



# Early classification of time series

## Cost-based optimization criterion and algorithms

Youssef Achenchabe<sup>1,2</sup> · Alexis Bondu<sup>2</sup> · Antoine Cornuéjols<sup>1</sup>  · Asma Dachraoui<sup>1</sup>

Received: 25 May 2020 / Revised: 5 March 2021 / Accepted: 19 March 2021 /  
Published online: 2 June 2021

© The Author(s), under exclusive licence to Springer Science+Business Media LLC, part of Springer Nature 2021

### Abstract

An increasing number of applications require to recognize the class of an incoming time series as quickly as possible without unduly compromising the accuracy of the prediction. In this paper, we put forward a new optimization criterion which takes into account both the cost of misclassification and the cost of delaying the decision. Based on this optimization criterion, we derived a family of non-myopic algorithms which try to anticipate the expected future gain in information in balance with the cost of waiting. In one class of algorithms, unsupervised-based, the expectations use the clustering of time series, while in a second class, supervised-based, time series are grouped according to the confidence level of the classifier used to label them. Extensive experiments carried out on real datasets using a large range of delay cost functions show that the presented algorithms are able to solve the earliness vs. accuracy trade-off, with the supervised partition based approaches faring better than the unsupervised partition based ones. In addition, all these methods perform better in a wide variety of conditions than a state of the art method based on a myopic strategy which is recognized as being very competitive. Furthermore, our experiments show that the non-myopic feature of the proposed approaches explains in large part the obtained performances.

**Keywords** Early classification of time series · Cost estimation · Sequential decision making

---

Editors: Ira Assent, Carlotta Domeniconi, Aristides Gionis, Eyke Hüllermeier.

---

✉ Antoine Cornuéjols  
antoine.cornuejols@agroparistech.fr

Youssef Achenchabe  
youssef.achenchabe@agroparistech.fr

Alexis Bondu  
alexis.bondu@orange.com

Asma Dachraoui  
asma.dach@gmail.com

<sup>1</sup> UMR MIA-Paris, AgroParisTech, INRAe, Université Paris-Saclay, 75005 Paris, France

<sup>2</sup> Orange Labs, 44 Avenue de la république, Châtillon, France

## 1 Introduction

In emergency wards of hospitals (Mathukia et al. 2015), in control rooms of national or international electrical power grids (Dachraoui et al. 2013), in government councils assessing emergency situations, in all kinds of contexts, it is essential to make timely decisions in absence of complete knowledge of the true outcome (e.g. should the patient undergo a risky surgical operation?). The issue facing the decision makers is that, usually, the longer the decision is delayed, the clearer is the likely outcome (e.g. the critical or not critical state of the patient) but, also, the higher the cost that will be incurred if only because earlier decisions allow one to be better prepared. How to optimize online the tradeoff between the earliness and the accuracy of the decision is the object of the early classification of time series problem and this is what is addressed in this paper.

The work presented takes the earliness versus accuracy trade-off at face value in the spirit of Dachraoui et al. (2015) and formalizes it in a generic way. In this paper, we extend this previous work as a generic framework for the early classification of time series.

Furthermore, this previous work described a method based on the clustering algorithm K-means. However, the use of an unsupervised approach to capture the relevant groups of time series leaves important information aside. This is why, we propose here to resort to supervised techniques in order to get better prediction performances.

Interestingly, we claim that the problem of deciding online whether a prediction, and the attendant actions, should be made, or if it should be delayed, can be cast in the LUPI (Learning Under Privileged Information) framework (Vapnik and Vashist 2009). During the learning phase, the learner has access to the full knowledge about the training time series in addition to their class, while at testing time, only the incoming, and incomplete, time series is known. And decisions as whether to wait for additional measurements or not have to be made from this incomplete knowledge. The resulting optimization criterion can be used in a *non-myopic* procedure where estimates of future costs can be compared to the cost incurred if the decision was made at the current time.

The second objective of this paper is to design efficient optimization algorithms which implement the presented framework. We define a set of design choices that allow the categorization of possible *non-myopic* approaches (see Table 1). This enables us to define three novel algorithms by varying these choices. They are then carefully evaluated and compared in experiments in order to identify the best approach. In addition, the three proposed *non-myopic* approaches outperform the best known approach in the literature (Mori et al. 2017) evaluated on the same collection of datasets proposed by the authors.

The rest of the paper is organized as follows. The next section presents important works related to the early classification of time series problem. Section 3 reformulates in a generic way the cost-based formalism leading to an optimization criterion. Then, Sect. 4 presents the new methods that allow finer estimations of this criterion. These methods vary depending on how they take advantage of the training set of complete time series to estimate future costs of decision. This gives rise to a set of questions as what are the characteristics that most drive the performance up. Section 5 presents the experimental setup to answer these questions while the results with analyses are reported in Sect. 6. Section 7 compares the proposed approaches with the state of the art competing approach presented in Mori et al. (2017). Finally, Sect. 8 concludes by underlying the main findings of this research and by discussing directions for future works.

**Table 1** Overview of the design choices of the different approaches: each approach differing from the previous one by only one design choice

Approaches	Partition type	Partitions number	Anticipation type
ECONOMY-K (Dachraoui et al. 2015)	<b>Unsupervised</b> ( <i>K-means</i> )	<b>1</b> partition of full-length time series.	<b>Weak model:</b> the continuations of time series are used in the groups.
ECONOMY-MULTI-K	<b>Unsupervised</b> ( <i>K-means</i> )	<b>T</b> partitions corresponding to the different time steps.	<b>Weak model:</b> the continuations of time series are used in the groups.
ECONOMY- $\gamma$ -LITE ( <i>binary classification</i> )	<b>Supervised:</b> quantiles of the classifiers' confidence levels are used	<b>T</b> partitions corresponding to the different time steps.	<b>Weak model:</b> the continuations of time series are used in the groups.
ECONOMY- $\gamma$ ( <i>binary classification</i> )	<b>Supervised:</b> quantiles of the classifiers' confidence levels are used	<b>T</b> partitions corresponding to the different time steps.	<b>Markov Chain</b> technique is used to anticipate missing measurements.

## 2 Related work

Formally, we suppose that measurements are made available over time in a time series which, at time  $t$ , is  $\mathbf{x}_t = \langle x_1, \dots, x_t \rangle$  where  $x_t$  is the current measurement and the  $x_{i(1 \leq i \leq t)}$  belong to some input domain (e.g. temperature and blood pressure of a patient). We suppose furthermore that each time series can be ascribed to some class  $y \in \mathcal{Y}$  (e.g. patient who needs a surgical operation or not). The task is to make a prediction about the class of an incoming time series as early as possible because a cost is incurred at the time of the decision, where the cost function increases with time.

If the measurements in a time series are supposed independently and identically distributed (i.i.d.) according to a distribution of unknown “parameter”  $\theta$ , then the relevant framework is the one of *sequential decision making* and *optimal statistical decisions* (DeGroot 2005; Berger 1985). In this setting, the problem is to determine as soon as possible whether the measurements have been generated by a distribution of parameter  $\theta_0$  (hypothesis  $H_0$ ) or of parameter  $\theta_1$  (hypothesis  $H_1$ ) with  $\theta_0 \neq \theta_1$ . One technique especially has gained a wide exposition: Wald’s Sequential Probability Ratio Test (Wald and Wolfowitz 1948). The log-likelihood ratio  $R_t = \log \frac{P(\langle x_1^t, \dots, x_t^t \rangle | y=-1)}{P(\langle x_1^t, \dots, x_t^t \rangle | y=+1)}$  is computed and compared with two thresholds that are set according to the required error of the first kind  $\alpha$  (*false positive error*) and error of the second kind  $\beta$  (*false negative error*). Extensions to non-stationary distributions have been put forward (see Liu and Li 2013; Novikov 2008).

In the early classification of time series problem, however, the successive measurements are not supposed to be i.i.d. To compensate for this weaker assumption, it is assumed that a labeled training set exists made of time series of finite length  $T$ :  $\mathbf{x}_T^i = \langle x_1^i, \dots, x_T^i \rangle$  together with their corresponding labels,  $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1 \leq i \leq m}$ . Each measurement  $x_j^i$  can be multivariate.

In the test phase, the scenario goes as follows. At each time step  $t < T$ , a new measurement  $x_t$  is collected and a decision has to be made as whether to make a prediction now or to defer the decision to some future time step. When  $t = T$ , a decision is forced.

As already mentioned, the problem of deciding online whether a prediction, and the attendant actions, should be made, or if it should be delayed, can be cast in the LUPI

framework (Vapnik and Vashist 2009). In the following, we examine previous works on the early classification problem in this light.

To the best of our knowledge, Alonso González and Diez (2004) is the earliest paper explicitly mentioning “classification when only part of the series are presented to the classifier”, and the main thrust of it is to show how the boosting method can be employed to the classification of incomplete time series.

For many researchers, the question to solve is *can we classify an incomplete times series while ensuring some minimum probability threshold that the same decision would be made on the complete input?* To answer this question several approaches have been put forward.

One is to assume that the time series are generated i.i.d. according to some probability distribution, and to estimate the parameters of the class distributions from the training set. Once  $p(\mathbf{x}_T|\mathbf{x}_t)$  the conditional probability of the entire time series  $\mathbf{x}_T$  given an incomplete realization  $\mathbf{x}_t$  is estimated, it becomes possible to derive guarantees of the form:

$$p(h_T(\mathbf{X}_T) = y|\mathbf{x}_t) = \int_{\mathbf{x}_T \text{ s.t. } h_T(\mathbf{x}_T)=y} p(\mathbf{x}_T|\mathbf{x}_t) d\mathbf{x}_T \geq \epsilon$$

where  $\mathbf{X}_T$  is a random variable associated with the complete times series,  $\epsilon$  is a confidence threshold, and  $h_T(\cdot)$  is a classifier learnt over the training set  $\mathcal{S}$  of complete times series. At each time step  $t$ ,  $p(h_T(\mathbf{X}_T) = y|\mathbf{x}_t)$  is evaluated and the prediction is triggered if this term becomes greater than some predefined threshold. Anderson et al. (2012), and Parrish et al. (2013) present this method and propose ways to make the required estimations, in particular the mean and the covariance of the complete training data, when the time series are generated by Gaussian processes. It so far applies only with linear and quadratic classifiers.

