

Adaptive Ensemble Active Learning for Drifting Data Stream Mining

Bartosz Krawczyk^{1*} and Alberto Cano¹

¹Department of Computer Science, Virginia Commonwealth University, Richmond VA, USA
 {bkrawczyk,acano}@vcu.edu

Abstract

Learning from data streams is among the most vital contemporary fields in machine learning and data mining. Streams pose new challenges to learning systems, due to their volume and velocity, as well as ever-changing nature caused by concept drift. Vast majority of works for data streams assume a fully supervised learning scenario, having an unrestricted access to class labels. This assumption does not hold in real-world applications, where obtaining ground truth is costly and time-consuming. Therefore, we need to carefully select which instances should be labeled, as usually we are working under a strict label budget. In this paper, we propose a novel active learning approach based on ensemble algorithms that is capable of using multiple base classifiers during the label query process. It is a plug-in solution, capable of working with most of existing streaming ensemble classifiers. We realize this process as a Multi-Armed Bandit problem, obtaining an efficient and adaptive ensemble active learning procedure by selecting the most competent classifier from the pool for each query. In order to better adapt to concept drifts, we guide our instance selection by measuring the generalization capabilities of our classifiers. This adaptive solution leads not only to better instance selection under sparse access to class labels, but also to improved adaptation to various types of concept drift and increasing the diversity of the underlying ensemble classifier.

1 Introduction

Velocity of data gave rise to the notion of data streams, potentially unbounded collections of data that continuously flood the system. We define data stream a sequence $\langle S_1, S_2, \dots, S_n, \dots \rangle$, where each element S_j is a new instance. In this paper, we work under the supervised learning scenario and thus we will define each instance as $S_j \sim p_j(x^1, \dots, x^d, y) = p_j(\mathbf{x}, y)$, where $p_j(\mathbf{x}, y)$ is a joint distribution of j -th instance, defined by a d -dimensional feature

space and originating from class y . Each instance in the stream is independent and randomly drawn from a stationary probability distribution $\Psi_j(\mathbf{x}, y)$. As new data is continuously arriving, storing a data stream is not a viable option. One needs to analyze new instances on-the-fly, incorporate the useful information into the classifier, and discard them. Both the prediction and classifier update steps cannot be of a high complexity, as instances arrive rapidly and bottlenecking must be avoided. Data streams are also subject to a phenomenon known as concept drift [Gama *et al.*, 2014], where the properties of stream are subject to a change over time.

Most of works dedicated to data stream mining assume a fully supervised learning process and thus unlimited access to class labels [Krawczyk *et al.*, 2017]. This assumption does not hold in real-life scenarios, where obtaining a ground truth for each instance is a costly and time-consuming process. Therefore, one needs to relax the requirements for class labels and work with sparsely labeled data streams. Active learning allows for selecting only the most valuable instances for label query and thus working under a strict budget.

In this paper, we propose a novel active learning method for drifting data streams, based on ensemble approach. We use a set of classifiers to decide which instances should be labeled, obtaining a more robust and accurate selection. We realize the instance selection procedure as a Multi-Armed Bandit problem, where each classifier is encoded as a bandit. We show how to conduct an adaptive instance selection for label query that learns from its past experiences and improves the performance over time. Additionally, we guide our instance selection not by popularly used classifier uncertainty metric, but by measuring how a new instance will affect the generalization capabilities of each base classifier in the ensemble. Our algorithm is a plug-in solution, capable of working with most of existing streaming ensemble classifiers. A thorough experimental study showcases that our method is capable of outperforming state-of-the-art active learning methods, is highly suitable for various concept drift scenarios, and has a beneficial influence on the diversity of the underlying ensemble, leading to increasing its drift adaptability.

2 Access to Ground Truth in Data Streams

Ever-growing volume, high-speed of instance arrival and concept drift presence are considered as the biggest challenges that must be handled by supervised learning algorithms in

*Contact Author

data stream mining. While there is a plethora of works dedicated to proposing new and efficient classifiers, many of them assume that the ground truth (class label y in the supervised classification case) is available immediately after a new instance S_j is processed. Therefore, these algorithms are designed to work with fully labeled data streams. However, such an assumption does not hold in the vast majority of real-life applications. Obtaining a class label requires either an access to an oracle-like source or a domain expert. The former case may hold for scenarios in which ground truth may be observed at no cost as a state of the environment (e.g., weather condition predictions). This is a rare case and is connected with delay at which labels become available. The latter case is a more frequent one and assumes a human annotator involved. As in data streams we deal with high-speed and massive volumes of instances, labeling all of them is costly, or even impossible due to human limitations in throughput. A scenario in which we can obtain labels only for a small subset of carefully selected instances is therefore a much more realistic one.

Active learning is an attractive solution to this problem, allowing for obtaining class labels at a limited budget [Lughofer, 2017]. Budget is defined as a number of label queries that can be made within a given time window and active learning strategies aim to optimize their usage to label the most important, diverse, or novel instances. While there exist a number of online active learning algorithms, they cannot be used for data streams as they do not possess any mechanisms for handling concept drift [Lughofer, 2017]. This is crucial for proper stream labeling, as the query strategy should adapt swiftly to changes in class characteristics [Ksieniewicz *et al.*, 2019]. At the same time, a proper sampling of instances coming from a new concept ensure a fast recovery rate after the concept drift has occurred. Most popular strategies use adaptive thresholds on classifier certainty (or support functions) [Zliobaite *et al.*, 2014], density information [Mohamad *et al.*, 2018a], error propagation [Mohamad *et al.*, 2018b], or combine active learning with drift detectors [Krawczyk *et al.*, 2018]. Recently, few works on ensemble learning for active learning from data streams emerged [Alabulrahman *et al.*, 2016; Xu *et al.*, 2016]. However, they treat the active learning compound as completely unrelated to the ensemble learner and thus should be considered as rather wrapper approaches, than core ensemble active learning.

