

# Anticipative and Dynamic Adaptation to Concept Changes

Ghazal Jaber<sup>1,2,3</sup>, Antoine Cornuéjols<sup>1,2</sup>, and Philippe Tarroux<sup>3</sup>

<sup>1</sup> AgroParisTech, UMR 518 MIA, F-75005 Paris, France

<sup>2</sup> INRA, UMR 518 MIA, F-75005 Paris, France

<sup>3</sup> Université de Paris-Sud, LIMSI, Bâtiment 508, F-91405 Orsay Cedex, France

**Abstract.** Learning from data streams is emerging as an important application area. When the environment changes, as is increasingly the case when considering unending streams and long-life learning, it is necessary to rely on on-line learning with the capability to adapt to changing conditions a.k.a. concept drifts. Previous works have focused on means to detect changes and to adapt to them. Ensemble methods relying on committees of base learners have been among the most successful approaches.

In this paper, we go one step further by introducing a second-order learning mechanism that is able to detect relevant states of the environment, to recognize recurring contexts and to *anticipate* likely concepts changes. Results of an empirical comparison with adaptive methods show that, for a very slight price in memory and computation load, the proposed algorithm always improves on, or at worst equals, the prediction performance of a mere adaptive approach.

## 1 Introduction

Recent years have witnessed the emergence of a whole new set of applications involving data streams made of pairs  $(\mathbf{x}_t, y_t)$ , where the “answer” or true label  $y_t$  is revealed (sometimes long) after the input  $\mathbf{x}_t$ . For instance, a set of customers can be submitted to ads or offers arriving on sequence, to which they answer with buying actions or not. The task for a forecaster is to predict the answer to each new incoming incentive  $\mathbf{x}_t$ , possibly in order to regulate the stocks, even before the outcome  $y_t$  is known. Because of varying economic conditions, or because of changes in the season or in the weather, the customers may modify their buying behavior. Sometimes it can even occur abruptly like when a big amendment in the economic policy is announced.

Learning from streams [9] is usually treated using online machine learning techniques which differ from classical batch learning methods in three major points. First, streaming data are not stored or reprocessed, due to memory constraints. Secondly, the prediction model should give answers in an any time fashion, while updating itself with each received information from the stream. Finally, online learning does not presuppose that the training data be independent and identically distributed. It is ready to adapt to changing conditions. This is why, even though most works in recent years have dealt with the computational issues raised by the demand for a small constant learning time and near constant memory resources, a stream of research has also focused on evolving

concepts in case of non-stationary environments, specially on *how to detect concept changes* and *how to best adapt* to them.

In this context, passive adaptation to concept changes may not be the best learning strategy. Indeed, a learner may profit from the information possibly conveyed by the very sequence of data. For instance, one can gain precious time and avoid costly incorrect predictions by being able to recognize a recurring situation or to anticipate the likely evolution to come along. This kind of second order learning is the object of the approach presented in this paper. The ADACC (*Anticipative Dynamic Adaptation to Concept Change*) method that we present deals both with the challenge of optimizing the stability-plasticity dilemma (keeping as much data as possible while at the same time recognizing when data points become obsolete and potentially misleading) and with the anticipation and recognition of incoming concepts. This is accomplished through an ensemble method that controls a pool of incremental learners. The use of a pool of learners enables the system to naturally adapt to the dynamics of the concept changes with few parameters to set. At the same time, a second order learning mechanism that observes and control the changes in the pool provides means for the anticipation of, and the quick adaptation to, the underlying modification of the context.

The rest of the paper is organized as follows. In Section 2, we discuss existing works on online learning in the context of concept changes. Section 3 presents our contribution, which is followed by the report of empirical results in Section 4. The last section presents conclusions and possible avenues for future work.

## 2 Relevant Works

Aside from the question of the severe constraints posed on the computational load of online learning (see [8] for a pioneering study and [11] for more references), the main issue in online learning is connected with the problem of learning in the presence of a changing environment. These changes can come in two guises according to which of the input distribution  $\mathbf{p}(\mathbf{x})$  or the conditional distribution  $\mathbf{p}(y|\mathbf{x})$ , is affected. If only  $\mathbf{p}(\mathbf{x})$  changes, it is said that a *virtual drift* or *covariate shift* occurs. If the changes in the environment concern the concept itself  $\mathbf{p}(y|\mathbf{x})$ , a *concept drift* occurs, which can be gradual or abrupt. In an abrupt drift, the target concept gives suddenly way to a new one (see Figure 1). In a gradual drift, however, the new target concept takes over the old one over a period  $\tau_{drift}$ . The most general type of changing environments involves changes both in  $\mathbf{p}(\mathbf{x})$  and in  $\mathbf{p}(y|\mathbf{x})$ . When the environment varies, it is important to be able to detect the changes and to modify the prediction model accordingly in order to maintain as best as possible the prediction performance.

Usually, the *detection of change* is done by monitoring either the variations in the distribution of the incoming data, or the variations in the prediction's performance of the system. In both cases, it is necessary to set an alarm threshold and to decide upon an adequate duration of observation in order to trigger justified alarms only. This requires some knowledge of the dynamic of the environment.