Xing et al. (2009) do not make assumptions about the form of the underlying distributions on the time series. They propose to use a 1NN classifier that chooses the nearest training time series  $\mathbf{x}_t^i \in \mathcal{S}$  to the incoming one  $\mathbf{x}_t = \langle x_1, \dots, x_t \rangle$  to make its prediction. To determine for which time step  $t$  it is appropriate to make the prediction, the method is based on the idea of the *minimum prediction length* (MPL) of a time series. For a time series  $\mathbf{x}_t^i$ , one finds the set of every training time series  $\mathbf{x}_t^j$  that have  $\mathbf{x}_t^i$  as their one nearest neighbor. The MPL of  $\mathbf{x}_t^i$  is then defined as the smallest time index for which this set does not change when the rest of the time series  $\mathbf{x}_t^i$  is revealed. In the test phase, at time step  $t$ , it is deemed that  $\mathbf{x}_t$  can be safely labeled if its 1NN =  $\mathbf{x}_t^i$  for which the MPL is  $t$ . The idea is that from this point on, the prediction about  $\mathbf{x}_t$  should not change. The authors found experimentally that this procedure, called ECTS (*Early Classification of Times Series*), leads to too conservative estimations of the earliest safe time step for prediction. They therefore proposed heuristic means to lower the estimated values. The stability criterion acts in a way as a proxy for a measure of confidence in the prediction. Similarly, Mori et al. (2015) proposes a method where the evolution of the accuracy of a set of probabilistic classifiers is monitored over time, which allows the identification of timestamps from whence it seems safe to make predictions.

The authors of Gupta et al. (2020) are concerned with early classification of multivariate time series where the variables are not equally sampled in time. They focus on making a decision when all sensors have been processed and a high enough confidence level is attained.

Another line of research is concerned with finding good descriptors of the time series, especially on their starting subsequences, so that early predictions can be reliable because they would be based on relevant similarities on the time series. For instance, in the works of Xing et al. (2011), Ghalwash et al. (2014) and He et al. (2015), the principle is to look

for shapelets, subsequences of time series which can be used to distinguish time series of one class from another, so that it is possible to perform classification of time series as soon as possible.

By contrast, there are methods for early classification of time series that do not take advantage of the complete knowledge available in the training set. For instance, in Parrish et al. (2013), Hatami and Chira (2013), Ghalwash et al. (2012), a model  $h_t(\cdot)$  is learnt for each early timestamp and various stopping rules are defined in order to decide whether, at time  $t$ , a prediction should be made or not. The price to pay for being outside the LUPI framework is that decisions are made in a myopic fashion which may prevent one from seeing that a better trade-off between earliness and accuracy is achievable in the future.

This is also the case for the work presented in Mori et al. (2019). In the paper, the authors recognize the conflict between earliness and accuracy, and instead of setting a tradeoff in a single objective optimization criterion (Mori et al. 2017), they propose to keep it as a multi-objective criterion and to explore the Pareto front of the multiple dominating tradeoffs. Accordingly, they propose a family of triggering functions involving hyperparameters to be optimized for each tradeoff. This contrasts with approaches whereby the decision is made solely on the basis of a given confidence threshold which should be attained. However, the optimization criterion put forward is heuristic, supposes that the cost of delaying a decision is linear in time, and involves a complex setup. Most importantly, again, it is a myopic procedure which does not consider the foreseeable future. For all these apparent shortcomings, this method has been found to be quite effective, beating most competing methods in extensive experiments. This is why it is used as a reference method for comparison in this paper, as is done also in Rußwurm et al. (2019) which compares several techniques for early classification of time series.

The authors of Schäfer and Leser (2020) propose a system, called TEASER, that combines three components: (i) slave classifiers that estimate the class probabilities of the incoming series, (ii) a master classifier which assesses the confidence that one can have in the class that has the higher probability according to the slave classifier at the current time step, and finally (iii), the TEASER system which outputs a class if the master classifier has vetted this class for at least the last  $v$  time steps. It is apparent that the cost of delaying decision is not explicitly taken into account. The authors propose to optimize the harmonic mean between accuracy and earliness which indirectly corresponds to a particular tradeoff and a particular cost of delaying the decision. Furthermore, the method is fairly empirical.

In Dachraoui et al. (2015), for the first time, the problem of early classification of time series is cast as the optimization of a loss function which combines the expected cost of misclassification at the time of decision plus the cost of having delayed the decision thus far. Besides the fact that this optimization criterion is well-founded, it permits also to apply the LUPI framework because the expected costs for an incoming subsequence  $\mathbf{x}_t$  can be estimated for future time steps and thus a non-myopic decision procedure can be used. These expectations can indeed be learned from the training set of  $m$  complete time series  $\mathcal{S} = \{(\mathbf{x}_t^i, y_i)\}_{1 \leq i \leq m}$ .

### 3 Early classification as a cost optimization problem

Classifying time series with as little measurements as possible implies optimizing a trade-off. The less data is available for classification, the lower is the attainable accuracy in general. But waiting for more data implies incurring higher delay costs. There is therefore an

optimization problem to be solved involving both the classification accuracy, which translates into a misclassification cost, and the cost associated with gathering measurements. In the framework of online classification, this optimization problem becomes an online sequential decision making problem, where at each new time step, and a corresponding new piece of data, the system must decide whether to output a label for the incoming time series or to wait for more measurements. If the decision is made solely on the basis of the currently available information, the approach is myopic. On the other hand, if the decision involves some sort of prediction about the expected value of the optimization criterion in the future, the approach is non-myopic. This is what is allowed by the LUPI framework. Such a perspective was presented in Dachraoui et al. (2015). We generalize it in this section and lay the ground for three novel cost-based optimization criteria that are the object of this paper.

This approach assumes that the user provides two cost functions:

- $C_m(\hat{y}|y) : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$  is the *misclassification cost function* that defines the cost of predicting  $\hat{y}$  when the true class is  $y$ .
- $C_d(t) : \mathbb{R} \rightarrow \mathbb{R}$  is the *delay cost function* which is non decreasing over time.

Both of these costs are expressed in the same unit (e.g. in dollars) and convey the characteristics of the application domain as known by experts.

The expected cost of a decision at time step  $t$ , when  $\mathbf{x}_t$  is the incoming time series, can be expressed as:

$$\begin{aligned}
 f(\mathbf{x}_t) &= \mathbb{E}^t [C_m|x_t] + C_d(t) = \sum_{(y,\hat{y}) \in \mathcal{Y}^2} P_t(\hat{y}, y|x_t) C_m(\hat{y}|y) + C_d(t) \\
 &= \sum_{y \in \mathcal{Y}} P_t(y|\mathbf{x}_t) \sum_{\hat{y} \in \mathcal{Y}} P_t(\hat{y}|y, \mathbf{x}_t) C_m(\hat{y}|y) + C_d(t)
 \end{aligned}
 \tag{1}$$

The expectation comes both from the misclassification probability  $P_t(\hat{y}|y, \mathbf{x}_t)$  which can be estimated by the confusion matrix of the classifier  $h_t(\cdot)$  applied at time  $t$ , and the posterior probability of each class given the input incomplete time series estimate  $P_t(y|\mathbf{x}_t)$ .

If the input time series was fully observed, this cost could be computed for all time steps  $t \in \{1, \dots, T\}$ , and the optimal time  $t^*$  for triggering the classifier’s prediction would be:

$$t^* = \underset{t \in \{1, \dots, T\}}{\text{ArgMin}} f(\mathbf{x}_t)
 \tag{2}$$

But of course, this would defeat the whole purpose of early classification, as one would have to observe the entire time series before knowing what would have been the optimal decision time! Then, instead of waiting until the entire time series is known, at each time  $t$ , one could “look into the future” and guess what will be the best decision time. And if the estimated best decision time matches the current time step  $t$ , then the decision must be made. For the incoming time series  $\mathbf{x}_t$ , the expected cost at  $\tau$  time steps in the future is:

$$f_\tau(\mathbf{x}_t) = \sum_{y \in \mathcal{Y}} P_{t+\tau}(y|\mathbf{x}_t) \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|y, \mathbf{x}_{t+\tau}) C_m(\hat{y}|y) + C_d(t + \tau)
 \tag{3}$$

where  $\mathbf{x}_{t+\tau}$  is the foreseen continuation of  $\mathbf{x}_t$ . Accordingly, the best expected decision time in the future becomes:

$$\tau^* = \underset{\tau \in \{0, \dots, T-t\}}{\text{ArgMin}} f_\tau(\mathbf{x}_t)
 \tag{4}$$

and if  $\tau^* = 0$  the decision is instantly requested, and  $\hat{t}^* = t$  denotes the trigger time. The problem now is *how to predict  $\mathbf{x}_{t+\tau}$  from the knowledge of  $\mathbf{x}_t$* . Can the LUPI framework help? Yes it can. Figure 1a provides an overview of the principle in the case of a univariate time series. The “envelope” of its foreseeable futures can be learned using the training dataset of complete time series  $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1 \leq i \leq m}$ .

Importantly, the solution chosen to guess the “envelope” of the  $\mathbf{x}_{t+\tau}$  will also provide a way to estimate the terms  $P_{t+\tau}(\hat{y}|y, \mathbf{x}_{t+\tau})$  because a confusion matrix can be learned on this envelope.

However, estimating the likely outcomes of the incoming time series using a probabilistic forecasting model involves making assumptions (e.g. assuming that the residuals distribution is Gaussian). Another way to facilitate the use of this cost-based formalism is taken in Dachraoui et al. (2015) which consists in learning typical groups of time series from the training set, and then in predicting the likely continuations of  $\mathbf{x}_t$  with regard to these groups (see Fig. 1b).

Let us note  $\mathbf{g}_k$  the  $k$ -th typical groups of time series, Eq. (1) then can be re-expressed as:

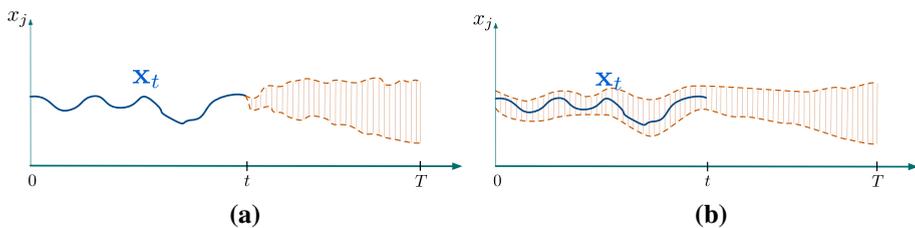
$$f(\mathbf{x}_t) = \sum_{\mathbf{g}_k \in \mathcal{G}} P_t(\mathbf{g}_k|\mathbf{x}_t) \sum_{y \in \mathcal{Y}} P_t(y|\mathbf{g}_k) \sum_{\hat{y} \in \mathcal{Y}} P_t(\hat{y}|y, \mathbf{g}_k) C_m(\hat{y}|y) + C_d(t) \tag{5}$$

And similarly, for Eq. (3):

$$f_\tau(\mathbf{x}_t) = \sum_{\mathbf{g}_k \in \mathcal{G}} P_t(\mathbf{g}_k|\mathbf{x}_t) \sum_{y \in \mathcal{Y}} P_t(y|\mathbf{g}_k) \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|y, \mathbf{g}_k) C_m(\hat{y}|y) + C_d(t + \tau) \tag{6}$$

Equation (6) can be easily interpreted by splitting it into two parts. The first term  $P_t(\mathbf{g}_k|\mathbf{x}_t)$  estimates the posterior probabilities of each group given  $\mathbf{x}_t$ . This term is estimated at time  $t$  and assumed to be constant over the time interval  $[t, t + \tau]$ . The next term expresses the expectations of the cost of misclassification over future possible continuations of  $\mathbf{x}_t$ . Namely, the second term  $P_t(y|\mathbf{g}_k)$  corresponds to the prior probabilities of class values within each group estimated at time  $t$ . And the third term  $P_{t+\tau}(\hat{y}|y, \mathbf{g}_k)$  estimates the probabilities of misclassification within each group, at time step  $t + \tau$ . The terms  $C_m(\hat{y}|y)$  and  $C_d(t + \tau)$  are the cost functions expressing properties of the domain of application.