All of existing active learning strategies are significantly impaired when the label query budget is small. In many real problems labeling even as little as 1% of instances from the stream may be too costly. Therefore, methods that allow obtaining class labels at no cost from drifting data streams are a promising research direction [Dyer *et al.*, 2014].

3 Ensemble Active Learning for Drifting Data Streams

In this section, we present the details of the proposed ensemble active learning algorithm dedicated to drifting data streams.

3.1 Active Learning for Data Streams

In classic AL scenario we assume having an access to classifier Ψ and a set of unlabeled instances \mathcal{U} :

$$\Psi : \mathcal{H} \times \mathcal{U} \rightarrow \mathbb{R}. \tag{1}$$

The role of active learning is to select the most useful instance from \mathcal{U} for label query, based on the current learning hypothesis (or model):

$$q = \arg \max_{x \in \mathcal{U}} \Psi(h, x). \tag{2}$$

Assuming an access to an oracle (or domain expert in real-world applications) that can provide a true class (over M -class problem) label for queried instance $\mathbf{o} : \mathcal{U} \rightarrow \{1, \dots, M\}$ and selected training algorithm A , we can formulate an incremental update of hypothesis h in i -th active learning iteration:

$$h_{i+1} = A(\{q_k, \mathbf{o}(q_k)\}_{k=1}^i), \tag{3}$$

where

$$q_i = \arg \max_{x \in \mathcal{U}_i} \Psi(h_i, x), \tag{4}$$

$$\mathcal{U}_{i+1} = \mathcal{U}_i \setminus \{q_i\}. \tag{5}$$

In classic AL scenario, hypothesis h depends on time (as it changes with every iteration), but Ψ remains constant. This however does not hold for a data stream scenario, where we need to constantly update the hypothesis and classification model, thus leading to a dynamic learning problem. Therefore, in the context of data streams we must assume that both h and Ψ may depend on time and adapt according to its past experiences:

$$q_i = \arg \max_{x \in \mathcal{U}_i} \Psi_i(h_i, x), \tag{6}$$

which is obtained by some adaptive learning process f :

$$\Psi_{i+1} = f((h_k)_{k=1}^i, (\Psi_k)_{k=1}^i, (q_k)_{k=1}^i, (\mathbf{o}(q_k))_{k=1}^i). \tag{7}$$

The considered function f used to update models over time could be any function that returns utility function - a popular problem in the context of data streams. In this paper, we propose to model this function as an ensemble learning problem controlled by multi-armed bandit adaptation.

3.2 Ensemble Active Learning as Multi-Armed Bandit Problem

Ensemble active learning for data streams is an emerging trend, showing high promise due to highly efficient label queries for newly arriving instances and higher robustness to concept drift. We assume that we have at our disposal a pool of L classifiers:

$$\Pi = \{\Psi_1, \dots, \Psi_L\}, \tag{8}$$

where each classifier in the pool fulfills the adaptability criteria over time (see Eq. 6) and their committee decision is used to decide on which instance should be queried.

We propose to achieve the adaptation over time by using a multi-armed bandit approach [Czarnecki and Podolak, 2014]. This allows us to conduct active learning in a different fashion than existing committee methods. While ensemble active learning is usually done via using multiple votes on whether new instance should be queried or not, the proposed approach selects dynamically the most competent classifier to be responsible for the query decision. Therefore, we are able to better utilize a pool of diverse classifiers in hope that at least one of them can anticipate the direction of changes in the stream better than the remaining ones. The idea behind this is similar to dynamic classifier selection [Almeida *et al.*, 2018; Zybiewski *et al.*, 2019], as we aim to exploit the unique competencies of base learners.

MAB assumes access to a finite number of processes (i.e., bandits) that are iteratively sampled (played) in order to obtain a value (reward r). MAB aims at maximizing the cumulative reward over time by observing the results of playing at a given machine. Of course, the issue related with MAB is the selection of a proper playing strategy. This can be seen as a minimization of regret function \mathcal{R} that can be formulated as a difference between the sum of obtained rewards using a selected strategy s and the sum of rewards obtained using a hypothetical optimal strategy:

$$\min_s \mathcal{R}_s = \sum_{k=1}^T r_k^{opt} - \sum_{k=1}^T r_k^s \iff \max_s \sum_{k=1}^T r_k^s, \quad (9)$$

where r_k^s is the reward after k -th iteration with s .

To use MAB for active learning from data streams, we consider each classifier in the pool as a separate machine and query them during playing (i.e., during data stream processing). After defining a reward function, we may use any existing MAB algorithm as our adaptation function f (see Eq. 7):

$$f((h_k)_{k=1}^i, (\Psi_k)_{k=1}^i, (q_k)_{k=1}^i, (\mathbf{o}(q_k))_{k=1}^i) := \text{MAB}((\Psi_k)_{k=1}^i, (r_k)_{k=1}^i), \quad (10)$$

where

$$r_i = r((h_k)_{k=1}^i, (\Psi_k)_{k=1}^i, (q_k)_{k=1}^i, (\mathbf{o}(q_k))_{k=1}^i). \quad (11)$$

As we deal with an ensemble of classifiers, we will denote the mean reward of j -th classifier as:

$$\bar{r}_i = \sum_{k \in P_l} \frac{r_k}{|P_l|}, \quad (12)$$

where $P_l = \{k : \Psi_k = \Psi_l\}$ is a set of iterations' indices where l -th classifier was used for label query.

3.3 MAB Strategy

In order to optimize our ensemble AL, we will use an efficient MAB algorithm. A thorough survey on MAB pointed out UCB1 as one of the most efficient approaches and thus we will focus on this strategy [Burtini *et al.*, 2015].

