

EXPERIMENTS WITH TWO APPROACHES FOR TRACKING DRIFTING CONCEPTS

Ivan Koychev

ABSTRACT. This paper addresses the task of learning classifiers from streams of labelled data. In this case we can face the problem that the underlying concepts can change over time. The paper studies two mechanisms developed for dealing with changing concepts. Both are based on the time window idea. The first one forgets gradually, by assigning to the examples weight that gradually decreases over time. The second one uses a statistical test to detect changes in concept and then optimizes the size of the time window, aiming to maximise the classification accuracy on the new examples. Both methods are general in nature and can be used with any learning algorithm. The objectives of the conducted experiments were to compare the mechanisms and explore whether they can be combined to achieve a synergetic effect. Results from experiments with three basic learning algorithms (kNN, ID3 and NBC) using four datasets are reported and discussed.

1. Introduction. Many machine learning applications employ algorithms for learning classifiers from examples. In those applications, new examples are obtained continuously and added to the training dataset. Then the classifier

ACM Computing Classification System (1998): I.2.6.

Key words: Machine Learning; Concept Drift; Forgetting Models

can be updated on the new set of examples, aiming to improve the classification accuracy using a larger training set. However, these applications often face the problem that real life concepts tend to change over time, i.e. a classifier learned on all previous examples will decrease its accuracy on the new data. Hence, some of the old observations that are out-of-date have to be ‘forgotten’. This problem is known as *concept drift* [18]. A promising example are the systems that learn from observations about user’s interests, aiming to help him/her to find items which are expected to be of interest to him/her. Usually, the user’s interests and preferences are inclined to change over time (e.g. [11], [14] and [15]). Another promising example is a typical data mining task – knowledge discovery in a market basket, where there has been also observed that the customers’ preferences tend to change over time and the association rules mined on historical data are likely to be out-of-date [19].

A number of forgetting mechanisms have been developed to cope with this problem. They can be divided in two major types depending on whether they are able to *forget only* or whether they are also able to keep old data/knowledge and *remember* them if they seem to be useful again.

The first ones have the advantage to being simpler, faster, and there is no need to maintain and search a large storage with old data or knowledge that could be useful. It performs well with concepts that do not change often and change gradually rather than abruptly, but it performs very poorly on frequently changing and recurring concepts.

The second ones have a clear advantage in case of recurring concepts (e.g. caused by season changes), where the old data or acquired knowledge can be very helpful (e.g. [11] and [15]). But it requires larger storage space for the old information and extra time for retrieval of the old relevant episodes of data or knowledge.

This study focuses on the first type of mechanisms. There are two basic approaches suggested to improve the performance of the pure time window. The first one, instead of considering all examples in the time window as equally important, provides *the examples with weights which decrease over time*. Thus it simulates a forgetting that is gradual, which seems to be a more natural one. The second one *adapts dynamically the size of the time window* according to the current concept behaviour. The paper reports results from experiments with two forgetting mechanisms, in particular the Gradual Forgetting as it is described in [10] and Time Window Optimization presented in [8]. The aims are to compare the approaches and to explore the possibility to combine them for further improvement of the performance.

The next section gives a short overview of the related work. The compared mechanisms are described in section 3. The research questions and hypotheses are raised in section 4. The experiment design is introduced in section 5. The results from the experiments are reported and discussed in section 6.

2. Related work. Different approaches have been developed to track changing (also known as shifting, drifting or evolving) concepts. Typically it is assumed that if the concept changes, then the old examples become irrelevant to the current period. The concept descriptions are learned from a collection of most recent examples called a time window. For example, Mitchell et al. [15] developed a software assistant for scheduling meetings, which employs machine learning to acquire assumptions about individual habits of arranging meetings, uses a time window to adapt faster to the changing preferences of the user. Widmer and Kubat [21] developed the first approach that uses adaptive window size. The algorithm monitors the learning process and if the performance drops under a predefined threshold it uses heuristics to adapt the time window size dynamically. Maloof and Michalski [14] have developed a method for selecting training examples for a partial memory learning system. The method uses a time-based function to provide each instance with an age. Examples that are older than a certain age are removed from the partial memory. Delany et al. [3] employ a case base editing approach that removes noise cases (i.e. the cases that contribute to the classification incorrectly are removed) to deal with concept drift in a spam filtering system. The approach is very promising, but is applicable for lazy learning algorithms only. To manage the problems with gradual concept drift and noisy data, the approach in [13] suggests the use of three windows: a small (with fixed size), a medium and a large (dynamically adapted by simple heuristics).