In this general framework, several choices can be made to implement this optimization criteria. Foremost is the determination of relevant groups  $\mathbf{g}_k$  of time series from the complete training set  $\mathcal{S}$ . In what follows, we propose four different alternatives to anticipate the expected future misclassification costs.



**Fig. 1** **a** Given an incomplete time series  $\mathbf{x}_t$ , the objective is to try to guess the “envelope” of its foreseeable futures. Various methods can be used to do so. **b** The incoming time series  $\mathbf{x}_t$  is viewed as a member of or close to some group(s) of times series, and this is used to guess the “envelope” of its foreseeable futures

## 4 Anticipating the future: a key to the optimization criterion

This section presents three novel non-myopic approaches to solve cost-based optimization problem. These are three ways of Learning Using Privileged Information (LUPI) and therefore being able to foresee likeliest future values for the optimization criterion. The characteristics of these approaches are summarized in Table 1.

The three proposed approaches seek to better extract information about the foreseeable future of incoming time series so as to offer a better basis for deciding whether to label the time series at the current time step or to wait for more measurements. As can be seen from Table 1, each proposed approach can be viewed as an incremental modification of a previous method, from ECONOMY-K presented in Dachraoui et al. (2015) to ECONOMY- $\gamma$ .

Where ECONOMY-K is computing partitions of the time series based on the complete training time series, through a clustering process, ECONOMY-MULTI-K partitions incomplete times series for each time step in order to increase adaptiveness to the incoming and incomplete time series (see Sect. 4.2). Both methods do not use the labels of the time series in order to compute the partitions.

By contrast, both ECONOMY- $\gamma$ -LITE and ECONOMY- $\gamma$  use the labels in order to predict the likely future of the incoming time series. ECONOMY- $\gamma$ -LITE uses the level of confidence of the classifier at time  $t$  on the incoming  $\mathbf{x}_t$  in order to define groups of time series (see Sect. 4.3), whereas ECONOMY- $\gamma$  uses Markov chains in order to anticipate the likely future measurements on the time series (see Sect. 4.4). The current implementations of ECONOMY- $\gamma$ -LITE and of ECONOMY- $\gamma$  only accommodate binary classification tasks, but extensions to multi-class problems are envisioned for future work.

The number of groups  $K$  is a hyper-parameter shared by all of these approaches. In practice, it can be tuned using cross validation as detailed in Sect. 5.4. An open-source code is available for full reproducibility of the experiments presented in this paper: [www.github.com/YoussefAch/Economy](http://www.github.com/YoussefAch/Economy). The following sub-sections provide the main operating principles and key ideas for each ECONOMY approach.

### 4.1 ECONOMY-K

ECONOMY-K has been introduced in Dachraoui et al. (2015). The idea is to first identify groups  $\mathbf{g}_k$  of times series using a clustering algorithm, here K-means with Euclidian distance. Then, given an incoming time series  $\mathbf{x}_t$ , the memberships  $P(\mathbf{g}_k|\mathbf{x}_t)$  are estimated using a logistic function of a distance between  $\mathbf{x}_t$  and the centers of the clusters  $\mathbf{g}_k$ . In order to estimate the terms  $P_t(\hat{y}|\mathbf{y}, \mathbf{g}_k)$  of the confusion matrix for each time step  $t = 1, \dots, T$ , a collection of classifiers  $\{h_t\}_{t \in \{1, \dots, T\}}$  is learned using training sets  $\{\mathcal{S}^t\}_{t \in \{1, \dots, T\}}$  of time series truncated to their first  $t$  measurements.

In the end, implementing the ECONOMY-K approach relies on the following choices: (i) a distance function for K-means, (ii) a distance between an incomplete time series and clusters, (iii) a membership function estimating  $P(\mathbf{g}_k|\mathbf{x}_t)$ .

As explained in Sect. 3, Eq. (6) is used to estimate the cost of deciding for future time steps  $t + \tau$  ( $0 \leq \tau \leq T - t$ ), and if  $\tau^*$  given by Eq. (4) is equal to zero or  $t = T$ , then a decision is triggered, otherwise a new measurement  $x_{t+1}$  is made, and the decision mechanism is called again.

**Algorithm 1** Learn ECONOMY-K**Input:** K: number of groups

- 1: build the groups  $\mathbf{g}_k$  by using K-means on a set of full-length time series  $\mathcal{S}$
- 2: **for all**  $t \in \{1, \dots, T\}$  **do**
- 3:   learn a classifier  $h_t(\cdot)$  from a set of truncated time series  $\mathcal{S}^t$
- 4:   **for all**  $g_k \in \mathcal{G}$  **do**
- 5:     estimate the confusion matrix of  $h_t(\cdot)$  in the group  $g_k$
- 6:   **end for**
- 7: **end for**

Algorithm 1 provides the pseudo-code which summarizes the learning stage of the ECONOMY-K approach. The next sections describe the algorithms 2, 3 and 4 emphasizing for each of them the single difference with the previous algorithm presented (see the italicized bold line in each algorithm).

**Algorithm 2** Learn ECONOMY-MULTI-K**Input:** K: number of groups

- 1: **for all**  $t \in \{1, \dots, T\}$  **do**
- 2:   *build the groups  $\mathbf{g}_k^t$  by using K-means on a set of truncated time series  $\mathcal{S}^t$*
- 3:   learn a classifier  $h_t(\cdot)$  from a set of truncated time series  $\mathcal{S}^t$
- 4:   **for all**  $g_k \in \mathcal{G}$  **do**
- 5:     estimate the confusion matrix of  $h_t(\cdot)$  in the group  $g_k$
- 6:   **end for**
- 7: **end for**

**4.2 ECONOMY-MULTI-K**

Instead of grouping time series using their full-length descriptions, an alternative consists in computing the clusters  $\mathbf{g}_k^t$  for each time step  $t$  using training sets  $\{\mathcal{S}^t\}_{t \in \{1, \dots, T\}}$  of truncated time series from the training set  $\mathcal{S}$  (see line 2 of Algorithm 2). Indeed, clustering time series on the fly, at each time step, may allow for a increased adaptiveness to the specifics of the the beginning of the time series. The term  $P(\mathbf{g}_k|\mathbf{x}_t)$  in Eq. 6 then becomes  $P(\mathbf{g}_k^t|\mathbf{x}_t)$ . The cost of potential future decisions is now estimated based on the terms  $P_{t+\tau}(\hat{y}|y, \mathbf{g}_k^t)$ .

**4.3 ECONOMY- $\gamma$ -LITE**

In the previous approaches, the confusion matrix with the term  $P_{t+\tau}(\hat{y}|y, \mathbf{g}_k)$  in Eq. (6), is computed using time series in  $\mathbf{g}_k$  and potentially aggregates all confidence levels of  $h_{t+\tau}$ , corresponding to all possible values of the conditional probability  $p(y = 1|\mathbf{x}_{t+\tau})$ . If this confusion matrix was instead computed over time series that share approximately the same confidence level in their classification, the estimation of future decision costs could be much more precise. This is the motivation behind the algorithms ECONOMY- $\gamma$  and ECONOMY- $\gamma$ -LITE.

In these methods, the groups  $\mathbf{g}_k^t$  are obtained by stratifying the time series by confidence levels<sup>1</sup> of  $h_t$  (see line 3 of Algorithm 3). At each time step  $t$ , the confidence level  $p(h_t(\mathbf{x}_t) = 1)$  of the classifier can take a value in  $[0, 1]$ . Examining the confidence levels for all time series in the validation set  $\mathcal{S}^t$  truncated to the first  $t$  observations, we can discretize the interval  $[0, 1]$  into  $K$  equal frequency intervals, denoted  $\{I_t^1, \dots, I_t^K\}$ . For instance, if  $K = 5$ , and  $|\mathcal{S}^t| = 1000$ , the intervals  $I_t^1 = [0, 0.30[$ ,  $I_t^2 = [0.30, 0.45[$ ,  $I_t^3 = [0.45, 0.58[$ ,  $I_t^4 = [0.58, 0.83[$ ,  $I_t^5 = [0.83, 1]$  could each correspond to 200 training time series. The discretization of confidence levels into equal frequency intervals corrects any bias in the calibration of  $h_t$ , in a similar way to isotonic calibration (Flach 2016).

Given an incoming time series  $\mathbf{x}_t$ , the classifier  $h_t$  is used to get an estimate of  $p(y = 1|\mathbf{x}_t)$  and determine the group  $\mathbf{g}_k^t$  to which  $\mathbf{x}_t$  belongs. The algorithm is the same as ECONOMY-MULTI-K, only with the groups  $\mathbf{g}_k^t$  obtained in a supervised way by leveraging the information about the membership to the classes.

One can notice that, in addition to the expected gain in performance due to a more informed grouping of time series than in the clustering-based approaches, this method as well as ECONOMY- $\gamma$ , does not require (i) the choice of a distance function for K-means, nor (ii) the determination of another distance between an incomplete time series  $\mathbf{x}_t$  and a cluster of full-length time series, and finally (iii) neither the choice of a membership function in order to estimate  $P(\mathbf{g}_k|\mathbf{x}_t)$ . The approach is therefore much simpler to implement.