Upper Confidence Bound (UCB1) [Kuleshov and Precup, 2014] is a popular algorithm for solving MAB problems that

approaches the minimal regret bound of $\Omega(\log T)$ when the constant variance of each bandit (in our case classifier) is assumed:

$$b = \arg \max_{l \in \{1, \dots, L\}} \left(\bar{r}_l + \sqrt{\frac{2 \log T}{|P_l|}} \right). \quad (13)$$

While UCB1 has gained popularity in solving various MAB problems, we will argue that it is not feasible for ensemble active learning. The assumption of identical variance of each bandit (i.e., classifier) is too strong and unrealistic for any ensemble composition. Therefore, we must assume that our classifiers will have varying variances within the pool Ψ .

In order to capture that, we propose to use a tuned version of UCB1 that takes into account the variance of each bandit [Burtini *et al.*, 2015]:

$$b = \arg \max_{l \in \{1, \dots, L\}} \left(\bar{r}_l + \sqrt{\frac{\log T}{|P_l|} \min \left(\frac{1}{4}, \max_{k \in P_j} \text{var}(r_k) + \sqrt{\frac{2 \log T}{|P_l|}} \right)} \right). \quad (14)$$

3.4 Reward Function

Vast majority of existing active learning methods for data streams base their decision regarding label query for a new instance on the uncertainty of a classifier regarding that instance. Therefore, one should select instances for labeling that are closest to the current classification boundary and have a high chance of being incorrectly labeled by the classifier:

$$r((h_k)_{k=1}^i, (\Psi_k)_{k=1}^i, (q_k)_{k=1}^i, (\mathbf{o}(q_k))_{k=1}^i) = r_{0/1}(h_i, q_i, \mathbf{o}(q_i)) = \mathbf{1}_{h_i(q_i) \neq \mathbf{o}(q_i)}. \quad (15)$$

While this uncertainty-based approach for label query is used by many successful active learning algorithms dedicated to data streams, its usefulness degrades with the increasing presence of concept drift. This is caused by the fact that concept drift may appear in regions associated with a high classifier certainty - therefore these regions will never be queried for new instances, prohibiting the classifier from adapting to such a change. Therefore, in this paper we propose to explore an alternative way of selecting instances for label query by measuring the increase in generalization capabilities of the classifier according to a metric m on a separate validation set V :

$$r_m(h_i, h_{i-1}, V) = m(h_i(V), \mathbf{o}(V)) - m(h_{i-1}(V), \mathbf{o}(V)). \quad (16)$$

Using this approach for instance selection leads to improved theoretical bound on the learning process from data streams. We will prove this using the UCB1 strategy for proof derivation.

Theorem 1 *Given an ensemble of classifiers Π suitable for learning from data streams, MAB-based active learning with UCB1 strategy, and reward function r_m , assuming that each classifier is stochastic, the generalization capabilities of the*

proposed ensemble active learning approach for used metric m is at most $\mathcal{O}(\log T)$ worse than the one obtained from an optimal selection strategy after T iterations.

Proof. Following the similar proofs for UCB1 MAB bounds in non-streaming environments, this theorem is a consequence of UCB1 MAB bounds on reward function and because the following:

$$\begin{aligned} \sum_{k=1}^T r_{m_k} &= \sum_{k=1}^T [m(h_k(V), \mathbf{o}(V)) - m(h_{k-1}(V), \mathbf{o}(V))] \\ &= m(h_T(V), \mathbf{o}(V)) - m(h_0(V), \mathbf{o}(V)) \end{aligned} \quad (17)$$

is the generalization capability of an underlying classification model that depends only on the initial hypothesis h_0 and the final hypothesis h_T . Therefore, this realizes the theoretical lower bound of MAB error of $\Omega(\log T)$ and in general cannot be asymptotically improved. \square

3.5 Practical Considerations

From the practical point of view, our proposed Ensemble Active Learning with Multi-Armed Bandit (EAL-MAB) can be seen as a plug-in approach working with vast majority of modern streaming ensembles. However, there are few practical considerations to be taken into an account when using EAL-MAB with a selected ensemble algorithm:

Validation set. Used ensemble learning algorithm must be capable of evaluating the generalization metric (see Eq. 16) for each base classifier on instances unseen by this classifier. In practice, this is realized by most of streaming ensemble methods, as they either train base classifiers using instance subsets (e.g., Online Bagging [Oza and Russell, 2001]) or using different chunks of data.

Classifier outputs. EAL-MAB requires for the base classifiers in ensemble to return continuous outputs (e.g., support functions) and not discreet labels. In practice, this is realized by most of online / streaming single classifiers.

Usage of labeling budget. EAL-MAB runs on each new chunk of data for T iterations to select instances, one per iteration. Thus, the given budget B for a window size of ω is equal to the number of iterations that EAL-MAB will perform: $T = B \times \omega$.

Usage of metric m . EAL-MAB may use any metric suitable for data streams [Krawczyk *et al.*, 2017].

4 Experimental Study

This experimental study was designed to answer four research questions that will provide an insight into the proposed EAL-MAB approach:

- **RQ1:** How does EAL-MAB compares to state-of-the-art active learning methods dedicated to drifting data streams and is it capable of more efficient utilization of the provided budget?
- **RQ2:** How does EAL-MAB reacts to concept drift and how well is it able to improve the recovery after a drift occurred?

- **RQ3:** How does the AL procedure within EAL-MAB affects the underlying ensemble diversity?
- **RQ4:** How does EAL-MAB work with different ensemble architectures and base classifiers?

4.1 Data Stream Benchmarks

For the purpose of evaluating our proposed algorithm, we generated 10 diverse and large-scale data stream benchmarks using MOA environment [Bifet *et al.*, 2010a], as well as two popular real-world data streams. By using data stream generators, we were able to fully control the nature and occurrence of concept drifts, which in turn leads to a more explainable experimental study. By analyzing how the proposed method behaves in a controlled environment, we may gain more in-depth insight into its strong and weak points. Details of used data streams are given in Table 1.