Likewise, *adaptation to changes* meets a dilemma regarding the adequate length of the subsequence of the data stream that must be trusted for providing relevant information about the current target concept. The earliest techniques used sliding windows over

the data stream [19]. The length of the window can be set beforehand, but more sophisticated approaches rely on some dynamic control of this length depending of the stability of the prediction performance. A second strategy uses a weighting scheme over past instances which is deemed to reflect the relevance of the data to the current context. Both these techniques, windowing and weighting, imply some a priori choices about threshold, decay factors and so on [12]. A third strategy tries to avoid these choices by relying instead on an ensemble-based learning technique, akin to bagging or boosting. The idea is to let a committee of base learners incrementally learn over the stream of data, and to replace the worst ones at certain times [16]. Thus, the hope is that the committee automatically eliminates the no longer relevant base learners while keeping and improving the most promising ones. In addition, for each incoming instance, the prediction to be made can result from a combined vote from the base learners, hopefully leading to more accurate answers, as in classical ensemble learning methods.

The committee based strategy still entails some choices regarding the base learners, the size of the ensemble or committee, the deletion mechanism, the introduction and initialization of new learners in the ensemble, and the decision voting process. Recent years have seen several proposals among these lines showing promising results on a variety of online learning tasks [1, 18, 4]. A limit of these algorithms, however, is that they passively wait for the changes to occur and then try to follow them as best as possible rather than proactively predict what is likely to happen.

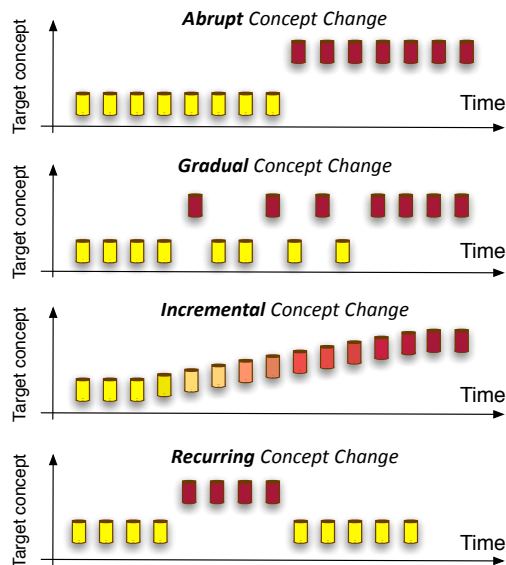
Very few works have confronted the anticipation of concept changes. Among them, the PreDet [5] algorithm uses decision trees as classifiers and anticipates future trees by predicting for each decision node the evaluation measure of each attribute, this value being used to determine which attribute will split the node. In the case of a leaf node, it predicts its class label distribution. Future changes are predicted using a linear regression model trained on a fixed size history. Another prediction system, RePro [21], stores the observed concepts in a Markov chain. In order to compute the transition matrices, RePro assumes that concepts repeat over time. Thus, when a change is detected, the Markov chain is used to predict the most likely concept to come.

The method we present below aims at benefiting from the knowledge of the past to predict future concept changes. However, contrary to other systems, it can be used with any base learner that produced parameterized hypotheses (e.g. neural networks) and is not limited to decision trees. Furthermore, unlike Markov chains, it does not necessitate that target concepts repeat over time in order to make anticipation.

### 3 Concept Changes: Adaptive and Anticipative

In online learning, the learning protocol is as follows. At each time step  $t$ , an input  $\mathbf{x}_t$  is received, the learning algorithm is asked to make a prediction  $\hat{y}_t$  over its label, and then only is revealed the true label  $y_t$ . The algorithm can then adapt its current view of the world, noted here  $H_t$ , and be ready for the next input  $\mathbf{x}_{t+1}$  to come.

The overall goal is to minimize the loss over time :  $\frac{1}{T} \sum_{t=1}^T \ell(H_t(\mathbf{x}_t, y_t))$  with  $T \rightarrow \infty$ . Usually, the loss function is simply the 0-1 loss function which counts the number of mistakes.



**Fig. 1.** Four different types of concept changes.


In the most general case, the target concept, which decides the label for the next input  $x_t$ , can change arbitrarily over time, and even be manipulated by an adversary. Then, the only existing theory is the online learning theory which characterizes the *regret* that can be achieved by a learning algorithm with respect to the best available “expert” in a given pool of experts [7]. If, however, more benign assumptions can be made about the dynamics of the environment, such as it evolves gradually or incrementally or it evolves with sudden changes but interspersed with stationary states (see Figure 1), then it becomes legitimate to try to assess the future performance of a classifier on the base of (a partial view of) its past history.

In this paper, we work on ways to anticipate future concepts when a transformation rule can be identified, and on the recognition of recurring concepts, whether the change between consecutive concepts happens gradually or suddenly.

In the following, we first introduce the adaptive architecture of the learning system we take as our basis (Section 3.1). Then, we present the core of this paper which is the anticipating mechanism that builds upon it (Section 3.2).