---

**Algorithm 3** Learn ECONOMY- $\gamma$ -LITE

---

**Input:**  $K$ : number of groups

- 1: **for all**  $t \in \{1, \dots, T\}$  **do**
  - 2:   learn a classifier  $h_t(\cdot)$  from a set of truncated time series  $\mathcal{S}^t$
  - 3:   *build the groups  $\mathbf{g}_k^t$  by discretizing the confidence level of  $h_t(\cdot)$  into  $K$  intervals*
  - 4:   **for all**  $g_k \in \mathcal{G}$  **do**
  - 5:     estimate the confusion matrix of  $h_t(\cdot)$  in the group  $g_k$
  - 6:   **end for**
  - 7: **end for**
- 

**4.4 ECONOMY- $\gamma$**

ECONOMY- $\gamma$  uses the ECONOMY- $\gamma$ -LITE principle to assign an incoming time series  $\mathbf{x}_t$  to a given group  $\mathbf{g}_k^t$ , but it tries to get better estimates of the future terms  $P_{t+\tau}(\hat{y}|y, \mathbf{g}_k^t)$  of the confusion matrices by replacing  $\mathbf{g}_k^t$  by a projection  $\mathbf{g}_k^{t+\tau}$  into the future as a probability distribution over the confidence intervals of  $h_{t+\tau}$ .

Let us call  $\vec{\gamma}_t = (\gamma_t^1, \dots, \gamma_t^K)^\top$  the real-value vector of  $K$  components  $\gamma_t^i$ , where each of the components is the probability that  $p(h_t(\mathbf{x}_t) \in I_t^i)$ . For instance, in Fig. 2,  $\vec{\gamma}_t = (0, 1, 0, 0, 0)^\top$ , all components are zero except  $\gamma_t^2 = 1$ .

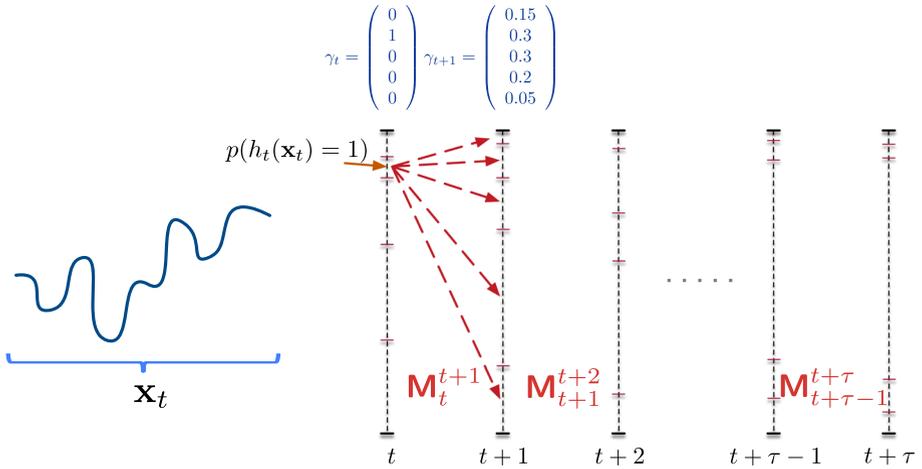
We would like to compute the vectors  $\vec{\gamma}_{t+\tau}$  ( $0 < \tau \leq T - t$ ) consisting of the components:

$$\gamma_{t+\tau}^j = p(h_{t+\tau}(\mathbf{x}_{t+\tau}) \in I_{t+\tau}^j) \tag{7}$$

In ECONOMY- $\gamma$ , we propose to estimate  $\gamma_{t+\tau}^j$  by using the  $K \times K$  transitions matrices  $\{\mathbf{M}_{t'}^{t'+1}\}_{t' \in \{1, \dots, T-1\}}$  from  $\vec{\gamma}_{t'}$  to  $\vec{\gamma}_{t'+1}$  (see line 4 of Algorithm 4) where each component of the matrix is estimated by:

---

<sup>1</sup> This restricts these methods to binary classification problems.



**Fig. 2** ECONOMY- $\gamma$ , computing the probability distribution  $p(\gamma_{t+\tau}|\gamma_t)$ . Here  $h_t(\mathbf{x}_t)$  falls in the second confidence level interval. Given a supposed learned transition matrix  $M_t^{t+1}$ , the next vector of confidence levels will be  $(0.15, 0.3, 0.3, 0.2, 0.05)^T$

$$m_{ij} = p(h_{t'+1}(\mathbf{x}_{t'+1}) \in I_{t'+1}^j \mid h_{t'}(\mathbf{x}_{t'}) \in I_{t'}^i) \tag{8}$$

given a validation set of time series. At time step  $t$ , and from  $\bar{\gamma}_t$  it then becomes possible to compute  $\bar{\gamma}_{t+\tau}$  by:

$$\bar{\gamma}_{t+\tau} = \bar{\gamma}_t^T \prod_{s=0}^{\tau-1} M_{t+s}^{t+s+1} \tag{9}$$

Like in Eq. (6), the future expected costs of decision are estimated through:

$$f_\tau(\mathbf{x}_t) = \underbrace{\sum_{j=1}^K \gamma_{t+\tau}^j}_{(1)} \sum_{y \in \mathcal{Y}} P(y|I_{t+\tau}^j) \underbrace{\sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|y, I_{t+\tau}^j) C_m(\hat{y}|y) + C_d(t+\tau)}_{(2)} \tag{10}$$

- (1): for all confidence intervals  $I_{t+\tau}^j$  of  $h_{t+\tau}$
- (2): probability of misclassification when  $h_{t+\tau}(\mathbf{x}_{t+\tau}) \in I_{t+\tau}^j$

Again, a decision is triggered at time  $\hat{\tau}^* = t$ , if  $\tau^* = \text{ArgMin}_{\tau \in \{0, \dots, T-t\}} f_\tau(\mathbf{x}_t)$  is found to be 0.

---

**Algorithm 4** Learn ECONOMY- $\gamma$

---

**Input:**  $K$ : number of groups  
 1: **for all**  $t \in \{1, \dots, T\}$  **do**  
 2:   learn a classifier  $h_t(\cdot)$  from a set of truncated time series  $\mathcal{S}^t$   
 3:   build the groups  $\mathfrak{g}_k^t$  by discretizing the confidence level of  $h_t(\cdot)$  into  $K$  intervals  
 4:   *estimate the transition matrix*  $\mathbf{M}_t^{t+1}$   
 5: **end for**

---

**4.5 Complexity analysis**

We provide here an analysis of the computational complexities of the proposed algorithms, first in relation to the learning stage, and then with regard to the inference phase.

These complexities are expressed in a generic way, i.e. regardless of the classifier (of complexity denoted *Learn*) and clustering (*Clustering*) algorithms used. The single difference between two successive algorithms is underlined using bold letters in the following expressions that relate to the **learning phase**.

ECO-K:            $\mathcal{O}(T.Learn + Clustering + T.|\mathcal{S}|.Predict + |Y|.K.|\mathcal{S}|)$   
 ECO-MULTI-K:    $\mathcal{O}(T.Learn + \mathbf{T}.Clustering + T.|\mathcal{S}|.Predict + |Y|.K.|\mathcal{S}|)$   
 ECO- $\gamma$ -LITE:    $\mathcal{O}(T.Learn + T.|\mathcal{S}|.\log(|\mathcal{S}|)) + T.|\mathcal{S}|.Predict + |Y|.K.|\mathcal{S}|)$   
 ECO- $\gamma$ :            $\mathcal{O}(T.Learn + T.|\mathcal{S}|.\log(|\mathcal{S}|)) + T.|\mathcal{S}|.Predict + |Y|.K.|\mathcal{S}| + |\mathcal{S}|^2.K^2)$

ECONOMY-K learns a collection of  $T$  classifiers and build a partition of full-length time series using a *Clustering* algorithm with a  $\mathcal{O}(T.Learn + Clustering)$  complexity. Then, confusion matrices are computed with a  $\mathcal{O}(T.|\mathcal{S}|.Predict)$  complexity and the prior probability for each class in each group is computed with a  $\mathcal{O}(|Y|.K.|\mathcal{S}|)$  complexity.

ECONOMY-MULTI-K has almost the same complexity and only differs by computing  $\mathbf{T}$  different partitions.

ECONOMY- $\gamma$ -LITE discretizes the outputs of the classifiers by sorting time series according to their confidence level with a  $\mathcal{O}(T.|\mathcal{S}|.\log(|\mathcal{S}|))$  complexity.

Finally, ECONOMY- $\gamma$  adds the estimation of the transition matrices computed with a  $\mathcal{O}(|\mathcal{S}|^2.K^2)$  complexity.

During the **inference phase**, a new measurement is received at each time step and the future costs must be estimated. All the approaches carry out this estimate in a  $\mathcal{O}(T. |Y|^2.K)$  complexity, except ECONOMY- $\gamma$  that uses a matrix-vector product for transition matrix to estimate future costs with a  $\mathcal{O}(T. |Y|^2.K + K^2)$  time complexity.

**5 Description of the experiments**

**5.1 Goal of the experiments**

The approaches presented all rely on the estimation of the best decision time based according to a cost-based criterion which expresses the expected misclassification cost for future time steps plus a delay cost. This is the basis of these non-myopic strategies.

A *first set of questions* regards the importance and *impact of the various design choices* that distinguish the four ECONOMY algorithms (see first part of the experiments in Sect. 6.2):

1. Time series are partitioned in an unsupervised way for ECONOMY-K and ECONOMY-MULTI-K and in a supervised mode for ECONOMY- $\gamma$ -LITE and ECONOMY- $\gamma$ . Is one approach better than the other?
2. On a finer grain, is it better to cluster series using their full-length descriptions as in ECONOMY-K or on their truncated description at each time step  $t$  as in ECONOMY-multi-K?
3. Is it useful to try to have a more precise anticipation of the future of the incoming time series as is done in ECONOMY- $\gamma$  compared with the simpler albeit coarser approach of ECONOMY- $\gamma$ -LITE?
4. How distant the cost incurred is from the ideal optimal cost that one would have paid if one had known the whole series and therefore the best decision time? This is akin to a *regret* for not having a perfect a posteriori knowledge.

A *second set of questions* is whether developing non-myopic approaches brings performance gains compared to state of the art approaches that are myopic?

The second part of experiments (Sect. 7.2), therefore, compares the best ECONOMY approach with Mori et al. (2017) which has state of the art performance, as confirmed by a recent paper (Rußwurm et al. 2019).

Our experiments are designed to answer these two sets of questions. This is why in Sect. 5.3, *two different collections of datasets* are used in the experiments dedicated to each set of questions.