The above approaches totally forget the examples that are outside the given window, or older than a certain age. The examples which remain in the partial memory are equally important for the learning algorithms. This is an abrupt forgetting of old examples, which probably does not reflect their rather gradual ageing. To deal with this problem it was suggested to weight training examples in the time window according to their appearance over time [10]. These weights make recent examples more important than older ones, essentially causing the system to gradually forget old examples. This approach has been explored further in [4], [7] and [12].

Some systems use different approaches to avoid loss of useful knowledge

learned from old examples. The CAP system [15] keeps old rules as long as they are competitive with the new ones. The FLORA system [21] also maintains a store of concept descriptions relevant to previous contexts. When the learner suspects a context change, it will examine the potential of previously stored descriptions to provide better classification. The COPL approach [11] employs a two-level schema to deal with changing and recurring concepts. On the first level the system learns a classifier from a small set of the most recent examples. The learned classifier is accurate enough to be able to distinguish the past episodes that are relevant to the current context. Then the algorithm constructs a new training set, ‘remembers’ relevant examples and ‘forgets’ irrelevant ones. As we explained in the introduction, the approach explored in this paper do not assume that old examples or models can be retrieved.

Widmer [20] assumes that the domain provides explicit clues to the current context (e.g. attributes with characteristic values). A two-level learning algorithm is presented that effectively adjusts to changing contexts by trying to detect (via meta-learning) contextual clues and, using this information to focus the learning process. Another two-level learning algorithm assumes that concepts are likely to be stable for some period of time [6]. This approach uses contextual clustering to detect stable concepts and to extract hidden context. The mechanisms studied in this paper do not assume that the domain provides some clues that can be discovered using a meta-learning level. They rather aim to get the best performance using a single learning level.

An adaptive boosting method based on dynamic sample-weighting is presented in [2]. This approach uses statistical hypothesis testing to detect concept changes. Gama et al. [5] also use a hypothesis testing procedure, similar to that used in control charts, to detect the concept drift, calculating on all of the data so far. The mechanism gives a warning at 2 standard deviations (approximately 95%) and then takes action at 3 standard deviations (approximately 99.7%). If the action level is reached then the start of the window is reset to the point at which the warning level was reached. However, for this mechanism, it can take quite a long time to react to changes and the examples that belong to the old concept are not always completely useless, especially when the concept drift is rather gradual.

To adapt the size of the window according to current changes in the concept, the approach presented in [7] uses a naïve optimisation approach, which tries all possible window sizes and selects the one with the smallest error rate. This work provides interesting results from experiments with different forgetting mechanisms applied for Support Vector Machines classifier, using a textual data

corpus.

To detect concept changes the TWO approach [8] uses a statistical test on selection population from the time window, which excludes its beginning and end. If a concept drift is detected an efficient optimisation algorithm is employed to detect the optimal window size. This approach is explained in more detail in the next section and studied in the conducted experiments.

3. Two approaches for tracking drifting concepts. Let us consider a sequence of examples. Each example is classified according to underlying criteria into one of a set of classes, which forms the training set. The task is to learn a classifier that can be used to classify the next examples in the sequence. However, the underlying criteria can subsequently change and the same example can be classified differently according to the time of its appearance, i.e. a concept drift takes place. As we discussed above to deal with this problem machine learning systems often use a time window, i.e., the classifier is not learned on all examples, but only on a subset of recent examples. The next section describes forgetting mechanisms that further develop this idea.