Abbr.	Generator	Instances	Features	Classes	Drift
HYP _{IF}	Hyperplane	1 000 000	10	2	incremental-fast
HYP _{IS}	Hyperplane	1 000 000	10	2	incremental-slow
LED _M	LED	1 000 000	24	10	mixed
LED _S	LED	1 000 000	24	10	sudden
RBF _B	RBF	1 000 000	100	5	blips
RBF _G	RBF	1 000 000	40	20	gradual
RBF _{GR}	RBF	6 000 000	20	10	gradual-recurring
SEA _G	SEA	3 000 000	3	4	gradual
SEA _S	SEA	3 000 000	3	4	sudden
TRE _S	RandomTree	2 000 000	10	6	sudden
ACT	Activity	1 048 570	3	6	unknown
SEN	Intel Sensor	2 219 804	5	54	unknown

Table 1: Properties of used data stream benchmarks.

4.2 Set-up

Here, we will present the details of the experimental study design.

Reference AL algorithms. In order to provide a fair competition for our proposed algorithm, we have selected three state-of-the-art AL methods dedicated specifically to drifting data streams:

- **Randomized Variable Uncertainty (R-VAR)** [Zliobaite *et al.*, 2014] is an AL approach that monitors the uncertainty of a classifier and dynamically adjust the label query threshold. Additionally, a randomized component is added, allowing for periodical query of feature space subsets characterized by a high classifier certainty. This allows for detection of concept drifts taking place far from the decision boundary.
- **Stream-based Active Learning (SAL)** [Mohamad *et al.*, 2018b] uses non-parametric Bayesian models to cope with the lack of prior knowledge about the data stream, realized as Dirichlet mixture models and the stick breaking process. This allows for querying instances that are most likely to reduce the classification error.
- **Bi-Criteria Active Learning Algorithm for Dynamic Data Streams (BIAL)** [Mohamad *et al.*, 2018a] uses a combination of label uncertainty criterion and density-based criterion. The latter criterion dynamically weights

incoming instances according to their underlying distributions, in order to bias the sampling process.

Ensemble learning algorithm. All used AL algorithms are plug-in tools that can be combined with any classifier. Reference methods were mainly designed for single classifiers, while the proposed EAL-MAB can be used with almost any stream-based ensemble learning algorithm. For the purpose of this comparison, we use all of them in connection with popular Leveraging Bagging [Bifet *et al.*, 2010b], as it is widely recognized as one of the most effective ensembles for data streams [Krawczyk *et al.*, 2017]. All AL methods use exactly the same initial ensemble. We restrict the size of Leveraging Bagging to 10 base classifiers (as suggested by authors) and use Hoeffding Trees as base learners. Additionally, we examined Online Bagging [Oza and Russell, 2001] and Accuracy Updated Ensemble [Brzezinski and Stefanowski, 2014] architectures with both Hoeffding Tree and Naïve Bayes base classifiers to examine if the proposed EAL-MAB algorithm can efficiently work with any underlying streaming classification model.

AL budgets. As we want for our experimental study to reflect a real-world scenario, we investigate small to medium budgets, from highly restrictive 1% to 30%, so $B \in \{0.01, 0.05, 0.1, 0.15, 0.2, 0.25, 0.3\}$.

Evaluation metrics. For comparing streaming classifiers, we use a prequential accuracy metric [Gama *et al.*, 2013]. In order to evaluate the diversity of ensembles, we use a streaming version of kappa interrater agreement metric [Brzezinski and Stefanowski, 2016].

Windows. We use a window size $\omega = 1000$ for calculating the budgets, prequential metrics, and training new classifiers for ensembles. For evaluating the generalization capabilities of base classifiers in the ensemble (see Eq. 16), we use a window $\omega_V = 50$ most recently labeled instances.

Statistical analysis. To examine the significance of the obtained results, we used McNemar’s test [Bifet *et al.*, 2015] for pairwise comparison and Bayesian signed rank test [Benavoli *et al.*, 2017] for comparison over multiple datasets.

Analyzing the concept drift. As we generated all 10 benchmark data streams, we have full information about the exact point when concept drift takes place. Therefore, we label these instances as drifting ones and monitor how many of them will be selected by evaluated AL algorithms for labeling.

4.3 Results and Discussion

Let us discuss the obtained results from the perspective of stated three research questions that guided this study.

Influence of AL algorithms on classification accuracy (RQ1). Figure 1 presents the prequential accuracies of examined AL methods with respect to varying budget sizes and benchmarks. Additionally, Figure 3 depicts the outcomes of Bayesian statistical analysis of result significance. In 10 out of 12 cases, the proposed EAL-MAB approach was capable of significantly outperforming all of AL reference methods,

Dataset	R-VAR	SAL	BIAL	EAL-MAB
HYP _{IF}	17.23±5.21	19.54±4.12	20.46±4.51	26.12 ±3.18
HYP _{IS}	18.65±4.26	22.54±3.95	21.89±4.26	28.81 ±3.52
LED _M	32.73±2.19	38.45±3.11	39.99±3.82	43.26 ±3.18
LED _S	27.41±1.86	29.45±2.11	29.88±3.28	33.47 ±1.68
RBF _B	21.09±2.76	24.98±2.98	29.72 ±3.07	26.54±3.01
RBF _G	36.44±4.98	38.72±6.11	40.07±5.28	45.28 ±5.39
RBF _{GR}	38.56±6.21	40.03±7.01	41.13±6.38	47.20 ±6.94
SEA _G	11.87±3.98	17.43±2.51	18.82 ±2.99	15.82±2.32
SEA _S	10.02±7.32	15.77±6.21	16.61±5.84	25.06 ±5.11
TRE _S	38.23±4.98	31.44±2.66	32.80±2.29	43.19 ±3.36
ACT	–	–	–	–
SEN	–	–	–	–

