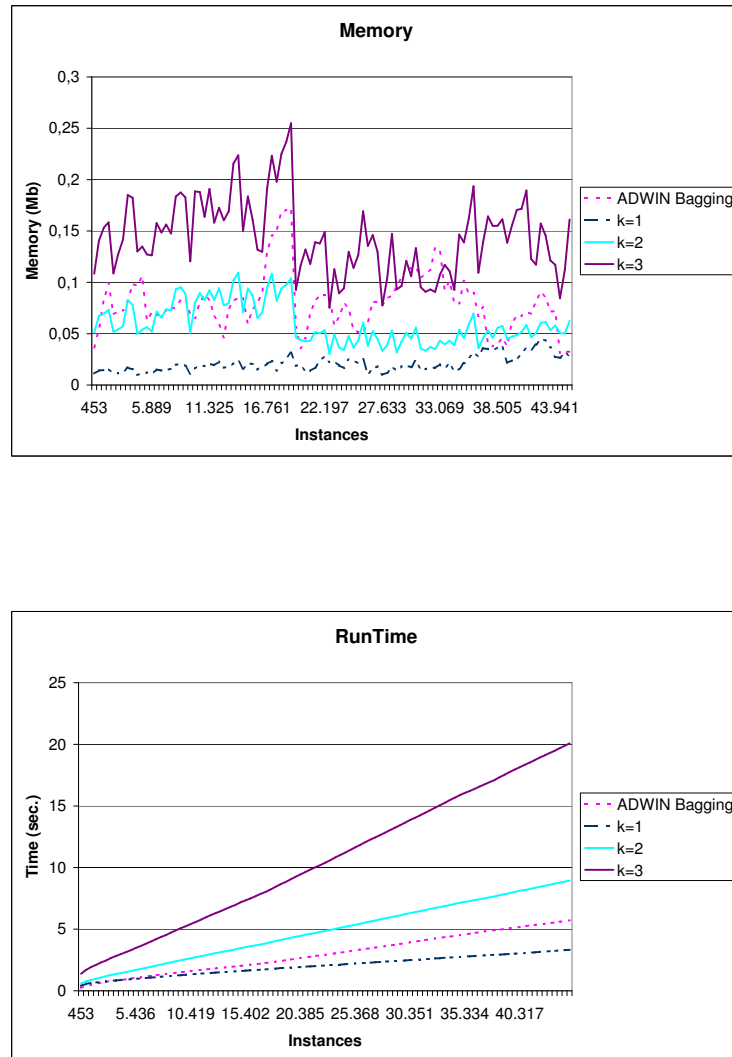


Fig. 2. Runtime and memory on the Electricity dataset.

that uses more RAM-Hours, but also the most accurate. As we increase the complexity of the model using higher values for k , we use more resources, but we improve the method's accuracy. ADWIN Bagging appears to use a similar amount of RAM-Hours as our method with $k = 2$, but its accuracy for this dataset is worse than with $k = 1$.

The results on the artificial data yield a somewhat different picture. On these data streams, larger values of k are necessary in most cases to obtain classification accuracy that is competitive with bagging. This is not surprising given the functional rela-