3.1. Gradual forgetting. This section describes a gradual forgetting mechanism earlier introduced in [10]. Just as the time window approach it assumes that when a concept tends to drift the newest observations better represent the current concept. Additionally, it aims to make the forgetting of old examples gradual. Let us define a gradual forgetting function $w = f(t)$, which provides weights for each instance according to its location in the course of time. The weights assign higher importance value to the recent examples. Earlier, Widmer [20] suggests an “exemplar weighting” mechanism, which is used for the IBL algorithm in METAL(IB), however it is not exploited for NBC in METAL(B). The researcher in area of boosting also faced the need of weighting examples. There are two ways, in which boosting employs the weights. The first one is known as boosting *by sampling* – the examples are drawn with replacement from the training set with a probability proportional to their weights. However, this approach requires a pre-processing stage where a new set of training examples should be generated. The better the sampling approximates the weights, the larger is the new training set. The second method, boosting *by weighting*, is used with learning algorithms that accept weighted training examples directly. In this case the weight is constrained as follows:

$$(1) \quad w_i \geq 0 \quad \text{and} \quad \sum_{i=1}^n w'_i = 1,$$

where n is the size of the training set. Most of the basic learning algorithms are designed to treat all examples as equally important. For kNN it can be easily implemented by multiplying the calculated distances with the weights of the examples. For other algorithms it does not look so straightforward. When there are no weights (i.e. all weights are the same) the formula (1) will provide weights $\forall w_i = \frac{1}{n}$. However, as the weights are utilized by multiplication, it seems to be better if we have $\forall w_i = 1$ in this boundary case. If we substitute that $w_i = nw'_i$ then the constraints in equation (1) will be transformed as follows:

$$(2) \quad w_i \geq 0 \quad \text{and} \quad \frac{\sum_{i=1}^n w_i}{n} = 1$$

Weights that obey the constraints (2), can be easily used in almost any learning algorithm requiring minor changes in the code i.e. every time when the algorithm counts an example it should be multiplied by its weight.

Various functions that model the process of forgetting and satisfy constraints (2) can be defined. For example, the following linear gradual forgetting function has been defined:

$$(3) \quad w_i = -\frac{2k}{n-1}(i-1) + 1 + k,$$

where: i is a counter of observations starting from the most recent observation and it goes back in time (as the function is forgetting $w_i \geq w_{i+1}$); $k \in [0, 1]$ is a parameter that represents the percent of decreasing the weight of the first and respectively increasing the weight of the last observation in comparison to the average. By varying k the slope of the forgetting function can be adjusted to reach better predictive accuracy.

The presented gradual forgetting mechanism is naturally integrated into a time window, by weighting the examples in it (i.e. $n = l$, where l is the size of the time window). The system will forget gradually inside the window. When $k = 0$ then all weights will be equal to 1, which means that we will have a “standard”

time window. When k is approaching 1 then the weights of the examples in the end on the window approach 0 and there is not an abrupt forgetting after the window's end.

3.2. Time window optimisation. This section presents an approach that learns an up-to-date classifier on drifting concept by dynamic optimisation of the time window size to gain maximum accuracy of prediction [8].

In case of dynamical adaptation of window size there are two important questions that have to be addressed in case of concept drift. The first one is *how to detect a change in the concept?* We are assuming that there is no background information about the process and we use the decrease in predictive accuracy of the learned classifier as an indicator of changes in the concept. Usually the developed detection mechanism uses a predefined threshold tailored for the particular dataset. However the underlying concept can change with different speeds and extend. To detect the changes the approach uses a statistical hypothesis test, which adapts the thresholds according to the recently observed concept deviation. The second important question is *how to adapt if a change is detected?* Other approaches use heuristics to decrease the size of the time window when changes in the concept are detected (e.g. [4] and [21]). This approach employs a fast 1-D optimisation to get the optimal size of the time window if a concept drift is detected. The next two subsections describe the algorithm in more details.

3.2.1. Detecting the concept drift. To detect concept changes the approach monitors the performance of the algorithm. For the presentation below we choose to observe the classification accuracy as a measure of the algorithm performance, but other measures, such as error rate or precision could have been used. The approach process the sequence of examples on small episodes (a.k.a. batches). On each step, a new classifier is learned from the current time window then the average accuracy of this classifier is calculated on the next batch of examples [15]. Then the presented approach suggests using a statistical test to check whether the accuracy of classification has significantly decreased compared with its historical level. If the average prediction accuracy of the last batch is significantly less than the average accuracy for the population of batches defined on the current time window, then a concept drift can be assumed.