Table 2: Percentage of instances labeled as drifting ones that were selected for label query by AL algorithms.

showing its capabilities for selecting more useful instances. This can be contributed to efficiently using the ensemble of classifiers for label query, which leads to reduced variance and improved evaluation on which instance may be useful not only for base classifiers, but for the ensemble as a whole. Additionally, reference methods are all based on classifier uncertainty, while the proposed EAL-MAB uses the generalization capabilities of base classifiers for selecting instances. This proves that stream-based AL should evaluate different selection criteria, as uncertainty itself loses information value in the presence of concept drift. This is especially limiting when drifting instances appear in regions with associated high classification certainty. While for high budgets BIAL reference method achieves similar performance to EAL-MAB, we should focus on realistic and small budgets. Here, EAL-MAB is capable of much better usage of the limited class labels, leading to much more stable and effective classification.

Capabilities of AL algorithms to adapt to concept drift (RQ2). As a second part of our experiment, we want to evaluate the capabilities of examined AL algorithms to handle concept drifts. We propose to realize this by checking how many instances labeled as drifting ones have been selected for label query. A good AL algorithm should select as many instances from the new concept as possible, leading to faster recovery and adaptation to the evolved environment. Table 2 presents the percentage of drifting instances that have been selected by examined AL algorithms. Here, we see that for 8 benchmarks on which EAL-MAB obtained the best performance, it also selected the highest number of drifting instances for labeling. This shows a relationship between the obtained predictive power and speed of adaptation to concept drifts. By using adaptive MAB approach with classifier generalization criterion, we are able to select most valuable instances during the drift.

Influence of AL algorithms on the ensemble diversity (RQ3). Let us check how the used AL algorithms influence the underlying ensemble diversity. Diversity metrics for non-streaming environments have been criticized for not being an efficient tool for ensemble forming and for not actually grasping the idea of what makes a good or bad diversity [Brown and Kuncheva, 2010]. However, in the context of data streams the situation has recently changed. Diversity