Table V. Performance for stacking using trees with $m-2$ and $m-3$ attributes each. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	$k=m-2$			$k=m-3$		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	86.22	82.51	67851.40		<i>DNC</i>	
ELECTRICITY	0.36	88.86	7.91	0.43	89.67	17.72
POKER-HAND	0.30	77.89	397.29	0.71	78.86	965.91
COVPOKELEC	219.18	78.41	220296.88		<i>DNC</i>	
RBF(0,0)	19.85	88.69 ± 0.15	1105.89	47.47	89.06 ± 0.11	2841.03
RBF(50,0.001)	0.25	52.82 ± 0.07	1070.34	0.59	53.88 ± 0.08	2567.74
RBF(10,0.001)	11.18	82.36 ± 0.11	1127.82	28.11	83.20 ± 0.05	2885.33
RBF(50,0.0001)	0.65	69.61 ± 0.37	1238.35	1.40	74.44 ± 0.33	2724.39
RBF(10,0.0001)	2.20	85.69 ± 0.12	1156.60	4.90	86.91 ± 0.08	2861.01
RT	39.81	92.61 ± 0.17	759.26	79.15	93.33 ± 0.20	1780.27
LED(50000)	5.56	73.06 ± 0.08	16113.59	45.41	73.04 ± 0.09	131258.32
Waveform(50)	95.07	85.28 ± 0.36	6458.53	481.69	85.26 ± 0.30	49611.36
HYP(10,0.001)	5.92	90.93 ± 0.24	528.76	13.95	91.30 ± 0.22	1220.77
HYP(10,0.0001)	13.78	91.62 ± 0.10	534.46	32.34	91.87 ± 0.08	1419.83
		81.45 Acc.			82.57 Acc.	
		43.25 RAM-Hours			39.97 RAM-Hours	

tionships underlying the data generation processes. Nevertheless, competitive performance can be obtained with $k = 3$ or $k = 4$.

Let us briefly consider our method in the case where each tree in the ensemble can access a large number of attributes—more specifically, close to the full number of attributes m . This case is interesting to consider because $\binom{m}{k} = \binom{m}{m-k}$, although using $m - k$ as subset size does introduce an extra factor of m in overall time complexity. Table V shows results for our method using $m - 2$ and $m - 3$ attributes respectively. Predictive performance on the artificial data streams is very good compared to the low-dimensional cases considered previously (where k was small) because these data streams benefit from modeling higher-dimensional interactions. However, using small values of k is clearly preferable on the real-world data, yielding higher accuracy with lower resource use.

4.1. Ensemble Diversity

We use the Kappa statistic κ [Margineantu and Dietterich 1997] to show how using the new classifier based on stacking with restricted Hoeffding trees, increases the diversity of the ensemble.

The Kappa statistic is a measure defined so that if two classifiers agree on every example then $\kappa = 1$, and if their predictions coincide purely by chance, then $\kappa = 0$. Given two classifiers h_a and h_b , and a data set containing m examples, the Kappa statistic is based on a contingency table where cell C_{ij} contains the number of examples for which $h_a(x) = i$ and $h_b(x) = j$.

The κ statistic is defined formally as:

$$\kappa = \frac{\Theta_1 - \Theta_2}{1 - \Theta_2}$$

The κ statistic uses Θ_2 for normalization, the probability that two classifiers agree by chance, given the observed counts in the table. If h_a and h_b make identical classifications on the data set, then all non-zero counts will appear along the diagonal. If h_a and h_b are very different, then there will be a large number of counts off the diagonal. This is used for Θ_1 . We define

$$\Theta_1 = \frac{\sum_{i=1}^L C_{ii}}{m}$$

Table VI. Comparing accuracy when using only the 10 or 40 best classifiers with trees of 2 attributes. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB.

	Top 10 trees			Top 40 trees		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	10.57	52.27	9100.09	10.57	54.46	9052.99
ELECTRICITY	0.14	85.44	5.70		<i>DNC</i>	
POKER-HAND	0.28	69.14	158.38	0.28	75.23	184.21
COVPOKELEC	17.89	55.25	53622.51	17.89	64.23	47917.17
RBF(0,0)	5.64	67.17 ± 0.33	435.28	5.65	78.19 ± 0.27	488.82
RBF(50,0.001)	0.14	28.91 ± 0.26	358.09	0.14	40.66 ± 0.09	422.18
RBF(10,0.001)	4.88	63.12 ± 0.32	422.03	4.88	72.06 ± 0.25	476.67
RBF(50,0.0001)	0.21	34.51 ± 0.89	378.15	0.21	59.13 ± 0.29	445.12
RBF(10,0.0001)	2.24	58.80 ± 1.21	407.90	2.30	73.17 ± 0.26	477.25
RT	124.33	14.29 ± 0.45	518.19	125.19	65.67 ± 0.53	549.01
LED(50000)	6.29	9.99 ± 0.03	3035.01	6.29	9.99 ± 0.03	3057.94
Waveform(50)	56.64	42.78 ± 10.02	1557.63	56.64	68.05 ± 8.20	1601.11
HYP(10,0.001)	2.19	71.06 ± 4.14	203.96	2.19	87.09 ± 1.07	226.83
HYP(10,0.0001)	5.32	73.58 ± 5.25	234.52	5.13	87.96 ± 0.82	257.55
	51.88 Acc. 4.52 RAM-Hours			64.30 Acc. 4.20 RAM-Hours		