## 5.2 Evaluation criterion

In order to compare the methods, it is important to consider a criterion which expresses its worth for the final user. We define a new evaluation criterion used in our experiments both to optimize  $K$  on a validation set and to evaluate the early classification approaches on a test set. In actual use, ultimately, the value of employing an early classification method corresponds to the average cost that is incurred using it. For a given dataset  $\mathcal{S}$ , it is defined as follows:

$$\text{AvgCost}(\mathcal{S}) = \frac{1}{|\mathcal{S}|} \sum_{(x_t, y) \in \mathcal{S}} (C_m(h_{\hat{t}^*}(x_{\hat{t}^*})|y) + C_d(\hat{t}^*)) \quad (11)$$

where  $\hat{t}^*$  is the decision time chosen by the method as the one optimizing the trade-off between earliness and accuracy. In our experiments, *AvgCost* is evaluated for each dataset and for each early classification method. Statistical tests allow us to detect significant difference in performance.

## 5.3 Datasets

Two distinct collections of datasets are used in our experiments.

*The datasets for the comparison of the ECONOMY approaches:* With respect to the first set of questions presented in Sect. 5.1, about the role of the various design choices, one cannot easily measure differences of performance if the datasets only include time series that are easy to classify very early or that are hard to classify even when the whole series are known. Indeed, if this happens, all methods either decide early to output a label or wait until the end, and their performances are almost indistinguishable. In order, then, to be able to measure differences between the various online decision methods, we excluded datasets with these characteristics.

All the selected datasets come from the UEA & UCR Time Series Classification Repository<sup>2</sup>. First, we removed potentially correlated datasets since it is important to select only independent datasets for the use of statistical tests. We identified almost identical dataset names and sizes. For instance, the datasets “Ford A” and “Ford B” contain the same number of time series with the same length. In this case, we keep only one dataset chosen at random. Then, we learned a collection of classifiers for the remaining datasets and manually removed those for which the successive classifiers did not improve their quality over time. More specifically, we plotted the Cohen’s  $\kappa$  score (Cohen 1960) for each possible lengths of the input time series. We removed the datasets with low and almost constant quality of classifiers over time. Please note that the removed datasets are overwhelmingly the smallest, as the low number of training examples generally leads to poor quality classifiers. The 34 selected datasets and their description are available at <https://tinyurl.com/ycmbxurq>.

*The datasets for comparisons with the state of the art methods:*

In order to be able to make direct comparison with the method described in Mori et al. (2017) we use the same datasets as they did. This benchmark consists of 45 datasets of variable sizes that come from a variety of application areas. This collection of datasets has also been used in Schäfer and Leser (2020) and Mori et al. (2019), making our experiences easily comparable to previous works.

*Dataset preparation:* First, the training and test sets were merged for each dataset to overcome the possibly unbalanced or biased split of the original data files. The remaining datasets were then transformed into binary classes since ECONOMY- $\gamma$  and ECONOMY- $\gamma$ -LITE are limited to binary classification. This was done by retaining the majority class and merging all the others. In order to reduce the computation time of the experiments and to compare datasets with time series of different lengths, we trained a classifier every 5% of the total length of the time series, instead of one classifier per time step, as done in Mori et al. (2017). Furthermore, for each dataset and for each possible length (i.e. 5%, 10%, 15% ... of the total length), we extracted 60 features<sup>3</sup> from the corresponding truncated time series in order to train the associated classifiers. To do this, we used the Time Series Feature Extraction Library (Barandas et al. 2020), which automatically extracts features on the statistical, temporal and spectral domains.

## 5.4 Experimental protocol

The datasets were divided by uniformly selecting 70% of the examples for the training set and using the remaining 30% for the test set. Furthermore, the training sets were divided into three disjoint subsets as follows: (subset a) 40% for training the various classifiers  $\{h_t\}_{t \in \{1, \dots, T\}}$ ; (subset b) 40% for learning the meta parameters; (subset c) 20% to optimize the number of groups in  $\mathcal{G}$ .

(subset a) *Learning the collection of classifiers:* for each dataset, the classifiers corresponding to the possible lengths of the input time series (i.e. every 5% of the total length) were learned. The XGboost Python library<sup>4</sup> was used, keeping the default values for the hyper-parameters.

<sup>2</sup> Available at : <http://www.timeseriesclassification.com>.

<sup>3</sup> More details are available in: [https://docs.google.com/spreadsheets/d/13u7L\\_5IX3XxFuq\\_SnbOZF1dXQfcBB0wR3PXhvevhPYA/](https://docs.google.com/spreadsheets/d/13u7L_5IX3XxFuq_SnbOZF1dXQfcBB0wR3PXhvevhPYA/).

<sup>4</sup> XGBoost is available in: <https://xgboost.readthedocs.io>.

(subset b) *Learning the meta-parameters*: they were learned for each ECONOMY approach, except the parameter  $K$  which is optimized using (subset c). For instance, a meta-model learned by the ECONOMY- $\gamma$  approach consists of: (i) the discretization into  $K$  intervals of the confidence level for each classifier (one for each possible length), and (ii) the transition matrices between a time step to the next one (i.e. every 5% of the time series length).

(subset c) *Optimizing the number of groups  $K$* : the ECONOMY algorithms were trained by varying the number of groups between 1 and 20 and evaluated by the  $AvgCost(.)$  criterion which represents the average cost actually paid by the user (see Eq. (11)). The value of  $K$  which minimizes the  $AvgCost(.)$  criterion has been kept.

*Costs setting*: the misclassification cost was set in the same way for all datasets:  $C_m(\hat{y}|y) = 1$  if  $\hat{y} \neq y$ , and  $= 0$  otherwise. The delay cost  $C_d(t)$  is provided by the domain experts in actual use case. In the absence of this knowledge, we define it as a linear function of time, with coefficient, or slope,  $\alpha$ :

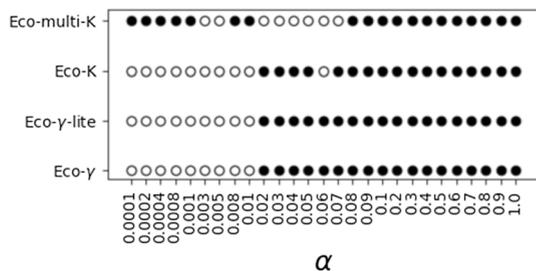
$$C_d(t) = \alpha \times \frac{t}{T} \tag{12}$$

The larger the  $\alpha$  coefficient, the more costly it is to wait for more measurements in the incoming time series. The delay cost  $C_d(t)$  is obviously of paramount importance to control the best decision time. If  $\alpha$  is very low, it does not hurt to wait for the whole time series to be known and  $t^* = T$ . If, on the contrary,  $\alpha$  is very high, the gain in misclassification cost obtained thanks to more observations cannot compensate for the increase of the delay cost, and it is better to make a decision at the beginning of the observations. Our experiments were run over a three ranges of values of  $\alpha$ : low time cost with  $\alpha \in [1e-04, 2e-04, 4e-04, 8e-04, 1e-03, 3e-03, 5e-03, 8e-03]$ ; medium time cost with  $\alpha \in [0.01, 0.02, 0.03, 0.04, 0.05, 0.06, 0.07, 0.08, 0.09]$ ; high time cost with  $\alpha \in [0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1.0]$ .

### 6 Best design choices for ECONOMY

This section presents the results of the experiments aimed at identifying the best design choices for the ECONOMY approaches (see Sect. 5.3).

**Fig. 3** Success of adapting the trigger times—Wilcoxon signed-rank test results for different values of  $\alpha$ : black dots indicate success and circles failures



## 6.1 Comparison of the ECONOMY approaches with a non adaptive baseline

As a first sanity test, it is interesting to see if the ECONOMY algorithms do indeed adapt the decision time to the incoming time series, or if they treat them all the same. In order to perform this test, each of the ECONOMY approaches is run once in its adaptive mode, and once made unable to adapt by forcing the number of groups  $K = 1$  (there is thus no difference made between the series).

The ECONOMY approaches are trained on the 34 selected datasets by varying the value of  $\alpha$ , and then, evaluated on the test sets using the *AvgCost* criterion. The Wilcoxon signed-rank test is used to assess whether the observed performance gap is significant. Figure 3 presents the results of the Wilcoxon signed-rank test for each ECONOMY approach, applied over the 34 datasets by varying the values of  $\alpha$ .

In the range  $\alpha \in [0.0001, 0.01]$ , ECONOMY-MULTI-K is the only approach that succeed in adapting its trigger times. By contrast, in the range  $\alpha \in [0.02, 0.1]$ , it appears that the ECONOMY approaches actually succeed in adapting their trigger times, with the exception of ECONOMY-MULTI-K which fails this test one-third of the time and behaves rather erratically when  $\alpha$  varies.

At the end, these approaches succeed in improving performance by adapting their trigger times, differing in their range of success.

## 6.2 Comparison of the ECONOMY approaches

### (a) Comparison with respect to the average decision cost

The *AvgCost* criterion was evaluated on the 34 test sets, and  $\alpha$  was adjusted for each dataset in order to reveal the greatest differences in performance between the best and worst approach (see Table 2 for more details). The Nemenyi test (Nemenyi 1962) was used to rank the different ECONOMY approaches in terms of average decision cost. The Nemenyi test consists of two successive steps. First, the Friedman test is applied to the average decision cost of competing approaches to determine whether their overall performance is similar. If not, the post-hoc test is applied to determine groups of approaches whose overall performance is significantly different from that of the other groups.

Figure 4a, reporting the results of the Nemenyi test, shows two groups of methods of which the performances are significantly different. Specifically, the ECONOMY- $\gamma$  and ECONOMY- $\gamma$ -LITE methods exhibit much better average decision costs than ECONOMY-K and ECONOMY-MULTI-K.

Figure 4b shows pairwise comparison using the Wilcoxon signed-rank test between the approaches. The small black squares identify pairs of approaches that do not differ significantly in performance. It is thus apparent that ECONOMY- $\gamma$  performs significantly better than ECONOMY- $\gamma$ -LITE.

### (b) Comparison with respect to the earliness of the decision time

In the following, the *earliness* of early classification approaches is evaluated using the median of the trigger times normalized by the length of the series, defined by  $earliness = med\{\hat{t}^*_{i^*}\}/T$  (see Table 2). Figure 5a shows that ECONOMY- $\gamma$ , on average, triggers its decision earlier than the competing methods, followed by ECONOMY- $\gamma$ -LITE. Furthermore, according to the Wilcoxon signed-rank test, this difference is significant compared to the other ECONOMY approaches (Fig. 5b).