### 3.1 Adapting to Concept Changes

Being able to recognize changes in the environment and to anticipate them requires some kind of second-order or meta learning. The learning system needs to be able to analyze and reflect upon its past experiences and to decide what is the best course of action or decision given the past. For this meta-learning to take place, the adaptive strategies based on ensemble of base classifiers are well suited because they naturally manage a set of potential models of the changing environment.



**Fig. 2.** *Left*, a concept represented by a circle. The examples are classified into one of two classes: inside, outside the circle. *Right*, the circle moving with time, creating concept drifts.

At any time, the current ensemble of base classifiers represents a kind of memory of the past, and offers the opportunity for second order learning. Furthermore, these ensemble methods generally adapt gracefully to a variety of changing conditions in the environment and do not require fine tuning of their few parameters.

The main idea is to maintain a pool of base learners  $\{h_t^i\}_{1 \leq i \leq N}$ , each of them adapting to the new input data, and to administer this pool or ensemble thanks to a deleting strategy and an insertion one. The main principles of these ensemble methods are the following:

- Each base learner in the pool continuously adapts with new incoming data until it is removed from the pool.
- Every  $\tau$  time steps, the base learners are *evaluated* on a window of size  $\tau_{eval}$ .
- Based on the results of this evaluation, the deletion procedure *chooses* a base learner to be removed.
- A new based learner is *created* and inserted in the pool. It is protected from possible deletion for a duration  $\tau_{mat}$ .
- For each new incoming instance  $\mathbf{x}_t$ , the prediction  $H(\mathbf{x}_t)$  results from a *combination* of the prediction of the individual based learners  $h_t(\mathbf{x}_t)$ .

Individual variations around this general framework lead to specific algorithms (for instance [17] [13]). In our studies, and after extensive testings, we converged on the following settings. The *base learners* can be any supervised predictors (in the reported experiments, they are (multi-layer) perceptrons or decision trees). The *evaluation* procedure simply counts the number of erroneous predictions on the last  $\tau_{eval}$  time steps. The *deletion strategy* randomly select one base learner from the worst half of the pool evaluated as above. The *global prediction* merely uses the prediction from the current best base learner (randomly chosen in case of ties).

This simple method offers a good trade-off on the plasticity-stability dilemma and leads to fast adaptation when the underlying concept changes.

### 3.2 Anticipating Concept Changes

In an evolving environment, two aspects should be considered when anticipating future concept changes: recurrence and predictability. *Recurrence* means that the same

concepts (or close approximations) might reappear with time, either in a cyclic manner (e.g. seasonal variations) or in an irregular manner (e.g. inflation rate, market mood). *Predictability* means that the concept evolves in a predictable manner, but without necessarily repeating over time. Here, “predictable” refers to an underlying prediction system, that is a learning system that, taking as input information about the past history of the concept evolution, is able to predict its (near) future. This, of course, implies that the past history be captured and represented. We explain below how a system can determine relevant states (called *snapshots*) of the evolution.

As an illustration, consider a concept defined in a 2-dimensional input space corresponding to a circle with a given radius. At any given time, the instances are labeled into two classes depending on whether they lie inside or outside the circle (see Figure 2). A concept drift would occur if the circle’s center moves, yielding a change in the labels of the examples “swept” by the moving concept. If the concept changes in some regular manner, say in a straight line with a constant speed, an appropriate learning system might be able to learn from a sequence of concept changes and predicts the likely future concepts. For instance, an Ellman’s neural network [20] could be trained on such a sequence in order to output predictions about the future changes. This is an instance where Markov chains would be ineffective since there would be no repetitions of past concepts.

In the following, we first describe a mechanism for the *recognition of the relevant states of the world* before presenting the *second-order learning mechanism* that works on these states in order to predict future models. Algorithm 1 offers a formal view of the steps and methods described below.

### **Recognition of significant models of the world**

In this step, we aim at recognizing the significant concepts encountered during the data streaming. Our goal is to take a snapshot (copy) of any recognized stable concept via the learned models or hypotheses in the adaptive ensemble. Note that we suppose therefore that the concept drift operates via a sequence of stable concepts in a gradual, abrupt or recurring manner (see Figure 1). A main challenge here is to decide when to take a snapshot. Given our reliance on an ensemble method, this is equivalent to decide when the ensemble reflects the existence of an underlying relevant concept.

One approach has been proposed in the ADWIN system [3]. It supposes a priori that  $m$  training instances are necessary for the system to stabilize after an abrupt concept change. Consequently, in this approach, a snapshot is taken  $m$  time steps after a change is detected. This method however presents limitations. First, it has difficulties detecting gradual concept changes resulting in snapshots that have often low relevance. Second, the choice of the value for the parameter  $m$  is difficult if one does not know in advance the properties of the incoming concept changes.

In this work, rather than making prior assumptions about the dynamics of the world, we aimed at finding an implicit way to identify periods of stability using the adaptive ensemble as a detector.