$$\Theta_2 = \sum_{i=1}^L \left(\sum_{j=1}^L \frac{C_{ij}}{m} \cdot \sum_{j=1}^L \frac{C_{ji}}{m} \right)$$

We could use Θ_1 as a measure of agreement, but in problems where one class is much more common than the others, all classifiers will agree by chance, so all pairs of classifiers will obtain high values for Θ_1 .

The Kappa-Error diagram consists of a plot where each point corresponds to a pair of classifiers. The x coordinate is the κ value for the two classifiers concerned. The y coordinate is their average error.

Figures 3 and 4 show the Kappa-Error diagram for the hyperplane data stream with 10 attributes changing at speed 0.0001 and the Electricity dataset respectively. As for $k = 2$, $k = 3$, $k = m - 2$ and $k = m - 3$, the number of classifiers is large, we only plot the points of the significant classifiers, the ones with higher perceptron weights. More specifically those that represent 90% of the total weight of the coefficients of the ensemble. The figures generally show that the new stacking method produces more diverse ensembles. The ADWIN bagging ensembles are tightly clustered whereas the stacking method produces more disagreement among classifiers in the ensemble. This is the case in particular for $k = 2$ and $k = 3$, as one would expect.

5. PRUNING RESTRICTED Hoeffding Trees

Although the results on the real-world data streams show that small values of k are often sufficient to obtain accurate classifiers, it is instructive to consider whether it is possible, at least in principle, to prune the ensemble by reducing the number of classifiers, without losing accuracy. To investigate this, we performed a further experiment, using only those trees for prediction that exhibited the largest average coefficients in the perceptrons, where the average was taken across class values. Table VI shows the results for reduced sets of 10 and 40 classifiers respectively, for $k = 2$. Even for 40 trees, performance is substantially below that of the full ensemble, indicating that a pruning strategy based on simply removing classifiers in this manner is not sufficient to maintain high accuracy.

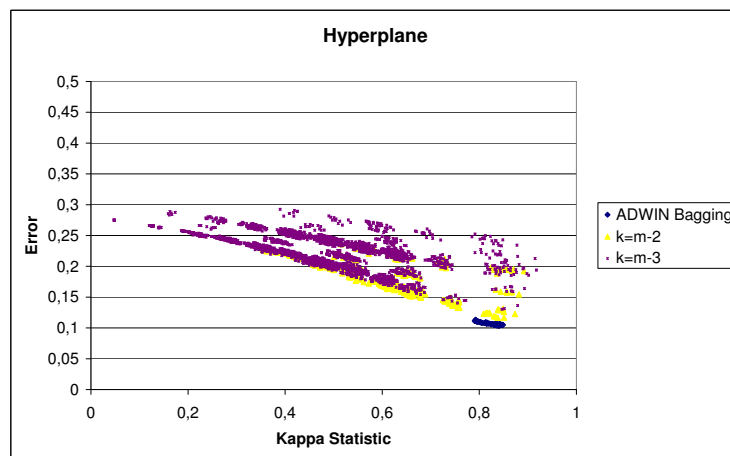
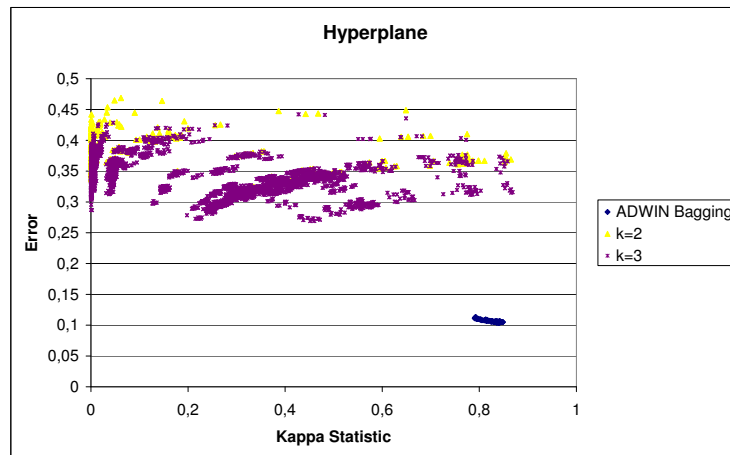