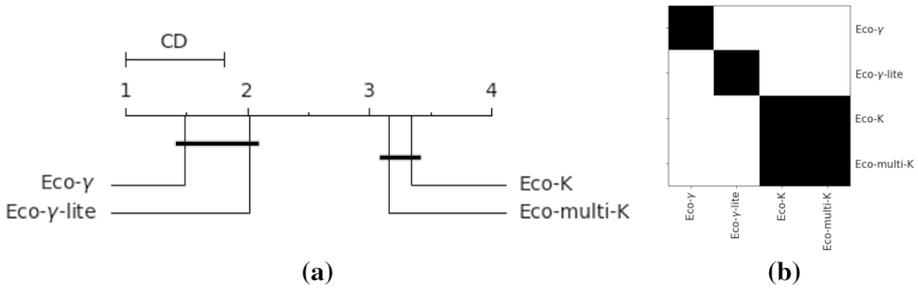
(c) Comparison with respect to the predictive performance of the algorithms The predictive performance is evaluated using the Cohen's *kappa* score (Cohen 1960) computed

**Table 2** Details of experimental results for each dataset

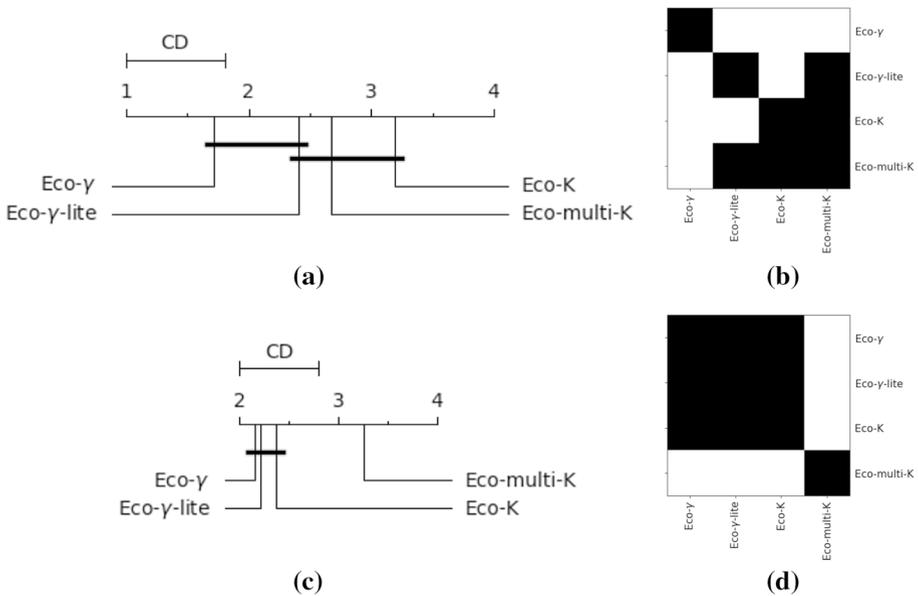
Datasets	$\alpha$	Nb groups: K				AvgCost				Median $\hat{t}^*$				Kappa					
		E- $\gamma$		E-K		E-m-K		Eco- $\gamma$		Eco-K		Eco-m-K		Eco- $\gamma$		Eco-K		Eco-m-K	
		E- $\gamma$	E- $\gamma$ -l	E-K	E-K	E-m-K	E-m-K	Eco- $\gamma$	Eco- $\gamma$ -l	Eco-K	Eco-K	Eco-m-K	Eco-m-K	Eco- $\gamma$	Eco- $\gamma$ -l	Eco-K	Eco-K	Eco-m-K	Eco-m-K
<i>CBF</i>	0.8	11	16	7	4	0.167	<b>0.163</b>	0.231	0.222	<b>0.09</b>	0.14	0.23	0.23	0.83	<b>0.89</b>	0.87	<b>0.89</b>	0.87	<b>0.89</b>
<i>ChlorineConcentration</i>	0.4	14	11	20	1	<b>0.252</b>	0.302	0.303	0.322	<b>0.10</b>	0.14	0.05	0.14	<b>0.59</b>	0.50	0.46	0.50	0.46	0.46
<i>CinCECGTorso</i>	0.1	20	15	1	1	0.014	<b>0.012</b>	0.021	0.021	0.05	0.05	0.05	0.05	<b>0.98</b>	<b>0.98</b>	0.96	<b>0.98</b>	0.96	0.96
<i>Crop</i>	0.06	10	19	19	18	<b>0.028</b>	0.031	0.031	0.042	0.04	0.04	0.04	0.04	0.70	0.68	0.65	0.70	0.65	<b>0.71</b>
<i>ECG5000</i>	0.5	18	15	1	19	<b>0.046</b>	0.051	0.065	0.064	0.05	0.05	0.05	0.05	<b>0.97</b>	0.96	0.92	<b>0.97</b>	0.96	0.92
<i>ECGFiveDays</i>	0.3	5	7	14	7	0.115	<b>0.097</b>	0.163	0.125	<b>0.09</b>	0.13	0.49	0.18	0.83	0.92	<b>0.95</b>	0.83	<b>0.95</b>	0.92
<i>ElectricDevices</i>	0.1	18	4	13	10	0.114	<b>0.109</b>	0.150	0.146	0.08	<b>0.04</b>	0.83	0.46	0.75	0.80	<b>0.82</b>	0.75	<b>0.82</b>	0.75
<i>FaceAll</i>	0.01	16	16	19	16	<b>0.002</b>	<b>0.002</b>	0.006	0.004	<b>0.05</b>	<b>0.05</b>	0.14	0.14	<b>0.99</b>	<b>0.99</b>	0.98	<b>0.99</b>	0.98	<b>0.99</b>
<i>FacesUCR</i>	0.5	11	2	13	18	<b>0.093</b>	0.111	0.123	0.133	<b>0.05</b>	<b>0.05</b>	0.14	0.14	<b>0.77</b>	0.67	0.69	<b>0.77</b>	0.67	0.72
<i>FiftyWords</i>	0.5	15	2	13	7	<b>0.104</b>	0.119	0.148	0.160	<b>0.05</b>	0.10	<b>0.05</b>	<b>0.05</b>	<b>0.68</b>	0.66	0.29	<b>0.68</b>	0.66	0.55
<i>Fora</i>	0.3	14	7	7	2	<b>0.114</b>	0.118	0.155	0.164	0.15	0.15	0.15	0.15	<b>0.88</b>	<b>0.88</b>	0.81	<b>0.88</b>	0.81	0.79
<i>FreezerRegularTrain</i>	0.2	17	4	7	17	<b>0.016</b>	0.021	0.024	0.025	0.05	0.05	0.05	0.05	<b>0.99</b>	0.98	0.98	<b>0.99</b>	0.98	0.98
<i>HamdOutlines</i>	0.01	17	4	4	10	<b>0.109</b>	0.113	0.124	0.160	<b>0.25</b>	0.40	0.50	1.00	<b>0.76</b>	<b>0.76</b>	0.73	<b>0.76</b>	0.73	0.66
<i>InsectWingbeatSound</i>	1	19	1	1	1	0.151	<b>0.130</b>	<b>0.130</b>	<b>0.130</b>	0.05	0.05	0.05	0.05	0.25	<b>0.29</b>	<b>0.29</b>	0.25	<b>0.29</b>	<b>0.29</b>
<i>ItalyPowerDemand</i>	0.5	9	14	1	1	<b>0.265</b>	0.302	0.349	0.349	<b>0.04</b>	0.08	0.33	0.33	0.61	0.54	<b>0.64</b>	0.61	<b>0.64</b>	<b>0.64</b>
<i>Mallat</i>	0.08	9	12	10	17	<b>0.016</b>	0.019	0.034	0.028	<b>0.05</b>	<b>0.05</b>	0.30	0.15	0.97	0.96	0.93	<b>0.97</b>	0.96	<b>0.99</b>
<i>MedicalImages</i>	0.07	12	11	19	20	<b>0.210</b>	0.238	0.240	0.279	<b>0.08</b>	0.16	0.12	0.69	<b>0.61</b>	0.57	0.55	<b>0.61</b>	0.57	0.54
<i>MelbournePedestrian</i>	0.8	12	3	1	11	0.126	<b>0.111</b>	0.134	0.122	0.04	0.04	0.04	0.04	0.71	0.85	0.70	0.71	0.85	<b>0.86</b>
<i>MixedShapesRegularTrain</i>	0.1	20	5	3	17	0.030	<b>0.020</b>	0.047	0.044	<b>0.05</b>	0.10	0.15	0.10	0.94	<b>0.97</b>	0.91	0.94	<b>0.97</b>	0.91
<i>NoteStrain</i>	0.4	7	18	11	18	<b>0.128</b>	0.142	0.156	0.178	<b>0.10</b>	0.14	0.14	0.29	<b>0.86</b>	0.84	0.80	<b>0.86</b>	0.84	0.82
<i>NonInvasiveFetalECGT2</i>	0.04	5	19	9	7	<b>0.011</b>	<b>0.011</b>	0.014	0.012	0.05	0.05	0.05	0.05	0.81	0.80	0.71	<b>0.81</b>	0.80	<b>0.83</b>
<i>PhalangesOutlinesCorrect</i>	0.2	5	16	4	1	<b>0.297</b>	0.348	0.320	0.316	<b>0.15</b>	0.25	<b>0.15</b>	<b>0.15</b>	<b>0.4</b>	0.29	0.36	<b>0.4</b>	0.29	0.35
<i>ProximalPhalanxOutlineCo</i>	0.5	6	4	7	17	<b>0.267</b>	0.300	0.300	0.306	<b>0.05</b>	0.10	<b>0.05</b>	0.20	<b>0.43</b>	0.39	0.36	<b>0.43</b>	0.39	0.55