The method is as follows. When the environment has been in a stable state for a sufficient time, the best hypotheses in the pool of base learners should converge toward the same, and near optimal, predictive performance. Therefore, their diversity, measured as below, should be low, while their error rate should decrease toward the

best achievable performance or close to it. This suggests to take into account both the diversity of the best hypotheses in the pool and their error rate as an index of the stability of the environment. In our study, we use the kappa statistics  $\mathcal{K}$  [6] in order to compute the *diversity*. This statistics measure evaluates the degree of agreement between the classification of a set of items by two classifiers<sup>4</sup>. In case of complete agreement,  $\mathcal{K} = 1$ . If there is no agreement other than what would be expected by chance,  $\mathcal{K} = 0$ . The *stability index* at time  $t$  is computed over the last  $\tau_s$  received examples:  $I_{stability} = agreement - performance$  where *agreement* and *performance* are computed over the best half of the current hypotheses in the pool, and are defined as:

$$agreement = \frac{\sum_{j=0}^{\tau_s-1} \sum_{\substack{i,j=1 \\ i \neq j}}^{N/2} \mathcal{K}_{h_i, h_j}(\mathbf{x}_{t-j})}{\tau_s * \frac{N}{2} * (\frac{N}{2} - 1)} \quad (1)$$

and:

$$performance = \frac{\sum_{j=0}^{\tau_s-1} \sum_{i=1}^{N/2} err(h_i, \mathbf{x}_{t-j}, y_{t-j})}{\tau_s * \frac{N}{2}} \quad (2)$$

where  $N$  is the size of the pool.

Each point in the stability index curve, over some predefined threshold  $\theta_I$ , is suggestive of a stable environment and is therefore a privileged moment to take a *snapshot* (copy) of the current best hypothesis in the pool of base learners, the one that seems to best represent the current state of the world.

One must however be careful not to store consecutive hypotheses that correspond to the same underlying state of the environment. Here again, the agreement statistics, in our case the kappa statistics, can be used to measure the agreement between a candidate snapshot  $h_t^*$  and the preceding one  $h_{t_k}^*$ . For this purpose, a set of  $n$  unlabeled data  $U = \{\mathbf{x}_i\}_{i=1}^n$  is generated at random using a uniform distribution over the input space  $\mathcal{X}$ , and the predictions of  $h_t^*$  are then compared with the predictions of the last stored snapshot  $h_{t_k}^*$ . If the estimated agreement is less than some predefined threshold  $\theta_d$ , the current candidate snapshot  $h_t^*$  is considered different enough from  $h_{t_k}^*$  and is therefore added to the list  $\mathcal{M}_{LT}$  of snapshots representing past stationary states of the environment.

### The long term memory and second order learning

The list of past snapshots  $\mathcal{M}_{LT} = \{C_1, C_2, \dots, C_k\}$ , ordered according to the snapshots' time appearance, is the basis of the second order learning mechanism. It serves two purposes. *First*, it provides a sequence of successive models of the environment that can be used by a learning algorithm in order to predict the most likely future state in the series. *Second*, it stores a memory of past successful models of the world, models that should be repeatedly tested against current data in case a recurring concept can be recognized.

In our experiments, we used Elman's recurrent neural networks as predictors because of their generality and the good performances reported for learning tasks similar to our's [20]. Accordingly, a network is trained on the pairs of consecutive concepts in

<sup>4</sup> Other agreement statistics should do as well.

---

**Algorithm 1:** Selection of snapshots by ADACC.

---

```

begin
   $E_0 \leftarrow \emptyset;$                                 /* Ensemble of experts */
   $\mathcal{M}_{LT} \leftarrow \emptyset;$                 /* List of snapshots */
   $k \leftarrow 1;$ 
  for  $t = 1$  to  $\infty$  do
    /* ----- */
    /* Adaptation */
    /* ----- */
     $(\mathbf{x}_t, y_t)$  is the current training instance;
     $[E_t, \tilde{y}_t] \leftarrow \text{AdaptationEnsemble}(E_{t-1}, \mathbf{x}_t, y_t);$ 
    /* ----- */
    /* Anticipation */
    /* ----- */

     $H = \{h_i\}_{i=1}^N / 2$  is the best half of experts in  $E_t$ ;

     $agr = \frac{1}{\tau_s * \frac{N}{2} * (\frac{N}{2} - 1)} \sum_{j=0}^{\tau_s-1} \sum_{\substack{i,j=1 \\ i \neq j}}^{N/2} \mathcal{K}_{h_i, h_j}(\mathbf{x}_{t-j});$ 

     $perf = \frac{1}{\tau_s * \frac{N}{2}} \sum_{j=0}^{\tau_s-1} \sum_{i=1}^{N/2} err(h_i, \mathbf{x}_{t-j}, y_{t-j});$ 

     $I_{stability} = agr - perf;$ 
    /* Detect Stable Concept */
    if  $I_{stability} \geq \theta_t$  then
       $h_t^* = \text{snapshot}(E_t);$ 
      /* Detect New Concept */
      if  $isEmpty(\mathcal{M}_{LT})$  then
         $C_k = h_t^*;$ 
         $\mathcal{M}_{LT} = add(\mathcal{M}_{LT}, C_k);$ 
      else if  $\mathcal{K}_{C_k, h_t^*} \leq \theta_d$  then
         $k = k + 1;$ 
         $C_k = h_t^*;$ 
         $\mathcal{M}_{LT} = \text{replace}(\mathcal{M}_{LT}, \tilde{C}_k, C_k);$ 
         $\tilde{C}_{k+1} = \text{predictNextConcept}(\mathcal{M}_{LT});$ 
         $\mathcal{M}_{LT} = add(\mathcal{M}_{LT}, \tilde{C}_{k+1});$ 
    end
  end