Fig. 3. Kappa-Error diagrams for ADWIN Bagging and stacking with restricted Hoeffding trees with $k = 2$, $k = 3$, $k = m - 2$ and $k = m - 3$ on the hyperplane data stream.

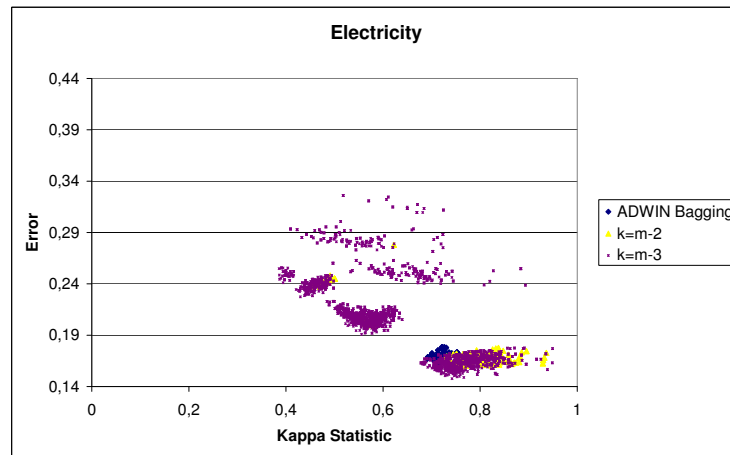
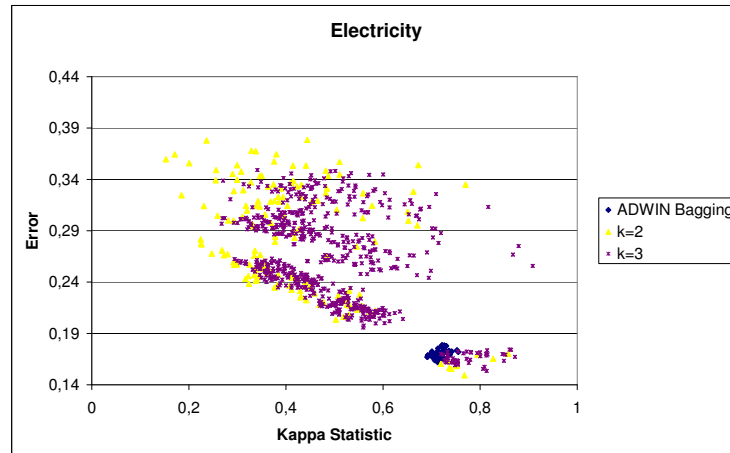


Fig. 4. Kappa-Error diagrams for ADWIN Bagging and stacking with restricted Hoeffding trees with $k = 2$, $k = 3$, $k = m - 2$ and $k = m - 3$ on the Electricity dataset.