**Table 2** (continued)

Datasets	$\alpha$	Nb groups: K						AvgCost						Median $\hat{f}^*$						Kappa					
		E- $\gamma$		E- $\gamma$ -l		E-K		E-m-K		Eco- $\gamma$		Eco- $\gamma$ -l		Eco-K		Eco-m-K		Eco- $\gamma$		Eco- $\gamma$ -l		Eco-K		Eco-m-K	
		E- $\gamma$	E- $\gamma$ -l	E- $\gamma$ -l	E-K	E-K	E-m-K	Eco- $\gamma$	Eco- $\gamma$ -l	Eco- $\gamma$ -l	Eco-K	Eco-K	Eco-m-K	Eco- $\gamma$	Eco- $\gamma$ -l	Eco- $\gamma$ -l	Eco-K	Eco-K	Eco-m-K	Eco- $\gamma$	Eco- $\gamma$ -l	Eco- $\gamma$ -l	Eco-K	Eco-K	Eco-m-K
<i>SengHandGenderCh2</i>	0.3	4	8	8	1	7	0.176	0.176	<b>0.174</b>	0.265	<b>0.15</b>	0.20	0.20	0.20	0.40	0.74	<b>0.77</b>	0.76	0.64						
<i>SonyAIBORobotSurface2</i>	0.8	7	10	17	15	15	<b>0.176</b>	0.210	0.228	0.208	<b>0.09</b>	0.18	0.18	0.18	0.18	0.84	<b>0.85</b>	0.82	0.84						
<i>StarLightCurves</i>	0.3	17	12	18	3	3	<b>0.062</b>	0.067	0.089	0.095	<b>0.05</b>	<b>0.05</b>	0.15	0.20	0.93	0.92	0.89	0.91							
<i>Strawberry</i>	0.6	2	9	17	8	8	0.195	<b>0.193</b>	0.251	0.201	0.09	0.14	<b>0.05</b>	0.19	0.78	0.77	0.59	<b>0.83</b>							
<i>Symbols</i>	0.2	15	15	15	4	4	0.034	<b>0.031</b>	0.057	0.042	0.05	0.05	0.05	0.05	0.97	0.98	0.89	<b>0.99</b>							
<i>TwoLeadECG</i>	0.9	15	16	19	19	19	<b>0.160</b>	0.168	0.197	0.185	<b>0.10</b>	<b>0.10</b>	0.15	0.15	0.88	<b>0.90</b>	0.86	0.86							
<i>TwoPatterns</i>	0.08	13	11	10	19	19	0.079	<b>0.065</b>	0.104	0.104	<b>0.56</b>	0.66	0.89	0.89	0.90	<b>0.95</b>	0.91	0.91							
<i>UWaveGestureLibraryX</i>	0.5	12	1	18	17	17	<b>0.130</b>	0.148	0.153	0.126	0.05	0.05	0.05	0.05	0.42	0.14	0.22	<b>0.62</b>							
<i>Wafer</i>	0.2	18	18	15	6	6	<b>0.010</b>	<b>0.010</b>	0.011	<b>0.010</b>	0.05	0.05	0.05	0.05	<b>1.00</b>	<b>1.00</b>	0.99	<b>1.00</b>							
<i>WordSynonyms</i>	0.6	19	20	20	1	1	<b>0.177</b>	0.218	0.209	0.307	<b>0.10</b>	<b>0.10</b>	<b>0.10</b>	0.14	<b>0.63</b>	0.53	0.54	0.36							
<i>Yoga</i>	0.03	14	8	8	8	20	<b>0.168</b>	0.246	0.183	0.193	<b>0.10</b>	0.49	0.49	0.49	<b>0.67</b>	0.54	0.66	0.65							



**Fig. 4** Evaluation based on *AvgCost*: **a** Nemenyi test applied to the 34 datasets; **b** pairwise comparison using the Wilcoxon signed-rank test, with black squares identifying non-significant comparisons



**Fig. 5** Earliness **a, b** and predictive performance **c, d** comparison of the ECONOMY approaches

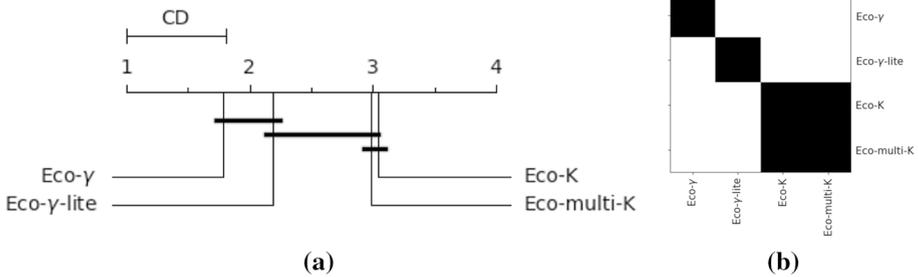
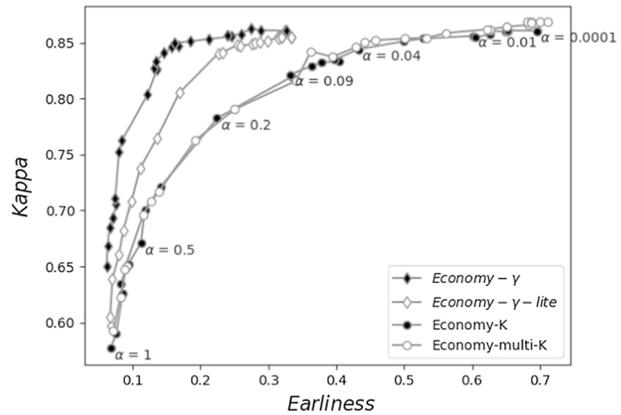
at  $\hat{t}^*$ , since this criterion properly manages unbalanced datasets (see Table 2). Again, the ECONOMY- $\gamma$  and ECONOMY- $\gamma$ -LITE dominate in terms of predictive performance, but here the difference is not statistically significant.

(d) *Pareto curves when varying the  $\alpha$  coefficient controlling the delay cost* In Fig. 6, the coordinates of each point are given by the average *Kappa* score and the average earliness obtained over the 34 datasets when the delay cost  $\alpha$  is chosen in the range  $[10^{-4}, 1]$ , and the Pareto curve is drawn for each of the approaches. The result is strikingly clear. For each value of  $\alpha$ , ECONOMY- $\gamma$  dominates all others approaches, even if ECONOMY- $\gamma$ -LITE is not far behind. The ECONOMY-K and ECONOMY-MULTI-K approaches yield much weaker results and are indistinguishable from each other.

(e) *Comparison with the best possible performance*

The approaches presented are able to adapt their decision time  $\hat{t}^*$  to the characteristics of the time series and to perform well in terms of average decision costs *AvgCost*, but to what

**Fig. 6** Average *Earliness* vs. Average *Kappa* score obtain over the 34 datasets by varying the slope of the time cost, such as  $\alpha \in [10^{-4}, 1]$



**Fig. 7** Evaluation of the quality of online decisions based on  $\Delta_{cost}$

extent these results differ from the optimal ones  $AvgCost^*$  computable *after* the entire time series is known? For each dataset,  $\Delta_{cost} = |AvgCost - AvgCost^*|$  was computed. Figure 7 shows that *ECONOMY- $\gamma$*  provides the best online decisions compared to the optimal ones, on average, followed by *ECONOMY- $\gamma$ -LITE*. According to the Wilcoxon signed-rank test, this difference is significant compared to the other *ECONOMY* approaches (Fig. 7b).

From all these results, several conclusions can be drawn.

1. The *ECONOMY* approaches that partition the time series using the learned classifiers (supervised-based methods) perform significantly better than those which exploit the K-means algorithm (unsupervised-based methods).
2. For the unsupervised approaches, partitioning the time series on full-length time series (*ECONOMY-K*), or on truncated ones (*ECONOMY-MULTI-K*) does not significantly affect the performances.
3. Regarding the supervised methods, using a more sophisticated anticipation mechanism of the incoming time series as done by *ECONOMY- $\gamma$*  is profitable and allows it to beat the less sophisticated *ECONOMY- $\gamma$ -LITE* method

## 7 Comparing ECONOMY- $\gamma$ and the state of the art

This section compares ECONOMY- $\gamma$  to the state of the art approach (Mori et al. 2017) and investigates the effects of the non-myopic property on its performance.

### 7.1 Comparing the performances

An important question is whether it is worth considering explicitly, in a single optimization criterion, earliness and accuracy, as in the ECONOMY approaches, and furthermore to adopt a non-myopic strategy with the modeling and computational costs involved. To assess this, we compared the ECONOMY methods with a competing algorithm, called SR presented in Mori et al. (2017) which is claimed to dominate all other algorithms over 45 benchmark datasets. The SR algorithm uses a trigger function to decide if the current prediction is reliable (output 1) or if it is preferable to wait for more data (output 0). Among several triggered functions, all of a heuristic nature, the most effective is:

$$Trigger(h_t(\mathbf{x}_t)) = \begin{cases} 0 & \text{if } \gamma_1 p_1 + \gamma_2 p_2 + \gamma_3 \frac{t}{T} \leq 0 \\ 1 & \text{otherwise} \end{cases} \quad (13)$$

where  $p_1$  is the largest posterior probability estimated by the classifier  $h_t$ ;  $p_1 = \text{ArgMax}_{y \in \mathcal{Y}}(\hat{p}(y|\mathbf{x}_t))$ ,  $p_2$  is the difference between the two largest posterior probabilities, defined as  $|\hat{p}(y = 1|\mathbf{x}_t) - \hat{p}(y = 0|\mathbf{x}_t)|$  in the case of binary classification problems, and where the last term  $\frac{t}{T}$  represents the proportion of the incoming time series that is visible at time  $t$ .

The parameters  $\gamma_1, \gamma_2, \gamma_3$  are real values in  $[-1, 1]$  to be optimized. In our experiments, these parameters were tuned using a grid-search over the set of values  $[-1, -0.95, -0.90, \dots, 0, 0.05, \dots, 0.90, 0.95, 1]$  in order to minimize the criterion *AvgCost*. The optimization was carried out for all possible time cost functions with a slope  $\alpha \in [10^{-4}, 1]$ .

After training, the *AvgCost* criterion was evaluated on the 45 test sets, and  $\alpha$  was adjusted for each dataset in order to find the most favorable setting for the SR algorithm, namely one maximizing  $\text{AvgCost}_{SR} - \text{AvgCost}_{Eco-\gamma}$ .