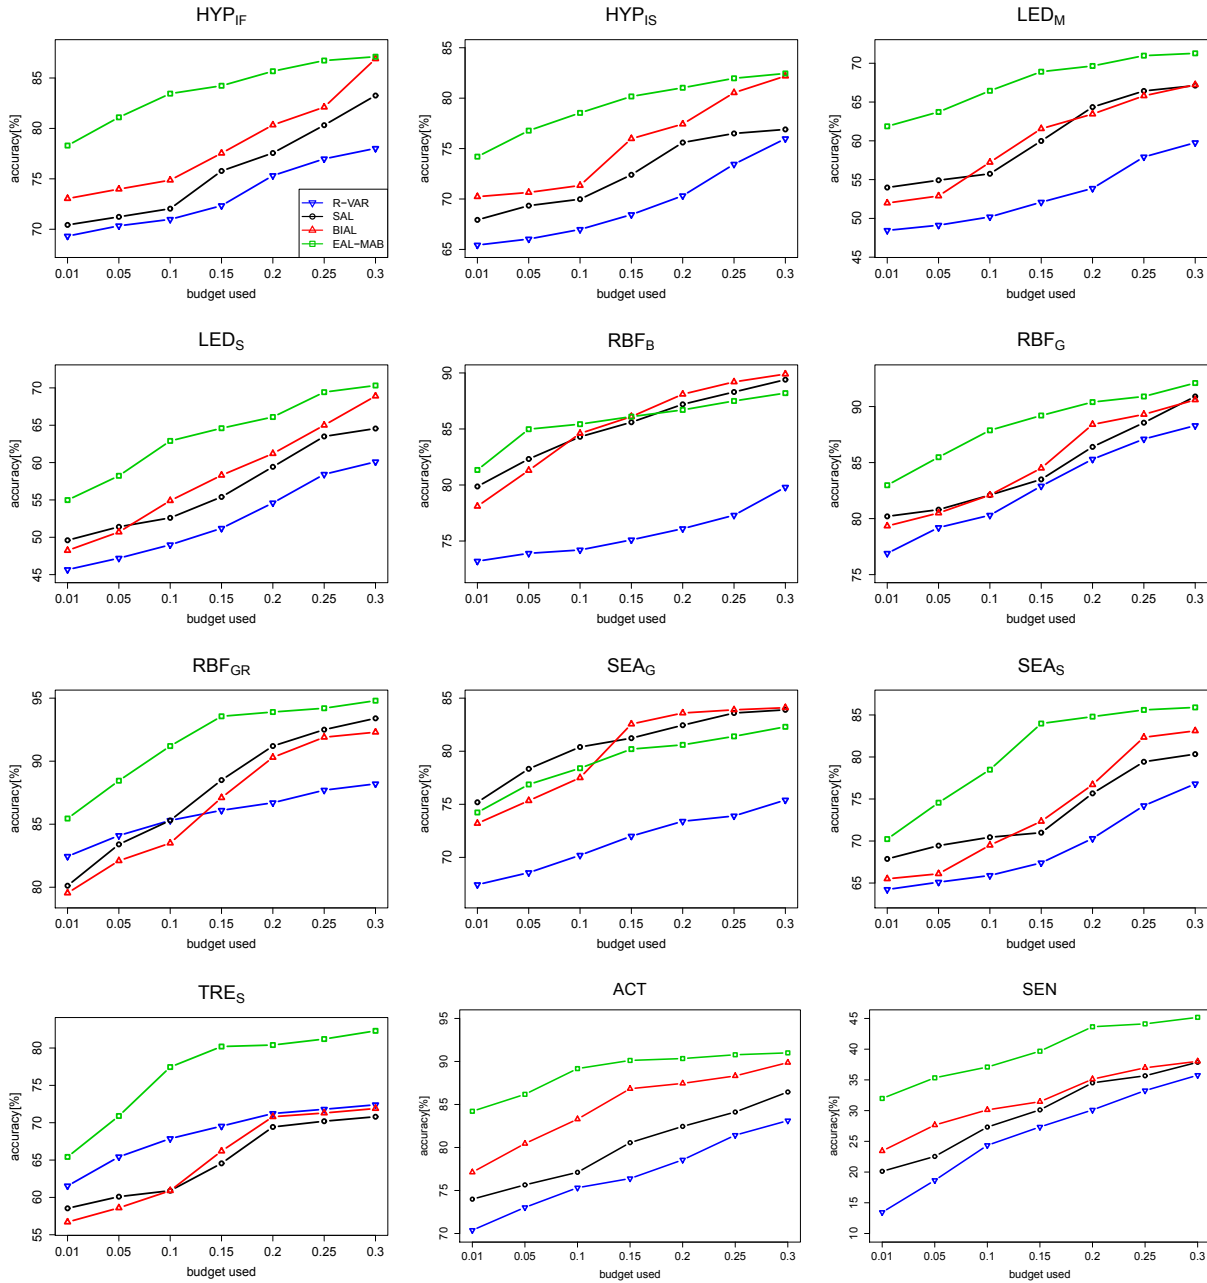


Figure 1: Averaged sequential accuracy of evaluated AL algorithms with respect to varying budget sizes.

has been identified as a crucial factor for ensemble adaptation to concept drift, as it allows for anticipating the direction of changes [Minku *et al.*, 2010]. Therefore, a good AL method should be (directly or indirectly) managing the ensemble diversity. Although there is a lack of success stories of applying diversity measures for ensemble forming from data streams [Krawczyk *et al.*, 2017], the same measures are useful for monitoring the ensemble performance over time [Brzezinski and Stefanowski, 2016]. Figure 2 presents the averaged diversity for each examined budget and benchmark. Here, we can see that none of the reference AL methods can manage

the diversity properly, especially for small budgets - as instances selected there are too similar to each other (they all are selected by classifier uncertainty metric). The proposed EAL-MAB is capable of much better diversity management, improving it significantly. This is especially visible in case of small budgets (0.01 - 0.1), where such enhanced diversity directly translates to improved predictive power and drift adaptation capabilities. EAL-MAB is capable of selecting useful and mutually complementary instances even under very sparse access to class labels. Please note that one cannot say what value of diversity metric leads to an improved ensemble