```

---

$\mathcal{M}_{LT}$ :  $\{(C_1, C_2), (C_2, C_3), \dots, (C_{k-1}, C_k)\}$  in order to predict the next likely snapshot  $\tilde{C}_{k+1}$ . To simplify the discussion, a snapshot  $C$  is represented as a vector of parameters of dimension  $n$ .

$$C = [c_1, c_2, \dots, c_n] \quad (3)$$

For instance, if the current concept is learnt by a base learner that is a neural network, the snapshot could be represented as the vector of the network's weight values. A pair consisting of two consecutive snapshots  $\delta_i = (C_i, C_{i+1})$  is called a *change sample*.



After recognizing  $k$  stable concepts, the  $k - 1$  change samples  $(\delta_1, \dots, \delta_{k-1})$  form the set of training examples input to the Elman’s network in order to predict the next snapshot  $\tilde{C}_{k+1}$ . The predicted snapshot is then temporally added to the list of snapshots  $\mathcal{M}_{LT}$ . It is replaced by the next snapshot  $C_{k+1}$  when this one is acquired.

Therefore, the list  $\mathcal{M}_{LT}$  contains snapshots that represent past stationary states of the environment (useful for the recognition of recurring concepts), plus a predicted future state (useful for anticipation).

Each snapshot in the list is then treated on an equal footing with the base learners in the pool of the adaptive ensemble method. The snapshots are therefore evaluated according to the evaluation strategy used by the adaptive ensemble to evaluate the base learners in the pool. A snapshot is used for prediction if its evaluation record is the best among all candidate hypotheses from both the pool of base learners and  $\mathcal{M}_{LT}$ .

While the pool of candidate base learners is managed according to the policy outlined in Section 3.1 and is therefore of a finite constant size, the list  $\mathcal{M}_{LT}$  of snapshots may a priori increase forever if new hypotheses are continually retained as worthy of storage. Fortunately, it is possible to keep this size under control by recognizing that the two roles of  $\mathcal{M}_{LT}$ : anticipation and memory for recurring concepts, ask for two different memory management systems. Indeed, since Elman’s networks are incremental learners, they do not need to keep the past history of snapshots at all. Regarding the memory for recurring concepts, it can be kept constant using various heuristics. One is to delete the oldest or the least recurring snapshots from the memory. Another one, more sophisticated, would be to store prototype snapshots instead of the original ones, using a hierarchical clustering technique. In our experiments, we did not rely on such memory management schemes since there was no need for them.


## 4 Empirical Results

The aim of the experiments was threefold. First, to test the performance of *the snapshot mechanism* and specially its ability to detect both abrupt and gradual concept changes and store the relevant target concepts with no, or limited, redundancy. Second, to examine the gain, if any, brought by *the anticipation scheme* compared to the mere adaptation mechanism. Obviously, this depends on the ability to anticipate the next state of the environment, and therefore on the underlying structure (if any) of the sequence of changes [2]. Thirdly, to test the mechanism for *the recognition of recurring concepts* and the gain it can bring.

In the worst case, where it is not possible to anticipate the next concept and when no recurring concept arises, the prediction performance of the system should fall back to the performance of its adaptation mechanism. Indeed in these cases, no snapshot in  $\mathcal{M}_{LT}$  does outperform the best base learners and the resulting behavior is the one of the adaptive system alone.

### 4.1 Experiments and Datasets

We conducted experiments on artificial and real data sets. The artificial sets were used to simulate *recurrent* and *predictable concept changes* while controlling the timing of the



**Fig. 3.** A typical evolution in the weight values of the hyperplane used in the artificial datasets.

change, its speed (abrupt, or more or less gradual) and its severity (amount of change) [14]. The real data set comes from video sequences taken with a mobile robot wandering in and out of rooms in a laboratory, creating *recurring contexts*.

In the *artificial data sets*, the input space  $\mathcal{X}$  is  $d$ -dimensional and the target concept is associated with a linear decision boundary (a hyperplane) described by the relation  $y(\mathbf{x}) = \text{sign}(\sum_{i=1}^d w_i x_i + w_0)$ . The experiments were carried out on streams with 7,150 time steps and hence data points.

In each stream, 12 concept changes were simulated by changing the weights  $\{w_i\}_{i=0}^d$  of the target hyperplane. The first 7 concepts evolved through the successive addition or subtraction of constant values (differing according to the experiments) to the weights. The idea was to look at the capacity of the anticipative mechanism to identify this regularity and therefore to predict likely future concepts. The last 6 concepts were recurring concepts, that is concepts already encountered in the past data stream (see Figure 3).

Three artificial data streams were generated, each involved a different level of severity in its concept changes: *low* or *medium* or *high* respectively involving changes in 1, 5 and 9 parameters out of the 11 that define the target concept, with respectively approximately 3%, 60% and 84% of the input space changing class between successive concepts. In each stream, the changes happened either suddenly or gradually, in a linear manner, between successive target concepts.