The significance test is done using a population from the time window that does not include the first one or a few batches from the beginning of the time window, because the predictive accuracy can be low caused by a previous concept drift. Also one or a few most recent batches are not included in the test

window, because if the drift is gradual the accuracy will be dropping slowly. Such a test that uses a test population from the core of the time window works well for both abrupt and gradual drift.

The confidence level for the significance test should be sensitive enough to discover concept drift as soon as possible, but not to mistake noise for changes in the concept. The experience from the conducted experiments shows that the “standard” confidence level of 95% works very well in all experiments. This drift detection level is rather sensitive and it assists the algorithm to detect the drift earlier. If a false concept drift alarm is triggered, it will activate the window optimising mechanism, but in practice, this only results in an insignificant decrease in the time window size.

*This mechanism works as follows: **If** concept drift is detected **then** the optimisation of the size of the time window is performed (see the next section) **otherwise**, the time window size is increased to include the new examples.*

3.2.2. Optimising the time window size. In general, if the concept is stable, the bigger the training set is (the time window), the more accurately classifier can be learned. However when the concept is changing, a big window will probably contain a lot of old examples, which will result in a decrease of the classification accuracy. Hence, the window size should be decreased to exclude the out-of-date examples and in this way to learn a more accurate classifier. But if the size of the window becomes too small, it will also lead to a decrease in accuracy. The shape of curve that demonstrates the relationship between the size of the time window and the accuracy of the classification is shown in Figure 1.

To adapt the size of the window according to current changes in the concept, the presented mechanism suggests using the Golden Section algorithm for one-dimensional optimization. The algorithm looks for an optimal solution in a closed and bounded interval $[a, b]$ – in our case the possible window sizes $X = [x_{\min}, x_c]$, where x_{\min} is a predefined minimum size of the window and x_c is the current size of the time window. It assumes that the function $f(x)$ is unimodal on X (i.e. there is only one max x^*) and it is strictly increasing on (x_{\min}, x^*) and strictly decreasing on (x^*, x_c) , which is the shape that can be seen in Figure 1. In our case the function $f(x)$ calculates the classification accuracy of the learned model using a time window with size x .

The basic idea of this algorithm is to minimize the number of function evaluations by trapping the optimum solution in a set of nested intervals. On each step the algorithm uses the golden section ($\tau = 0.618$) to split the interval into three subintervals, as shown in Figure 1, where $l = b - \tau(b - a)$ and $r = a + \tau(b - a)$.

If $f(l) > f(r)$ then the new interval chosen for the next step is $[a, r]$ else $[l, b]$. The length of the interval for the next iteration is $\tau(b - a)$. Those iterations continue until the interval containing the maximum reaches a predefined minimum size. x^* is taken to lie at the centre of the final interval.

The Golden Section algorithm is a very efficient way to trap the x^* that optimizes the function $f(x)$. After n iterations, the interval is reduced to 0.618^n times its original size. For example if $n = 10$, less than 1% of the original interval remains. Note that, due to the properties of the golden section, each iteration requires only one new function evaluation.




Fig. 1. A sample shape of the correlation between the window size and accuracy of the learned classifier

In conclusion, if we can assume that the classification accuracy in relation to the time window is a unimodal function then the golden section algorithm can be used as an efficient way to find the optimal size of the time window. It is possible to find datasets for which the unimodal assumption is not true – e.g. when the concept changes very often and abruptly. In such cases, we can use other optimization methods that do not assume a unimodal distribution, however they are much more expensive in time. The trade-off that we have to take into consideration is to accept that we can occasionally be trapped in a local maximum, but have a fast optimization; or find a global maximum, but have significantly slower optimization.

4. Research questions and hypothesis. The aim of the conducted experiments is twofold:

- To compare the performance of the above mechanisms;

- To explore whether the present forgetting mechanisms can be combined to achieve a synergetic effect (i.e. improving performance).