Table VII. Comparison using ADWIN Bagging* of 10 and 100 trees, with new strategy of replacing classifiers that exhibit decreased accuracy with new ones. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	ADWIN Bagging* 10 Trees			ADWIN Bagging* 100 Trees		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	0.50	83.62	136.58	5.19	83.82	1770.60
ELECTRICITY	0.11	83.58	2.57	1.18	83.76	23.42
POKER-HAND	0.09	75.18	62.26	0.85	75.49	502.19
COVPOKELEC	0.68	78.48	318.27	6.79	78.72	4993.54
RBF(0,0)	6.10	87.64 ± 0.48	166.02	59.99	88.39 ± 0.30	1788.53
RBF(50,0.001)	0.08	50.88 ± 0.10	196.16	0.80	51.32 ± 0.08	1901.54
RBF(10,0.001)	2.62	79.66 ± 0.33	165.88	23.17	80.75 ± 0.17	1907.64
RBF(50,0.0001)	0.19	59.24 ± 0.27	207.26	1.66	59.38 ± 0.26	1976.48
RBF(10,0.0001)	0.66	76.94 ± 0.57	188.78	5.15	78.63 ± 0.27	2168.74
RT	20.34	91.13 ± 0.14	137.77	202.62	91.25 ± 0.13	1496.53
LED(50000)	1.40	73.11 ± 0.07	171.74	9.34	73.12 ± 0.08	1960.94
Waveform(50)	2.91	84.14 ± 0.54	196.41	27.15	84.52 ± 0.52	2303.07
HYP(10,0.001)	2.07	90.05 ± 0.26	88.52	19.13	90.48 ± 0.29	934.18
HYP(10,0.0001)	4.10	91.21 ± 0.19	95.66	41.36	91.68 ± 0.18	1130.24
	78.92 Acc. 0.02 RAM-Hours			79.38 Acc. 2.73 RAM-Hours		

Table VIII. Comparison using sets of 2 or 3 attributes with ADWIN strategy of replacing the worst classifier in the ensemble. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	$k = 2$ replacing worst classifier			$k = 3$ replacing worst classifier		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	72.43	82.02	19749.91	0.00	<i>DNC</i>	0.00
ELECTRICITY	0.47	83.32	6.39	1.16	84.90	13.80
POKER-HAND	0.56	74.16	220.08	2.44	77.63	788.86
COVPOKELEC	131.88	75.44	151330.51	0.00	<i>DNC</i>	0.00
RBF(0,0)	5.62	73.89 ± 0.43	444.65	21.91	81.24 ± 0.32	1643.85
RBF(50,0.001)	0.14	31.72 ± 1.60	363.17	0.51	34.21 ± 0.66	1413.85
RBF(10,0.001)	5.28	68.65 ± 0.46	449.17	20.68	74.95 ± 0.18	1550.86
RBF(50,0.0001)	0.29	42.81 ± 1.63	409.55	1.17	56.24 ± 0.35	1396.60
RBF(10,0.0001)	5.49	69.01 ± 0.50	414.10	22.00	76.25 ± 0.25	1616.54
RT	119.42	59.90 ± 0.51	507.61	268.19	77.92 ± 0.32	1369.13
LED(50000)	6.59	58.78 ± 1.95	3000.86	79.36	68.60 ± 0.10	43172.92
Waveform(50)	53.36	83.61 ± 1.50	1525.19	311.42	85.18 ± 0.60	16085.90
HYP(10,0.001)	10.82	78.72 ± 1.85	245.69	40.72	82.21 ± 2.20	839.50
HYP(10,0.0001)	14.83	88.52 ± 0.91	260.20	53.17	89.88 ± 0.63	885.63
	69.33 Acc. 20.73 RAM-Hours			74.10 Acc. 15.80 RAM-Hours		

6. USING STACKING AND THE NEW ADWIN STRATEGY WITH BAGGED TREES

To conclude our experimental evaluation, we consider empirically if the components of the new method presented in this paper can be easily applied to standard bagged decision trees, specifically:

- the new strategy of replacement of ensemble classifiers: when one of the ADWIN estimators detects change, the classifier monitored by this estimator is replaced by a new one
- the use of stacking instead of majority voting as a combination mechanism for obtaining the ensemble predictions

We also considered the effect of averaging instead of stacking in our method, and the effect of using the old ADWIN-based classifier replacement strategy in our method.