The transition between consecutive concepts took from 0 to 200 time steps and changes would start happening every 400 to 700 time steps. We did not observe any effect of the dimension up to more than one hundred and therefore only report results for the 10-dimensional case. The base learners in the adaptation ensemble were perceptrons with 10 input units and one output unit, involving 11 weights (10 + 1 for the bias). The Elman's networks took as input the 11 weights of a snapshot and gave as output the 11 weights of the next predicted snapshot.

The *real data set* was issued from the COLD database of the Saarbrücken laboratory [15], a benchmark for vision-based localization systems. It contains sequences of images recorded by a mobile robot under different variations of illumination and weather: sunny, cloudy and night. We worked on the dataset captured in sunny conditions. The

images were labeled into one of four classes: *corridor*, *one-person office*, *printer area* and *classroom*, and the total length of the data sequence was 753. The robot visited the rooms in the following order: *corridor*, *bathroom*, *corridor*, *one-person office*, *corridor*, *printer* and *corridor*. It stayed in the same room between 45 and 284 time steps. Images were first pre-processed into a 128-dimensional space using the Self-Organizing Map described in [10]. In the experiments, we used decision trees (as implemented in Matlab) as base learners in the adaptation ensemble.

One important goal of the experiments was to compare the performances achieved with the combined anticipative and adaptive mechanism, with the ones of a purely adaptive mechanism.

The *anticipative meta-learning* system itself involves three parameters that all pertain to the detection of relevant snapshots. They are the *stability threshold*  $\theta_i$ , the *decision threshold*  $\theta_d$  and the *duration for the evaluation* of candidate snapshots  $\tau_s$ . They were set respectively to  $\theta_i = 0.9$  and  $\theta_d = 0.8$  while  $\tau_s = 100$  was chosen for artificial data streams and  $\tau_s = 25$  for robotics in order to cope with a faster dynamics. In the base version with no sophisticated management of the snapshot list  $\mathcal{M}_{LT}$ , there are no additional parameters.

The remaining parameters concern the adaptive mechanism, and we tried to optimize these in order to not unfairly attribute gains to the anticipative process. The parameters for the ensemble method for *adaptive online learning* include the size of the pool  $N$ , the maturity age  $\tau_{mat}$  and the evaluation size  $\tau_{eval}$ . After extensive experiments, they were set as follows.

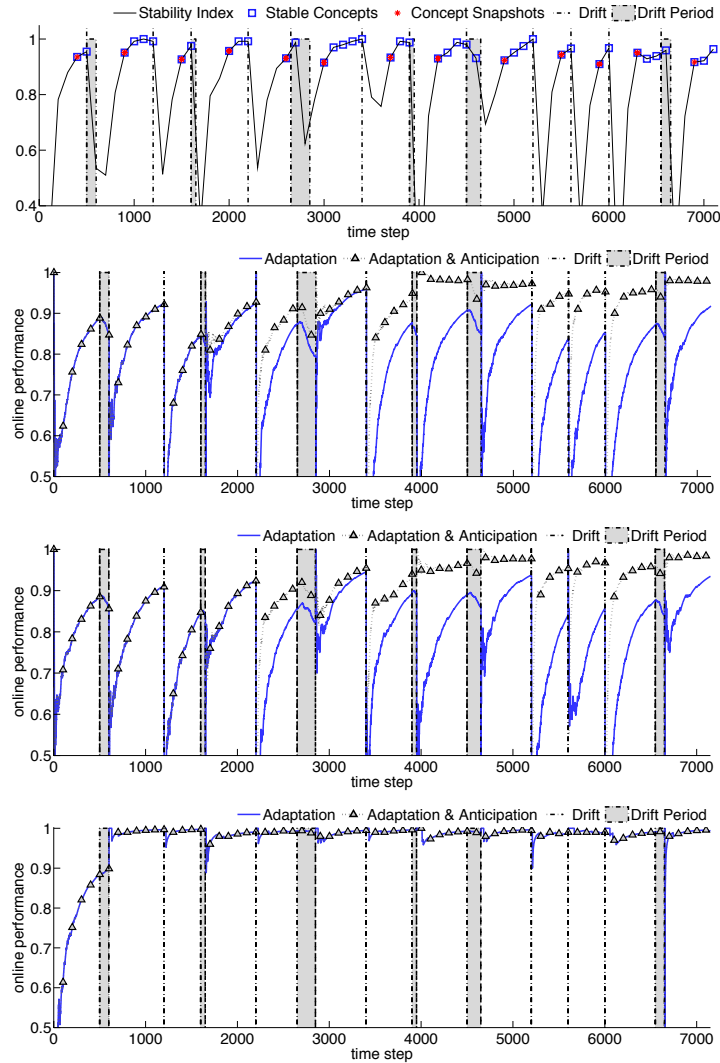
The pool comprised  $N = 20$  base learners for the artificial data sets ( $N = 15$  for the robotics data). In order to be compared, base learners were evaluated on the most recent  $\tau_{eval} = 20$  data points (time steps) ( $\tau_{eval} = 15$  for the robot). The duration for maturity  $\tau_{mat}$  was equally set to 20 time steps ( $\tau_{mat} = 10$  for the robot).