In the conducted experiments we decided to measure the performance of the algorithms in classification accuracy on a test set that is formed by the next batch in the sequence. It could be measured using other statistics such as error rate, precision, etc.

The experiments conducted in [8] provided strong evidence that the Time Window Optimisation mechanism is able to improve the prediction accuracy significantly in comparison with Fixed-size Time Window. The gradual forgetting approach was developed earlier [9], [10] and was originally tested on an artificial dataset and data from a recommender system. It was further explored in [4], [7], etc., but there are not extensive tests with different algorithms and data sets on whether it is able to improve significantly the predictive accuracy in comparison with the pure Fixed-size Time Window. Therefore the first hypothesis that we want to test is:

Hypothesis 1: *Fixed-size Time Window with Gradual Forgetting is better than pure Fixed-size Time Window, measured in classification accuracy.*

It was also interesting to compare both forgetting mechanisms. Since the Time Window Optimisation mechanism comes across more promising than the approaches with fixed parameters the second hypothesis was formulated as follows:

Hypothesis 2: *Time Window Optimisation is better than Fixed-size Time Window with Gradual Forgetting measured in classification accuracy.*

Furthermore it can be supposed that if both mechanisms are combined (i.e. adaptive time window with gradual forgetting inside it) a synergetic effect will appear. Thus if Hypothesis 2 is true then the third hypothesis was formulated as follows:

Hypothesis 3: *Time Window Optimisation with Gradual Forgetting is better than Time Window Optimisation measured in classification accuracy.*

In case Hypothesis 2 fails then Hypothesis 3 will be reformulated to take into account the results from the experiments.

5. Experiment design. All experiments were designed to run iteratively, in this way simulating the process of mechanisms’ utilisation. For this reason the data streams were chunked on episodes/batched. On each iteration, a concept description is learned from the examples in the current time window. Then the learned classifier is tested on the next batch (see Figure 2).




Fig. 2. Iterations’ cross-validation design

To facilitate the presentation below the basic forgetting mechanisms are abbreviated as follows:

- FTW – Fixed-size Time Window,
- GF – Gradual Forgetting,
- TWO – Time Window Optimisation.

For each data set, the window size for the FTW was chosen to approximate the average time window size obtained in the experiments with the TWO mechanism on the same dataset. This is extra help for the FTW that would not be available in a real situation where the forthcoming sequence of events is unknown. The aim here is to allow the FTW approach to show its best performance.

Each hypothesis was tested on four experimental datasets using three basic Machine Learning algorithms. The rest of this section describes them in more details. The results from the experiments are reported and discussed in the next subsection.

5.1. Datasets. This subsection explains how the data streams used in the experiments were generated.

5.1.1. STAGGER problem. The first experiments were conducted with an artificial learning problem that was defined and used by Schlimmer and Granger [18] for testing STAGGER, probably the first concept drift tracking system. Much of the subsequent work dedicated to this problem used this dataset for testing purposes (e.g. [1], [4], [6], [10], [11], [14], [19] and [21]).

The STAGGER problem is defined as follows. The instance space of a simple blocks world is described by three attributes: *size* = {*small*, *medium*, *large*}, *color* = {*red*, *green*, *blue*}, and *shape* = {*square*, *circular*, *triangular*}. There is a sequence of three target concepts: (1) – *size* = *small* and *color* = *red*; (2) – *color* = *green* or *shape* = *circular*; and (3) – *size* = (*medium* or

large). 120 training instances are generated randomly and classified according to the current concept. The underlying concept is forced to change after every 40 training examples in the sequence: (1)–(2)–(3). The setup of the experiments with the STAGGER dataset was done exactly in the same way as in other similar works. The retraining step is 1, however there is a test set with size 100, generated randomly and classified according to the current concept. This differs from the other experiments where the retraining step and the test set are the same – a batch. The size of the FTM is set up to 25, which approximates the average size of the optimised windows.