Table IX. Performance for $k = 2$ and $k = 3$ when probability estimates are simply averaged. Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	$k = 2$			$k = 3$		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	10.57	36.59	7992.86	240.83	36.67	224105.06
ELECTRICITY	0.14	44.53	6.23	0.27	44.54	13.90
POKER-HAND	0.28	50.16	187.53	0.72	50.15	580.98
COVPOKELEC	17.89	44.53	66759.11		<i>DNC</i>	
RBF(0,0)	5.63	73.08 ± 0.08	436.23	21.79	82.03 ± 0.10	1621.30
RBF(50,0.001)	0.14	26.24 ± 0.08	366.55	0.38	26.26 ± 0.09	1359.20
RBF(10,0.001)	4.89	67.56 ± 0.27	415.29	17.97	75.29 ± 0.14	1507.81
RBF(50,0.0001)	0.21	27.00 ± 0.12	413.91	0.68	27.43 ± 0.10	1405.66
RBF(10,0.0001)	2.20	62.59 ± 1.82	423.59	5.21	62.99 ± 4.85	1556.24
RT	122.97	45.01 ± 0.42	510.39	249.81	52.21 ± 0.96	1342.32
LED(50000)	6.29	72.82 ± 0.06	2931.12	68.95	72.82 ± 0.07	42748.55
Waveform(50)	59.35	79.84 ± 2.76	1530.27	212.41	81.33 ± 2.87	16469.82
HYP(10,0.001)	2.33	53.05 ± 1.27	210.50	8.01	53.95 ± 1.64	702.58
HYP(10,0.0001)	5.22	73.38 ± 4.28	234.68	17.54	78.58 ± 2.83	779.17
		54.03 Acc.			57.25 Acc.	
		5.32 RAM-Hours			67.40 RAM-Hours	

First, we implemented ADWIN Bagging*, an extension to ADWIN Bagging with the same strategy of replacing the classifiers as our new stacking method with restricted Hoeffding Trees. Table VII shows the results of this extension using 10 and 100 base classifiers. If we compare these results with the results for standard ADWIN Bagging in Table IV, we observe that average accuracy increases for bagging with 100 trees and decreases with bagging with 10 trees. Looking at these results, it appears that the best strategy for a small number of trees is to replace the worst classifier when a change is detected in the accuracy of one of the learners. For large ensembles of classifiers, the best strategy is to replace the classifier that exhibited change. However, there are exceptions to this rule. Overall, the differences appear comparatively small.

Second, we tested the use of the other ADWIN strategy, replacing the worst classifier when change is detected, using the new method proposed in this paper, stacking with restricted Hoeffding trees. Table VIII shows the results. Comparing this with the results in Tables II and III shows that the new replacement strategy is essential for stacked ensembles of restricted Hoeffding trees.

Third, we investigated whether it is necessary to train a perceptron to combine the trees' predictions. To this end, we show results for simple averaging of the trees' probability estimates in Table IX. We see that using the simple averaging mechanism we use the same resources, but obtain lower accuracy. Hence, combining predictions adaptively is clearly essential. This is due to the fact that not all attribute subsets (and, thus, Hoeffding trees) are equally relevant for prediction.

Finally, we tested using stacking to combine the bagged trees' predictions, and we show results in Table X. Comparing these results with standard ADWIN Bagging in Table IV or ADWIN Bagging* in Table VII, we observe that using a perceptron to combine the prediction results of the base learners improves the global accuracy of the ensemble. It is interesting to observe that using 100 trees in standard ADWIN Bagging did not improve global accuracy over using 10 trees, however, using the perceptron, as the contribution of every learner prediction is not the same, we improve global accuracy of the ensemble. On real-world datasets, we note that looking at Tables II and III, limited-interaction trees still have an edge.

Considering the artificial datasets, the two best results overall reported in this paper are for our new method with $k = m - 3$ in Table V and for stacking 100 bagged trees

Table X. Comparison using stacking with perceptron for 10 and 100 bagged decision trees . Accuracy is measured as the final percentage of examples correctly classified over the full test/train interleaved evaluation. Time is measured in seconds, and memory in MB. The best individual accuracies are indicated in boldface.