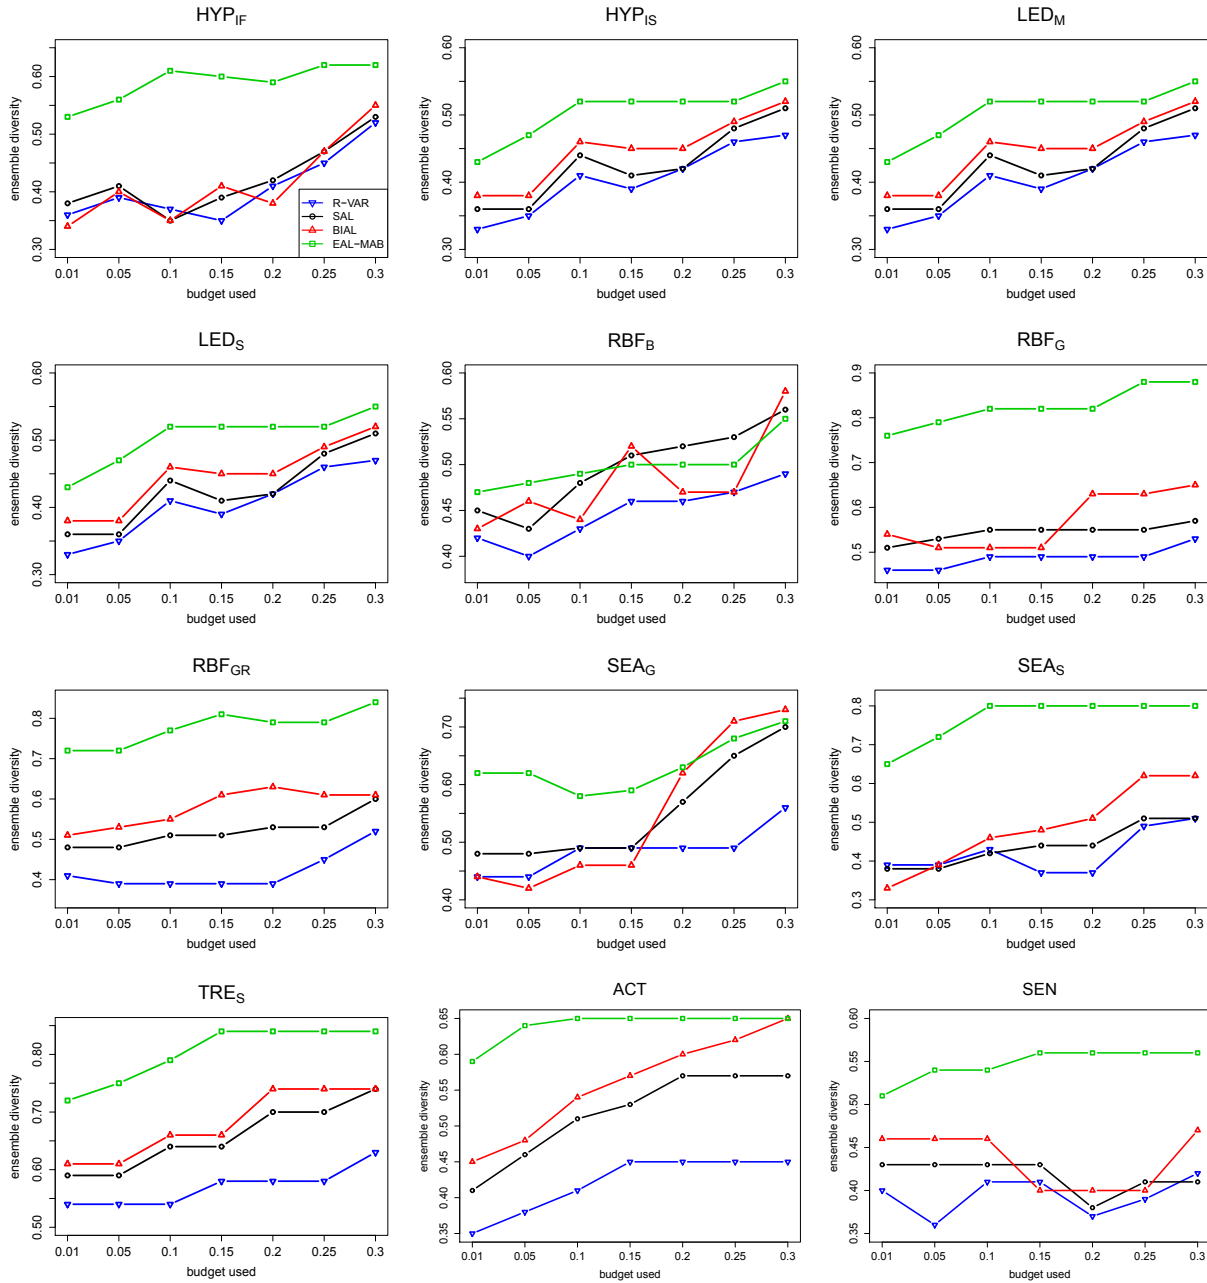


Figure 2: Diversity of ensembles measured with kappa interrater agreement metric with respect to varying budget sizes.

ble performance. However, what is important is to observe that the diversity of ensemble reacts to even small changes in the number of provided instances (i.e., for small budget increases). That means that increased budget was used to choose complementary instances that introduce new knowledge to base classifiers and that these classifiers do not converge to a similar model. As we can see, for three reference AL algorithms, the diversity metric is stable for small to medium budget values, showing that the new instances either do not introduce new knowledge, or that they were supplied to all base classifiers at the same time. This is not the

case with the EAL-MAB, where each classifier takes an individual part in the instance selection. Thus, it is possible to select instances that will improve the local (for a given base classifier) or global performance (of the entire ensemble).

Influence of different ensemble architectures and base classifiers (RQ4). Finally, let us analyze how does EAL-MAB relies on the underlying ensemble model and base classifier. This is necessary to examine, as EAL-MAB aims to be a flexible plug-in AL method that can work with any ensemble scheme that follows the practical considerations dis-




Figure 3: Visualization of sample posterior distribution from correlated Bayesian sign test for comparison with (left) R-VAR, (center) SAL, and (right) BIAL. A higher concentration of points in the right region of the plot corresponds to EAL-MAB outperforming the reference method.




Figure 4: Comparison of EAL-MAB and reference active learning algorithms over different ensemble architectures and base classifiers. Results presented with respect to the number of wins (green), ties (yellow), and losses (red) over 84 cases (12 benchmark datasets and 7 different budgets). A tie was considered when McNemar’s test rejected the significance of difference between tested algorithms.

cussed in Section 3.5. Figure 4 depicts the results over three different ensemble models with two different base classifiers. Results are presented according to pairwise McNemar’s test for statistical significance over 84 test cases (12 benchmarks with 7 budget sizes). We can see that while there is a small variance in results, the general trends are clearly preserved with EAL-MAB outperforming all three reference AL algorithms. This shows the flexibility of the proposed AL framework and proves that it can be used as an efficient plug-in tool for making ensembles work under labeling constraints for drifting data stream mining.

5 Conclusions and Future Works

In this paper, we have discussed the issue of learning from drifting data streams under a restricted access to class labels. To address this challenging, real-world problem, we proposed a new ensemble-based active learning scheme. We treated base classifiers as bandits and realized the instance query procedure as a Multi-Armed Bandit problem. This allowed us to efficiently use the pool of classifiers to select such instances that will improve both individual learners and the

ensemble as a whole. We proposed to guide the selection process not by a popular uncertainty measure, but by measuring how each new instance influences the generalization capabilities of each classifier. This has led to a much better selection tool that is capable of handling drifting data streams.

Obtained results encourage us to pursue further works in this directions. We plan to investigate novel measures for selecting instances that are based on local instance hardness, as well as adapt our approach to binary and multi-class imbalanced and drifting data streams.

References

[Alabdulrahman *et al.*, 2016] Rabaa Alabdulrahman, Herna Viktor, and Eric Paquet. An active learning approach for ensemble-based data stream mining. In *KDIR*, pages 275–282. SciTePress, 2016.

[Almeida *et al.*, 2018] Paulo R. L. Almeida, Luiz S. Oliveira, Alceu S. Britto Jr., and Robert Sabourin. Adapting dynamic classifier selection for concept drift. *Expert Syst. Appl.*, 104:67–85, 2018.