5.1.2. German credit dataset. This subsection presents the results from the experiments conducted with the German credit dataset, from the UCI Machine Learning Repository¹. The dataset contains 1000 Instances of bank credit data which are described by 20 attributes. The examples are classified in two classes as either “*good*” or “*bad*”. To simulate hidden changes in the context the dataset was sorted by an attribute then this attribute was removed from the dataset for the experiments. Using an attribute to sort the data set and in this way simulate changing context is a commonly used approach to set up experiments to study concept drift. Two sorted datasets were created using the German credit dataset: The first one was sorted by a continuous attribute: “*age*”, which would produce a gradual drift of the “*class*” concept. The second one was sorted by the attribute “*checking_status*”, which has three discrete values. We aimed in this way to create abrupt changes of the “*class*” concept. Thus using this source dataset, two data streams are generated for the experiments. The datasets were divided into a sequence of batches, each of them containing 10 examples. The size of the FTM is set to 200, which approximates the average size of the optimised windows.

5.1.3. Spam dataset. Experiments have also been conducted with the Spam dataset from the UCI Machine Learning Repository. Spam is an unsolicited email message. The dataset consists of 4601 instances, 1813 (39.4%) of which are spam messages. The dataset is represented by 54 attributes that represent the occurrence of a pre-selected set of words in each of the documents plus three attributes representing the number of capital letters in the e-mail. To simulate the changing hidden context the examples in the dataset are sorted according to the “*capital_run_length_total*”, which is the total number of capital letters in the e-mail. This attribute and the related two attributes “*capital_run_length_average*” and “*capital_run_length_longest*” are removed from the dataset, because they can

¹ <http://www.ics.uci.edu/~mllearn/MLRepository.html>

provide explicit clues for the concept changes. The sorted dataset was divided into a sequence of batches with a length of 10 examples each. For this dataset the fixed window size was set to 400 – an approximation of the average window size for this dataset used by the TWO mechanism.

5.2. Machine learning algorithms. The experiments were conducted using three basic machine learning algorithms:

- k Nearest Neighbours (kNN) – also known as Instance Based Learning (IBL) [1]. $k = 3$ was the default setting for the experiments reported below except for experiments with STAGGER dataset, where $k = 1$ was chosen, because it produces a more accurate classification than $k = 3$;
- Induction of Decision Trees (ID3) [17] (using an attribute selection criteria based on the χ^2 statistics);
- Naïve Bayesian Classifier (NBC) [16].

6. Experimental results and discussion. The results from the conducted experiments are presented in Tables 1, 2 and 3 below. In all these tables, rows present the used learning algorithms: kNN , ID3 and NBC. The columns present the results from the experiments with different datasets. For a more compact presentation in the tables the dataset are referred with numbers as follows:

- 1** – STAGGER dataset;
- 2** – German dataset – sorted by the attribute “*checking_status*”;
- 3** – German dataset – sorted by the attribute “*age*”;
- 4** – Spam dataset – sorted by the attribute “*capital_run_length_total*”.

We used the paired t-tests with 95% confidence level to see whether there is significant difference in the accuracy of learned classifiers. The pairs are formed by comparing the algorithms’ accuracies on the same iteration. In the tables below, we are reporting the results from the experiments testing the above formulated hypothesis. The cells represent the results from the significance test comparing the results from the runs using the different forgetting mechanisms. The used basic learner is presented in the row and the used dataset stated in the column. The notation is as follows:

- ✓ – denotes that the second algorithm performs significantly better then the first one;

\times – denotes that the first algorithm performs significantly better than the second one;

\div – denotes that there is no difference in the performance of the compared algorithms.