	Stacking 10 Bagged Trees			Stacking 100 Bagged Trees		
	Mem.	Acc.	Time	Mem.	Acc.	Time
COVTYPE	0.67	86.11	198.15	6.25	88.01	2526.41
ELECTRICITY	0.10	86.97	4.76	1.22	88.78	38.28
POKER-HAND	0.08	76.35	89.82	0.86	78.64	832.82
COVPOKELEC	0.86	80.13	520.95	8.01	82.34	8479.41
RBF(0,0)	5.85	87.84 ± 0.25	224.52	57.57	89.41 ± 0.11	2583.72
RBF(50,0.001)	0.10	50.42 ± 0.14	278.55	0.91	55.53 ± 0.06	2567.34
RBF(10,0.001)	3.39	80.80 ± 0.12	245.35	32.71	83.19 ± 0.06	2738.54
RBF(50,0.0001)	0.19	62.73 ± 0.31	291.49	1.74	70.30 ± 0.40	2895.87
RBF(10,0.0001)	0.69	82.07 ± 0.38	258.99	5.55	86.52 ± 0.15	2674.73
RT	20.34	91.12 ± 0.09	165.91	201.34	92.15 ± 0.13	1732.35
LED(50000)	0.22	73.08 ± 0.08	259.77	2.11	73.09 ± 0.08	3038.83
Waveform(50)	2.36	84.74 ± 0.37	291.24	21.44	85.15 ± 0.31	3289.44
HYP(10,0.001)	1.03	89.91 ± 0.21	117.59	9.08	91.11 ± 0.22	1191.91
HYP(10,0.0001)	1.20	90.61 ± 0.16	128.85	9.42	91.82 ± 0.16	1335.41
	80.21 Acc. 0.03 RAM-Hours			82.57 Acc. 3.49 RAM-Hours		

in Table X. It is interesting to note that they have quite similar accuracy and an order of magnitude difference in the use of resources. Clearly the stacking procedure merits further investigation.

7. CONCLUSIONS

This paper has presented a new method for data stream classification that is based on combining an exhaustive ensemble of limited-attribute-set tree classifiers using stacking. The approach generates tree classifiers for all possible attribute subsets of a given user-specified size k . For $k = 2$, this incurs an extra factor of k in time complexity compared to building a single tree grown from all attributes. For larger values of k , the method is only practical for datasets with a small number of attributes (unless k is chosen to be close to the number of attributes). However, even if limited to $k = 2$, the method appears to be a promising candidate for high-accuracy data stream classification on practical data streams. Our results on real-world datasets show excellent classification accuracy.

Part of our method is an approach for resetting the learning rate for the perceptron meta-classifiers that are used to combine the trees' predictions. This method, based on the ADWIN change detector, is employed because the meta-data stream is not identically distributed: the trees' predictions improve over time. When change is detected, the learning rate is reset to its initial, larger value, so that meta learning quickly adapts to the improved predictions of the base models.

The ADWIN change detector is also used to reset ensemble members when their predictive accuracy degrades significantly, in which case the learning rate is reset also. This makes it possible to use the ensemble classifier for evolving data streams where the distribution changes over time. Promising results are demonstrated for applying the new classifier replacement strategy introduced in this paper to bagging. When dealing with a large number of trees the best strategy is generally to replace the classifier that exhibits an increased error rate.

We have empirically verified that it is important to use stacking to learn how to combine the tree classifiers' predictions in an appropriate manner: simple averaging yields very poor accuracy. This is not surprising because not all attribute subsets are equally relevant. We have also shown that stacking can be beneficial for bagging. Our initial

experiments with ensemble pruning, discarding the least important ensemble members, did not prove successful. In future work, we would like to investigate smarter pruning techniques and methods for pre-selecting attribute subsets, e.g. based on using a diversity-based measure, or based on maximizing pairwise Hamming distance between indicator vectors. We also plan to investigate whether using stacking to combine ensemble predictions is beneficial in non-streaming settings.