## 4.2 Evaluation Measures and Methodology

In the experiments, we evaluated *the snapshots* stored by the system with respect to the known target concepts. Ideally, there would be one snapshot exactly for each encountered state during the data stream. For instance, in the top of Figure 4, candidate snapshots are indicated with small squares and the retained ones appear as red (or black) squares.

We also evaluated the *gain in prediction errors* resulting from the use of the anticipation mechanism over the use of the adaptation scheme alone. Likewise, we measured the gain (if any) due to the recognition of a recurring concept. The gain is simply the number of errors of prediction that were avoided with respect to the use of the adaptive strategy only (see Table 1 below).

Finally, the graphs (see Figure 4, the bottom three graphs) report at each time step the current *online predictive performance*, i.e. the mean number of instances correctly classified so far. In order to better visualize the gain after each concept change, the performances of the adaptation mechanism and of the adaptation + anticipation mechanisms were reset to 0.5 (the chance prediction rate). One can then observe, for instance, that the gains due to the anticipation mechanism start only to show after the fourth concept change, which is unsurprising.



**Fig. 4.** Curves for 10-dimensional artificial data streams. The *top* plot shows the evolution of the stability index in case of the medium severity concept changes. Small squares indicate the time steps where candidate snapshots are considered, and (red (or black) squares) when they are retained. Each concept change gives rise to exactly one snapshot. The three other plots (*high severity* concept change at the top, *medium severity* in the middle and *low severity* at the bottom) show the online predictive performance of the adaptive learning strategy (continuous (blue) line) and with the second order learning taking place. They are averaged over 10 repeated experiments. The beginning/end of concept changes are indicated as vertical dotted lines. In case of gradual concept changes, the transition period between consecutive concepts is colored in gray. The online predictive performance is reset every time a transition is complete.

### 4.3 Results

Stream		Adaptation			Anticipation	Total gain		Due to predict.	Due to recurrence
name	size	base learner	mean err.	std-dev	predictor	mean	std-dev	mean	mean
10-D Low	7,150	perceptron	<b>107.2</b>	7.7	Elman net	<b>1.9</b>	1.7	0.0	1.9
10-D Med.	7,150	perceptron	<b>803.8</b>	43.1	Elman net	<b>296.2</b>	23.9	63.3	232.9
10-D High	7,150	perceptron	<b>912.1</b>	46.5	Elman net	<b>345.9</b>	38.9	101.4	244.5
Robot	753	decis. tree	<b>43.0</b>	2.6	-	<b>9.0</b>	1.9	-	9.0

**Table 1.** Summary of the experiments and the measured gains in prediction errors wrt. an adaptive only strategy.

Table 1 sums up the experimental results, averaged over 10 experiments. The table shows the mean predictive error of the adaptive learning strategy, and the gain of using the anticipation mechanism, in both predictability and recurrence. The gain is measured as the difference between the number of prediction errors made by the adaptive ensemble and the number of prediction errors made by the anticipation mechanism. In the artificial data streams, we highlight the gain brought by the *predictability* of the first 7 concept changes, and the gain brought by the last 6 *recurring* concepts.

Figure 4 illustrates the mechanism for the selection of snapshots on one data stream and it shows the evolutions of the prediction performance over 10 repeated experiments according to the severity of the concept changes.

#### Detection of concept changes and selection of snapshots

As can be seen in Figure 4, the value of the stability index closely mirrors the concept changes. As soon as the appearance of a new concept is detected by the system and the corresponding candidate snapshot sufficiently differs from the previously stored ones, it is stored away in  $\mathcal{M}_{LT}$ . That policy enables the fast detection of novel target concept. In our experiments, 100 % of all new concepts (2,600 altogether) that were introduced in the data streams triggered the storage of a new snapshot. There was no redundancy (no more than one snapshot per concept) in the artificial data streams. Some redundancy appeared for the robotic data because of the variation within each concept (e.g. *office*).

#### Second order learning

For concept changes of *low severity* (Figure 4, bottom), the adaptive strategy is able to follow the variation of the environment as soon as enough candidate hypotheses are good enough, which happens at the end of the first concept (circa 400 time steps). Therefore, the anticipation strategy does not bring an advantage there. The situation is significantly altered, however, when the concept changes are of *medium* or *high severity*.

In our experiments, even though the concept changes occur at varying dates and with varying speed, the anticipation mechanism is able to predict relevant foreseeable target concepts that, in turn, are quickly recognized as the best for labeling the incoming

examples. This brings significant gains in the online performance starting already after the 4th change of concept, and the gain increases thereafter with each new concept change.

Table 1 shows that the gain in the number of labeling errors attains more than  $296/803 = 36\%$  for concept changes of medium severity, and approximately 38% in the case of high severity. These gains are impressive in face of a difficult learning task. It is unlikely that they could be obtained without a second order learning mechanism working over the adaptive one.