- [Benavoli *et al.*, 2017] Alessio Benavoli, Giorgio Corani, Janez Demsar, and Marco Zaffalon. Time for a change: a tutorial for comparing multiple classifiers through bayesian analysis. *Journal of Machine Learning Research*, 18:77:1–77:36, 2017.
- [Bifet *et al.*, 2010a] Albert Bifet, Geoff Holmes, Richard Kirkby, and Bernhard Pfahringer. MOA: massive online analysis. *Journal of Machine Learning Research*, 11:1601–1604, 2010.
- [Bifet *et al.*, 2010b] Albert Bifet, Geoffrey Holmes, and Bernhard Pfahringer. 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*, pages 135–150, 2010.
- [Bifet *et al.*, 2015] Albert Bifet, Gianmarco De Francisci Morales, Jesse Read, Geoff Holmes, and Bernhard Pfahringer. Efficient online evaluation of big data stream classifiers. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, pages 59–68, 2015.
- [Brown and Kuncheva, 2010] Gavin Brown and Ludmila I. Kuncheva. "good" and "bad" diversity in majority vote ensembles. In *MCS*, volume 5997 of *Lecture Notes in Computer Science*, pages 124–133. Springer, 2010.
- [Brzezinski and Stefanowski, 2014] Dariusz Brzezinski and Jerzy Stefanowski. Reacting to different types of concept drift: The accuracy updated ensemble algorithm. *IEEE Trans. Neural Netw. Learning Syst.*, 25(1):81–94, 2014.
- [Brzezinski and Stefanowski, 2016] Dariusz Brzezinski and Jerzy Stefanowski. Ensemble diversity in evolving data streams. In *Discovery Science - 19th International Conference, DS 2016, Bari, Italy, October 19-21, 2016, Proceedings*, pages 229–244, 2016.
- [Burtini *et al.*, 2015] Giuseppe Burtini, Jason Loepky, and Ramon Lawrence. A survey of online experiment design with the stochastic multi-armed bandit. *CoRR*, abs/1510.00757, 2015.
- [Czarnecki and Podolak, 2014] Wojciech M. Czarnecki and Igor T. Podolak. Adaptive active learning as a multi-armed bandit problem. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 989–990, 2014.
- [Dyer *et al.*, 2014] Karl B. Dyer, Robert Capo, and Robi Polikar. COMPOSE: A semisupervised learning framework for initially labeled nonstationary streaming data. *IEEE Trans. Neural Netw. Learning Syst.*, 25(1):12–26, 2014.
- [Gama *et al.*, 2013] João Gama, Raquel Sebastião, and Pedro Pereira Rodrigues. On evaluating stream learning algorithms. *Machine Learning*, 90(3):317–346, 2013.
- [Gama *et al.*, 2014] João Gama, Indre Zliobaite, Albert Bifet, Mykola Pechenizkiy, and Abdelhamid Bouchachia. A survey on concept drift adaptation. *ACM Comput. Surv.*, 46(4):44:1–44:37, 2014.
- [Krawczyk *et al.*, 2017] Bartosz Krawczyk, Leandro L. Minku, João Gama, Jerzy Stefanowski, and Michal Woźniak. Ensemble learning for data stream analysis: A survey. *Information Fusion*, 37:132–156, 2017.
- [Krawczyk *et al.*, 2018] Bartosz Krawczyk, Bernhard Pfahringer, and Michal Wozniak. Combining active learning with concept drift detection for data stream mining. In *IEEE International Conference on Big Data, Big Data 2018, Seattle, WA, USA, December 10-13, 2018*, pages 2239–2244, 2018.
- [Ksieniewicz *et al.*, 2019] Pawel Ksieniewicz, Michal Wozniak, Boguslaw Cyganek, Andrzej Kasprzak, and Krzysztof Walkowiak. Data stream classification using active learned neural networks. *Neurocomputing*, 353:74–82, 2019.
- [Kuleshov and Precup, 2014] Volodymyr Kuleshov and Doina Precup. Algorithms for multi-armed bandit problems. *CoRR*, abs/1402.6028, 2014.
- [Lughofer, 2017] Edwin Lughofer. On-line active learning: A new paradigm to improve practical useability of data stream modeling methods. *Inf. Sci.*, 415:356–376, 2017.
- [Minku *et al.*, 2010] Leandro L. Minku, Allan P. White, and Xin Yao. The impact of diversity on online ensemble learning in the presence of concept drift. *IEEE Trans. Knowl. Data Eng.*, 22(5):730–742, 2010.
- [Mohamad *et al.*, 2018a] Saad Mohamad, Abdelhamid Bouchachia, and Moamar Sayed Mouchaweh. A bi-criteria active learning algorithm for dynamic data streams. *IEEE Trans. Neural Netw. Learning Syst.*, 29(1):74–86, 2018.
- [Mohamad *et al.*, 2018b] Saad Mohamad, Moamar Sayed Mouchaweh, and Abdelhamid Bouchachia. Active learning for classifying data streams with unknown number of classes. *Neural Networks*, 98:1–15, 2018.
- [Oza and Russell, 2001] Nikunj C. Oza and Stuart J. Russell. Online bagging and boosting. In *Proceedings of the Eighth International Workshop on Artificial Intelligence and Statistics, AISTATS 2001, Key West, Florida, USA, January 4-7, 2001*, 2001.
- [Xu *et al.*, 2016] Wenhua Xu, Fengfei Zhao, and Zhengcai Lu. Active learning over evolving data streams using paired ensemble framework. In *ICACI*, pages 180–185. IEEE, 2016.
- [Zliobaite *et al.*, 2014] Indre Zliobaite, Albert Bifet, Bernhard Pfahringer, and Geoffrey Holmes. Active learning with drifting streaming data. *IEEE Trans. Neural Netw. Learning Syst.*, 25(1):27–39, 2014.
- [Zyblewski *et al.*, 2019] Pawel Zyblewski, Pawel Ksieniewicz, and Michal Wozniak. Classifier selection for highly imbalanced data streams with minority driven ensemble. In *Artificial Intelligence and Soft Computing - 18th International Conference, ICAISC 2019, Zakopane, Poland, June 16-20, 2019, Proceedings, Part I*, pages 626–635, 2019.