REFERENCES

- BIFET, A. AND GAVALDÀ, R. 2007. Learning from time-changing data with adaptive windowing. In *SIAM International Conference on Data Mining*.
- BIFET, A., HOLMES, G., KIRKBY, R., AND PFAHRINGER, B. 2010. MOA: Massive Online Analysis. *Journal of Machine Learning Research* 11, 1601–1604.
- BIFET, A., HOLMES, G., AND PFAHRINGER, B. 2010. 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.
- BIFET, A., HOLMES, G., PFAHRINGER, B., AND FRANK, E. 2010. Fast perceptron decision tree learning from evolving data streams. In *Advances in Knowledge Discovery and Data Mining, 14th Pacific-Asia Conference, PAKDD 2010, Hyderabad, India, June 21-24, 2010. Proceedings. Part II*. 299–310.
- BIFET, A., HOLMES, G., PFAHRINGER, B., AND GAVALDÀ, R. 2009. Improving adaptive bagging methods for evolving data streams. In *First Asian Conference on Machine Learning ACML*. 23–37.
- BIFET, A., HOLMES, G., PFAHRINGER, B., KIRKBY, R., AND GAVALDÀ, R. 2009. New ensemble methods for evolving data streams. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 139–148.
- BREIMAN, L. 2001. Random forests. In *Machine Learning*. 5–32.
- BREIMAN, L., FRIEDMAN, J. H., OLSHEN, R. A., AND STONE, C. J. 1984. *Classification and Regression Trees*. Wadsworth.
- DOMINGOS, P. AND HULTEN, G. 2000. Mining high-speed data streams. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 71–80.
- FRANK, A. AND ASUNCION, A. 2010. UCI machine learning repository.
- GAMA, J., MEDAS, P., CASTILLO, G., AND RODRIGUES, P. P. 2004. Learning with drift detection. In *SBIA Brazilian Symposium on Artificial Intelligence*. 286–295.
- GAMA, J., ROCHA, R., AND MEDAS, P. 2003. Accurate decision trees for mining high-speed data streams. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 523–528.
- GAMA, J., SEBASTIÃO, R., AND RODRIGUES, P. P. 2009. Issues in evaluation of stream learning algorithms. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 329–338.
- HARRIES, M. 1999. Splice-2 comparative evaluation: Electricity pricing. Tech. rep., The University of South Wales.
- HO, T. K. 1998. The random subspace method for constructing decision forests. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 20, 8, 832–844.
- HOLMES, G., KIRKBY, R., AND PFAHRINGER, B. 2005. Stress-testing Hoeffding trees. In *Knowledge Discovery in Databases: PKDD 2005, 9th European Conference on Principles and Practice of Knowledge Discovery in Databases, Porto, Portugal, October 3-7, 2005, Proceedings*. 495–502.
- HULTEN, G., SPENCER, L., AND DOMINGOS, P. 2001. Mining time-changing data streams. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 97–106.
- KIRKBY, R. 2007. Improving hoeffding trees. Ph.D. thesis, University of Waikato.
- KOLTER, J. Z. AND MALOOF, M. A. 2007. Dynamic weighted majority: An ensemble method for drifting concepts. *J. Mach. Learn. Res.* 8, 2755–2790.
- MARGINEANTU, D. D. AND DIETTERICH, T. G. 1997. Pruning adaptive boosting. In *Proceedings of the Fourteenth International Conference on Machine Learning (ICML 1997), Nashville, Tennessee, USA, July 8-12, 1997*. 211–218.
- OZA, N. C. AND RUSSELL, S. J. 2001a. Experimental comparisons of online and batch versions of bagging and boosting. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 359–364.
- OZA, N. C. AND RUSSELL, S. J. 2001b. Online bagging and boosting. In *Artificial Intelligence and Statistics 2001*. 105–112.
- WOLPERT, D. H. 1992. Stacked generalization. *Neural Networks* 5, 2, 241–259.