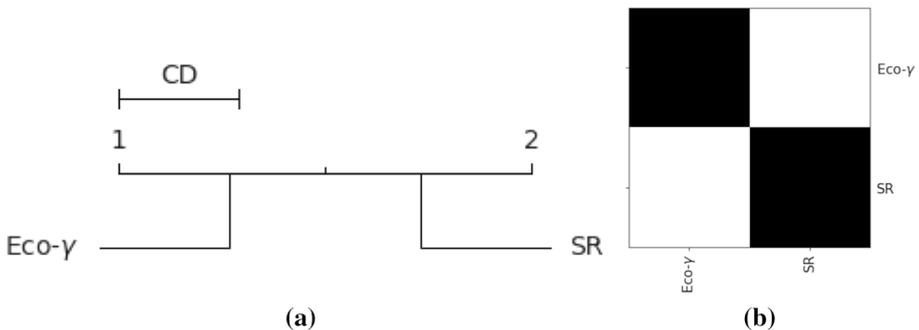
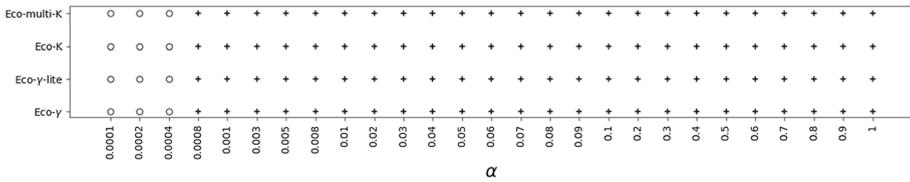
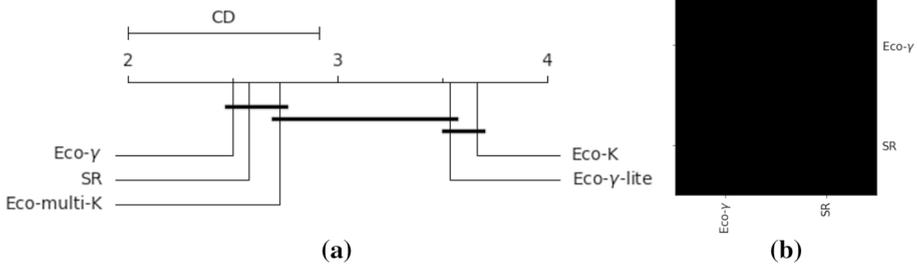


Fig. 8 SR vs. ECONOMY- $\gamma$ : evaluation based on *AvgCost*



**Fig. 9** Mori vs. ECONOMY approaches : evaluation based on AvgCost using the Wilcoxon signed-rank test, for different values of  $\alpha$ : “+” indicate success of ECONOMY approaches and “o” insignificant difference in performance



**Fig. 10** Evaluation of the myopic version of Economy approaches and SR approach based on AvgCost: **a** Nemenyi test applied to the 45 datasets; **b** pairwise comparison using the Wilcoxon signed-rank test, with black squares identifying non-significant comparisons

Figure 8a reports the results of the Nemenyi test and demonstrates that even in these situations favoring the SR algorithm, ECONOMY- $\gamma$  reaches significantly better performances. The Wilcoxon signed-rank test presented in Fig. 8b reinforces this conclusion.

We also carried out the Wilcoxon signed-rank test to compare the SR approach with the four ECONOMY approaches, for each value of  $\alpha \in [10^{-4}, 1]$ . The results (see Fig. 9) shows forcibly that the ECONOMY approaches perform significantly better than the SR approach, regardless of the value of  $\alpha$ ; except for  $\alpha \in \{10^{-4}, 2 \cdot 10^{-4}, 4 \cdot 10^{-4}\}$  where this difference is not significant.

**7.2 Measuring the effect of the non-myopic property of ECONOMY**

One important feature that differentiates the ECONOMY approaches from the state of the art methods is being non-myopic. Where standard methods decide whether to make a prediction at the current time based only on currently available information, the Economy algorithms look at future instants in order to predict the best decision time.

This section presents experiments aimed at answering the following question: Is it better to be non-myopic for an online decision system?

To answer this question, four myopic versions of the proposed ECONOMY approaches were implemented by limiting the horizon to only one measurement in the future, instead of looking at all the future time steps. The experiments performed with these myopic approaches are similar to those described in Sect. 7. The results are reported in Fig. 10.

Figure 10a reports the results of the Nemenyi test between the myopic ECONOMY approaches and the SR approach of Mori et al. (2017). They are cast in the same group.

Moreover, the Wilcoxon signed-rank test presented in Fig. 10b shows that there is no statistical significant difference in performance. Consequently, it appears that the non-myopic feature is a key property required to obtain better results.

## 8 Conclusions

An increasing number of applications require the ability to recognize the class of an incoming time series as quickly as possible without unduly compromising the accuracy of the prediction. In this paper, we reformulated in a generic way an optimization criterion put forward in Dachraoui et al. (2015) which takes into account both the cost of misclassification and the cost of delaying the decision.

This generic framework has been technically declined, leading to the design of three new “non-myopic” algorithms—i.e. able to anticipate the expected future gain in information in balance with the cost of waiting. In one class of algorithms, unsupervised-based, the expectations use the clustering of time series, while in a second class, supervised-based, time series are grouped according to the confidence level of the classifier used to label them.

We have defined a new evaluation criterion that represents the average cost incurred when the method is applied over a set of labelled time series. This criterion makes it possible to evaluate both earliness and predictive performance as a single objective, with respect to the ground truth. It offers a well-grounded framework widely applicable for the comparison of methods.

Extensive experiments carried out on real datasets using a large range of delay cost functions show that the presented algorithms are able to satisfactorily solve the earliness vs. accuracy trade-off, with the supervised partition based approaches generally faring better than the unsupervised partition based ones. In addition, all these methods perform better in a wide variety of conditions than the state of the art competitive method of Mori et al. (2017). The non-myopic feature of the ECONOMY approaches is required for this good achievement. We have shown that the non-myopic property of the ECONOMY approaches plays a key role for these good performances.

Given the merit of the proposed approach, we envision several extensions. One is to allow the ECONOMY- $\gamma$ -LITE and ECONOMY- $\gamma$  approaches, which use the confidence level of a binary classifier, to solve multi-classes problems. A second one is to use a supervised clustering technique to compute groups of times series (see Lemaire et al. 2020) in the ECONOMY-K and ECONOMY-MULTI-K approaches. Finally, we are working on the adaptation of these methods to the on-line detection of anomalies in a data stream.

**Acknowledgements** We thank Vincent Lemaire and Fabrice Clérot (Orange Labs) for their advice and interesting discussions. We also thank Orange Labs for supporting this research.

## References

- Alonso González, C. J., & Díez, J. J. R. (2004). Boosting interval-based literals: Variable length and early classification. In *Data mining in time series databases*, World Scientific (pp. 149–171).
- Anderson, H. S., Parrish, N., Tsukida, K., & Gupta, M. (2012). Early time-series classification with reliability guarantee. Sandria Report.
- Barandas, M., Folgado, D., Fernandes, L., Santos, S., Abreu, M., Bota, P., Liu, H., Schultz, T., & Gamboa, H. (2020). Tsfel: Time series feature extraction library. *SoftwareX* 11:100456, <https://github.com/fraunhoferportugal/tsfel>.

- Berger, J. O. (1985). *Statistical decision theory and Bayesian analysis*. Springer.
- Cohen, J. (1960). A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1), 37–46.
- Dachraoui, A., Bondu, A., & Cornuéjols, A. (2013). Early classification of individual electricity consumptions. In *RealStream2013 (ECML)* (pp. 18–21).
- Dachraoui, A., Bondu, A., & Cornuéjols, A. (2015). Early classification of time series as a non myopic sequential decision making problem. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases* (pp. 433–447). Springer.
- DeGroot, M. H. (2005). *Optimal statistical decisions*, (Vol. 82). Wiley.
- Flach, P. A. (2016). Classifier calibration. *Encyclopedia of Machine Learning and Data Mining* (pp. 1–8).
- Ghalwash, M. F., Ramljak, D., & Obradović, Z. (2012). Early classification of multivariate time series using a hybrid hmm/svm model. In *2012 IEEE International Conference on Bioinformatics and Biomedicine* (pp. 1–6). IEEE, .
- Ghalwash, M. F., Radosavljevic, V., & Obradovic, Z. (2014). Utilizing temporal patterns for estimating uncertainty in interpretable early decision making. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining, ACM*, (pp. 402–411).
- Gupta, A., Gupta, H. P., Biswas, B., & Dutta, T. (2020). An early classification approach for multivariate time series of on-vehicle sensors in transportation. *IEEE Transactions on Intelligent Transportation Systems*.
- Hatami, N., & Chira, C. (2013). (2013) Classifiers with a reject option for early time-series classification. In *IEEE symposium on computational intelligence and ensemble learning (CIEL)* (pp. 9–16). IEEE.
- He, G., Duan, Y., Peng, R., Jing, X., Qian, T., & Wang, L. (2015). Early classification on multivariate time series. *Neurocomputing*, 149, 777–787.
- Lemaire, V., Alaoui Ismaili, O., Cornuéjols, A., & Gay, D. (2020). Predictive k-means with local models. In *Workshop LDRC-2020 (Workshop on Learning Data Representation for Clustering) in PAKDD-2020 (The 24th Pacific-Asia Conference On Knowledge Discovery and Data Mining), Singapore*, 11-16 May 2020.
- Liu, Y., & Li, X. R. (2013). Performance analysis of sequential probability ratio test. *Sequential Analysis*, 32(4), 469–497.
- Mathukia, C., Fan, W., Vadyak, K., Biege, C., & Krishnamurthy, M. (2015). Modified early warning system improves patient safety and clinical outcomes in an academic community hospital. *Journal of Community Hospital Internal Medicine Perspectives*, 5(2), 26716.
- Mori, U., Mendiburu, A., Dasgupta, S., & Lozano, J. (2015). Early classification of time series from a cost minimization point of view. In *Proceedings of the NIPS Time Series Workshop*.
- Mori, U., Mendiburu, A., Dasgupta, S., & Lozano, J. A. (2017). Early classification of time series by simultaneously optimizing the accuracy and earliness. *IEEE Transactions on Neural Networks and Learning Systems*, 29(10), 4569–4578.
- Mori, U., Mendiburu, A., Miranda, I. M., & Lozano, J. A. (2019). Early classification of time series using multi-objective optimization techniques. *Information Sciences*, 492, 204–218.
- Nemenyi, P. (1962). Distribution-free multiple comparisons. *Biometrics*, 18(2), 263.
- Novikov, A. (2008). Optimal sequential tests for two simple hypotheses based on independent observations. *International Journal of Pure and Applied Mathematics*, 45(2), 291–314.
- Parrish, N., Anderson, H. S., Gupta, M. R., & Hsiao, D. Y. (2013). Classifying with confidence from incomplete information. *The Journal of Machine Learning Research*, 14(1), 3561–3589.
- Rußwurm, M., Lefevre, S., Courty, N., Emonet, R., Körner, M., & Tavenard, R. (2019). End-to-end learning for early classification of time series. arXiv preprint [arXiv:1901.10681](https://arxiv.org/abs/1901.10681).
- Schäfer, P., & Leser, U. (2020). Teaser: Early and accurate time series classification. *Data Mining and Knowledge Discovery*, 34(5), 1336–1362.
- Vapnik, V., & Vashist, A. (2009). A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5–6), 544–557.
- Wald, A., & Wolfowitz, J. (1948). Optimum character of the sequential probability ratio test. *The Annals of Mathematical Statistics* (pp. 326–339).
- Xing, Z., Pei, J., & Philip, S. Y. (2009). Early prediction on time series: A nearest neighbor approach. In *IJCAI, Citeseer* (pp. 1297–1302).
- Xing, Z., Pei, J., Philip, S. Y., & Wang, K. (2011). Extracting interpretable features for early classification on time series. *SDM, SIAM*, 11, 247–258.