For example: the sign \checkmark in the first column and first row in Table 1 denotes that the experiments conducted using kNN as a basic algorithm and using the STAGGER dataset (corresponds to the column 1), shows that FTW with GF is significantly better than the plain FTW, measured in accuracy of classification.

algorithm	dataset	1	2	3	4
kNN		\checkmark	\checkmark	\checkmark	\times
ID3		\checkmark	\checkmark	\checkmark	\div
NBC		\checkmark	\checkmark	\checkmark	\div

Table 1. Comparison of FTW versus FTW+GF forgetting mechanisms

Table 1 shows the results from the experiment testing the first hypothesis (i.e. FTW with GF provides a better classification accuracy in comparison to pure FTW). The results show that for three of the four datasets, for any of the algorithms the FTW with GF provides significantly better average classification accuracy in comparison to pure FTW, however the GF fails to improve the accuracy on the fourth data set. In comparison with other datasets, the main difference of this dataset is its sparseness. Therefore, we can assume that probably this is the reason for the lack of improvement. Perhaps, using a feature selection and an adapted similarity measure for kNN which have been used for textual dataset, will diminish this problem. However, further studies are needed to find the correct answer to these questions.

Despite the lack of success on the fourth dataset, the results from experiments with other dataset fully support the first hypothesis.

Table 2 presents the results from the experiment testing the second hypothesis (i.e. FTW with GF provides significantly better predictive accuracy in comparison to pure FTW). The results show that in all experiments the TWO mechanisms achieve significantly better classification accuracy than GF, except one where there is no difference in performance. The results from experiments provide strong evidence in support of the second hypothesis.

algorithm	dataset	1	2	3	4
kNN		✓	÷	✓	✓
ID3		✓	✓	✓	✓
NBC		✓	✓	✓	✓

Table 2. Comparison of FTW+GF versus TWO forgetting mechanisms

algorithm	dataset	1	2	3	4
kNN		÷	✓	✓	✓
ID3		✓	✓	÷	✓
NBC		✓	÷	✓	÷

Table 3. Comparison of TWO versus TWO+GF forgetting mechanisms

Table 3 shows the result from the experiment testing the third hypothesis: Applying GF with TWO improves the predictive accuracy in comparison to pure TWO i.e. both mechanisms are having synergetic effect. The results from the experiments show that there is no difference in the average performance in the compared mechanisms. Thus the experiments fail to confirm the third hypothesis.

7. Conclusion. The paper presents an experimental study of two forgetting mechanisms for dealing with the concept drift problem. Both are general in nature and can be added to any relevant algorithm. Experiments were conducted with three learning algorithms using four datasets. The results from the experiments show that: applying Gradual Forgetting inside a fixed-size Time Window leads to a significant improvement of the classification accuracy on drifting concept; when comparing both forgetting mechanisms the advantage clearly is on the side of the self-adapting algorithm – Time Window Optimisation; finally, applying both mechanisms together usually does not lead to an improvement in the classification accuracy.

Further controlled experiments using common patterns from the space of different type of drift (i.e. frequent – rare; abrupt – gradual; slow – fast; permanent – temporary; etc.) will provide clearer ideas of what kind of forgetting mechanism is most appropriate for which pattern.

REFERENCES

- [1] AHA D., D. KIBLER, M. ALBERT. Instance-Based Learning Algorithms. *Machine Learning* **6** (1991), 37–66.
- [2] CHU F., C. ZANIOLO. Fast and light boosting for adaptive mining of data streams. Proceedings of the 8th Pacific-Asia Conference on Knowledge Discovery and Data Mining. Lecture Notes in Computer Science, Vol. **3056**, Springer, 2004, 282–292.
- [3] DELANY S. J., P. CUNNINGHAM. An Analysis of Case-base Editing in a Spam Filtering System. In: Advances in Case-Based Reasoning, Proceedings of Seventh European Conference on Case-Based Reasoning, ECCBR-04 (Eds P. Funk, Gonzales, P. Calero.), Lecture Notes in Artificial Intelligence, Vol. **3155**, Springer, 2004, 128–141.
- [4] DUCATEL G., A. NÜRNBERGER. iArchive: An Assistant To Help Users Personalise Search Engines. In Chapter 2. User Profiles in Information Retrieval, Enhancing the Power of the Internet. Series: Studies in Fuzziness and Soft Computing (Eds. M. Nikravesh, B. Azvine, R. Yager, L. Zadeh) Vol. **139**, Springer, 2004, VIII, 406.
- [5] GAMA J., P. MEDAS, G. CASTILLO, P. RODRIGUES. Learning with Drift Detection. In: Proceedings of the 17th Brazilian Symposium on Artificial Intelligence. (Eds C. Ana, S. Bazzan, Labidi) Lecture Notes in Computer Science, Vol. **3171**, Springer, 2004, 286–295.
- [6] HARRIES M., C. SAMMUT. Extracting Hidden Context. *Machine Learning* **32** (1998), 101–126.
- [7] KLINKENBERG R. Learning Drifting Concepts: Example Selection vs. Example Weighting. In Intelligent Data Analysis. *Special Issue on Incremental Learning Systems Capable of Dealing with Concept Drift*, Vol. **8**, No. 3, (2004), 281–300.
- [8] KOYCHEV I., R. LOTHIAN. Tracking Drifting Concepts by Time Window Optimisation. In: Research and Development in Intelligent Systems XXII Proceedings of AI-2005, the 25th SGAI Int. Conference on Innovative Techniques and Applications of Artificial Intelligence. (Eds Max Bramer, Frans Coenen, Tony Allen) Springer, London, 46–59.