Table 1 distinguishes furthermore between the gain due to the predictability and the gain due to the fast recognition of a recurring concept. Predictability brings significant gain in the medium and high severity settings for the artificial data sets. In the case of the robotics data, the gain is totally due to the fast recognition of recurring concepts which outperforms the anticipation mechanism.

As expected, there is never a negative gain. As noted earlier, because the ensemble methods is based on a continual competition between base learners from the adaptive mechanism and base learners from the anticipative one, second order learning can never be detrimental to the overall prediction performance as compared to the adaptive only policy.

## 5 Conclusions and Future Work

The ability to make predictions when data arrives continuously in stream, possibly from a non stationary environment, is becoming increasingly important. Significant research works in recent years have brought new techniques to cope with these learning conditions. In this paper, we presented a general framework to endow adaptive online learning systems based on an ensemble approach with second order learning capacity.

Our method provides means (i) to identify significant stationary states of the world, (ii) to make anticipation about likely future states, and (iii) to recognize recurring concepts if they ever arise. Few parameters are involved in the second-order learning scheme and they appear to not need to be finely tuned.

The empirical evaluation explored various conditions for evolving data streams. It showed that as soon as the concept changes are significant (medium or high severity), second order learning yields substantial gains in prediction performance over a mere adaptation policy. Furthermore, second order learning can only improve and never deteriorate the prediction performance, at a small cost in memory and computation.

Our future plans include to carry out experiments with very long streams of data ( $\sim 10^5$  time steps) in order to test possible strategies for the management of the memory of snapshots.

## Acknowledgments

Two Grants from the French National Research Agency (ANR): “Holyrisk” (ANR “Blanc-2009” NT09.439612) and “Coclico” (ANR ”MN-2012”) supported in part this work.

## References

1. M. Baena-García, J. del Campo-Ávila, R. Fidalgo, A. Bifet, R. Gavaldà, and R. Morales-Bueno. Early drift detection method. *Fourth International Workshop on Knowledge Discovery from Data Streams*, 2006.
2. P. Bartlett, S. Ben-David, and S. Kulkarni. Learning changing concepts by exploiting the structure of change. *Machine Learning*, 41(2):153–174, 2000.
3. A. Bifet and R. Gavaldà. Learning from time-changing data with adaptive windowing. In *SIAM International Conference on Data Mining*, pages 443–448, 2007.
4. A. Bifet, G. Holmes, B. Pfahringer, R. Kirkby, and R. Gavaldà. New ensemble methods for evolving data streams. 2009.
5. M. Bottcher, M. Spott, and R. Kruse. Predicting future decision trees from evolving data. In *Data Mining ICDM'08. Eighth IEEE International Conference on*, pages 33–42, 2008.
6. J. Carletta. Assessing agreement on classification tasks: the kappa statistic. *Computational linguistics*, 22(2):249–254, 1996.
7. N. Cesa-Bianchi and G. Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
8. P. Domingos and G. Hulten. Mining high-speed data streams. In *Proceedings of the sixth ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 71–80. ACM, 2000.
9. J. Gama. *Knowledge discovery from data streams*. Citeseer, 2010.
10. H. Guillaume, M. Dubois, P. Tarroux, and E. Frenoux. Temporal bag-of-words: A generative model for visual place recognition using temporal integration. In *Proceedings of the International Conference on Computer Vision Theory and Applications*, 2011.
11. R. Kirkby. *Improving hoeffding trees*. PhD thesis, The University of Waikato, 2007.
12. R. Klinkenberg. Learning drifting concepts: Example selection vs. example weighting. *Intelligent Data Analysis*, 8(3):281–300, 2004.
13. J. Z. Kolter and M. A. Maloof. Dynamic weighted majority: An ensemble method for drifting concepts. *The Journal of Machine Learning Research*, 8:2755–2790, 2007.
14. L. Minku, A. White, and X. Yao. The impact of diversity on online ensemble learning in the presence of concept drift. *Knowledge and Data Engineering, IEEE Transactions on*, 22(5):730–742, 2010.
15. A. Pronobis and B. Caputo. Cold: Cosy localization database. In *The International Journal of Robotics Research*, 28(5), 2009.
16. K. Stanley. Learning concept drift with a committee of decision trees. *Informe técnico: UT-AI-TR-03-302, Department of Computer Sciences, University of Texas at Austin, USA*, 2003.
17. K. O. Stanley. Learning concept drift with a committee of decision trees. *UT-AI-TR-03-302, Department of Computer Sciences, University of Texas at Austin, USA*, 2003.
18. A. Tsymbal, M. Pechenizkiy, P. Cunningham, and S. Puuronen. Dynamic integration of classifiers for handling concept drift. *Information Fusion*, 9(1):56–68, 2008.
19. G. Widmer and M. Kubat. Learning in the presence of concept drift and hidden contexts. *Machine learning*, 23(1):69–101, 1996.
20. W. Wilson. A comparison of alternatives for recurrent networks. In *Proceedings of the sixth Australian conference on Neural Networks (ACNN'93)*, pages 189–192, 1993.
21. Y. Yang, X. Wu, and X. Zhu. Mining in anticipation for concept change: Proactive-reactive prediction in data streams. *Data mining and knowledge discovery*, 13(3):261–289, 2006.