- [9] KOYCHEV, I., I. SCHWAB. Adaptation to Drifting User's Interests. Proceedings of ECML2000 Workshop: Machine Learning in New Information Age, Barcelona, Spain, 39–46.
- [10] KOYCHEV I. Gradual Forgetting for Adaptation to Concept Drift. Proceedings of ECAI 2000 Workshop on Current Issues in Spatio-Temporal Reasoning, Berlin, (2000), 101–107.
- [11] KOYCHEV I. Tracking Changing User Interests through Prior-Learning of Context. In: Adaptive Hypermedia and Adaptive Web Based Systems (Eds P. de Bra, P. Brusilovsky, R. Conejo) Lecture Notes in Computer Science, Vol. **2347**, Springer-Verlag (2002), 223–232.
- [12] KUKAR M. Drifting Concepts as Hidden Factors in Clinical Studies. In: Proceedings of 9th Conference on Artificial Intelligence in Medicine in Europe, AIME 2003, Protaras, Cyprus, October 18–22, 2003 (Eds D. Dojat, Elpida Keravnou, Pedro Barahona) Lecture Notes in Computer Science, Vol. **2780**, Springer-Verlag (2003), 355–364.
- [13] LAZARESCU M., S. VENKATESH, H. BUI. Using Multiple Windows to Track Concept Drift. *Intelligent Data Analysis Journal* **8**, No 1 (2004), 29–59.
- [14] MALOOF M., R. MICHALSKI. Selecting examples for partial memory learning. *Machine Learning* **41** (2000), 27–52.
- [15] MITCHELL T., R. CARUANA, D. FREITAG, J. MCDERMOTT, D. ZABOWSKI. Experience with a Learning Personal Assistant. *Communications of the ACM* **37**, No 7 (1994), 81–91.
- [16] MITCHELL T. Machine Learning. McGraw-Hill, (1997).
- [17] QUINLAN R. Induction of Decision Trees. *Machine Learning* **1** (1986), 81–106.
- [18] SCHLIMMER J., R. GRANGER. Incremental Learning from Noisy Data. *Machine Learning* **3** (1986), 317–357.
- [19] SPILIOPOULOU M., S. BARON. Temporal Evolution, Local Patterns. In: Local Pattern Detection, International Seminar, Dagstuhl Castle, Germany, April 12–16, 2004, Revised Selected Papers. (Eds. K. Morik, JF Boulicaut, A. Siebes) Lecture Notes in Computer Science, Vol. **3539**, Springer, 2005.

- [20] WIDMER G. Tracking Changes through Meta-Learning. *Machine Learning* **27** (1997), 256–286.
- [21] WIDMER G., M. KUBAT. Learning in the presence of concept drift and hidden contexts. *Machine Learning* **23** (1996), 69–101.

Department of Information Research
Institute of Mathematics and Informatics
Bulgarian Academy of Science
Acad. G. Bonchev Str., Bl. 8
1113 Sofia, Bulgaria
e-mail: ikoychev@math.bas.bg

Received May 11, 2006
Final Accepted February 23, 2007