



HAL
open science

Early and Revocable Time Series Classification

Youssef Achenchabe, Alexis Bondu, Antoine Cornuéjols, Vincent Lemaire

► **To cite this version:**

Youssef Achenchabe, Alexis Bondu, Antoine Cornuéjols, Vincent Lemaire. Early and Revocable Time Series Classification. International Joint Conference on Neural Networks (IJCNN), Jul 2022, Padua, Italy. 10.1109/IJCNN55064.2022.9892391 . hal-03798026

HAL Id: hal-03798026

<https://hal.science/hal-03798026>

Submitted on 5 Oct 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Early and Revocable Time Series Classification

Youssef Achenchabe
Orange Labs
Paris-Saclay University
Paris, France

youssef.achenchabe@universite-paris-saclay.fr

Alexis Bondu
Orange Labs
Châtillon, France
alexis.bondu@orange.com

Antoine Cornuéjols
Paris-Saclay University
AgroParisTech
Paris, France
antoine.cornuejols@agroparistech.fr

Vincent Lemaire
Orange Labs
Lannion, France
vincent.lemaire@orange.com

Abstract—Many approaches have been proposed for early classification of time series in light of its significance in a wide range of applications including healthcare, transportation and finance. Until now, the early classification problem has been dealt with by considering only irrevocable decisions. This paper introduces a new problem called *early and revocable* time series classification, where the decision maker can revoke its earlier decisions based on the new available measurements. In order to formalize and tackle this problem, we propose a new cost-based framework and derive two new approaches from it. The first approach does not consider explicitly the cost of changing decision, while the second one does. Extensive experiments are conducted to evaluate these approaches on a large benchmark of real datasets. The empirical results obtained convincingly show (i) that the ability of revoking decisions significantly improves performance over the irrevocable regime, and (ii) that taking into account the cost of changing decision brings even better results in general.

I. INTRODUCTION

Consider Eisenhower, in June 1944, having to decide when to launch the landing on the French coast [10]. He had an imperfect knowledge of the weather conditions. The longer he waited, the more precise they became, allowing for a more informed decision: to launch the landing today or wait for another day, but the more difficult it became to ensure that all arrangements would be met and that the enemy remained unaware of the danger. Eisenhower was faced with a very common problem, even if dramatic here, to have to optimize a trade-off between the earliness of a decision and its potential cost. Note that once the decision to launch operation Overlord was made, it was *irrevocable*. There was no way it could be halted.

In many situations, however, one can take a decision and then decide to *change* it after some new pieces of

information become available. The change may be costly but still warranted because it seems likely to lead to a much better outcome. This can be the case for instance when an outdoor event is canceled due to a dramatic change in the weather forecast, or when a doctor revises what now seems a misdiagnosis.

The problem now is to identify the optimal sequences of decisions given an incoming series of measurements, and the various existing costs defined above.

Notations

More formally, we assume that there exists a data set $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1 \leq i \leq m}$ of *complete* time series $\mathbf{x}_T = \langle x_1, \dots, x_T \rangle$ each of which is associated with a label $y \in \mathcal{Y}$ (e.g. patient who needs a surgical operation or patient who does not). The measurements x_i ($1 \leq i \leq T$) belong to some input space \mathcal{X} and can be univariate as well as multivariate. At each time step t , the decision-maker gets to know the time series measured so far: $\mathbf{x}_t = \langle x_1, \dots, x_t \rangle$ and must decide either to make a prediction \hat{y}_t about the class of the incoming time series or to postpone the decision.

In the *irrevocable regime*, once a decision has been taken, it cannot be changed and the decision-maker endures a cost which is the sum of the misclassification cost $C_m(\hat{y}_t|y)$ plus the cost of having delayed the decision until time t : $C_d(t)$. Whereas, in the *revocable regime*, the decision-maker can change its prediction several times before the time limit T . Let us call \mathcal{D}_ℓ , the sequence of the ℓ successive predictions $\langle \hat{y}_{t_1}, \dots, \hat{y}_{t_\ell} \rangle$ made at times t_1, \dots, t_ℓ in the time interval $\{1, \dots, T\}$. It is assumed that each decision change from \hat{y}_{t_i} to $\hat{y}_{t_{i+1}}$ entails a cost $C_{cd}(\hat{y}_{t_{i+1}}|\hat{y}_{t_i})$ that is greater or equal to 0.

Contributions of this paper

While the early classification of time series, in the *irrevocable* regime, has been addressed in several papers in the last few years, we do not know of similar works for the *revocable* regime. Yet, our work on a large set of datasets and a wide spectrum of misclassification, delay and revision costs, shows that intelligently identifying revocations instants can yield significant gains. Indeed, even with state of the art early classification methods, there exists situations where additional knowledge of the unfolding time series warrants to change decisions, even in face of increasing delay cost and additional cost of decision changes. We thus found that, depending on these relative costs, between 3% to 8% of the times series would benefit if changes of decision be made. Furthermore, a revocable strategy which takes into account the cost of changing decision almost always beats a naive yet *non-myopic*¹ revocable strategy that changes decision without considering the decision change cost.

The impact such an intelligent revocable strategy could have on prediction maintenance, intensive care units, autonomous cars and many more application domains where decisions have to be made optimizing costs of mistaken decisions and delay costs is quite significant.

The contribution of this paper is threefold. *First*, it formalizes the optimization problem associated with the revocable regime for the early classification problem (see challenge #9 raised in [5]). *Second*, it proposes two approaches to tackle this problem and we introduced an extended notion of *non-myopia*. Both approaches are *non-myopic* in that, to make their decisions, they take into account expectancies of the cost likely to incur in the foreseeable future:

- 1) The first approach is *conventionally* non-myopic, in the sense that it is only aware of the delay and misclassification costs: it is ready to revoke decision as soon as this seems reasonable, without considering the cost of changing decision.
- 2) The second approach is non-myopic of *second order*, as it estimates the future expected cost of a decision by taking into account the risk of revocation itself, which is not trivial (see challenge #10 raised in [5]). Specifically, a decision that will probably be revoked afterward should be delayed due to this risk. Conversely, a decision which promises to be sustainable should be anticipated.

Third, extensive experiments are presented that show that it is actually better to be able to revise decisions than

to implement an irrevocable decision strategy and, in addition that it is worth considering the non-myopic of *second order* approach.

For the clarity of this paper, we consider a simple case where the input is in the form of a univariate time series whose measurements are observed over time (i.e., equivalent to a single sensor). But the framework and approaches presented in this paper can be adapted to multivariate time series directly. It all depends on having a classifier that is able to use multivariate series as input.

This paper is organized as follows. Section II provides an overview of classical *early classification* approaches, all of which deal with the irrevocable regime. Section III focuses on a non-myopic framework which is designed for the irrevocable regime. The *early and revocable classification* problem is defined in Section IV. Then, two new approaches are proposed, which are evaluated through extensive experiments in Section V. Perspectives and future work are discussed in Section VI.

II. STATE OF THE ART ON EARLY CLASSIFICATION OF 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 with heuristic means to assess the confidence of the prediction at any one time [2], [11], [12], [14], [18]. It is noticeable that, in these works, 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.

In [16], the authors recognize the conflict between earliness and accuracy, and instead of setting a tradeoff in a single objective optimization criterion [15], they propose to keep it as a multi-objective criterion and to explore the Pareto front of the multiple dominating trade-offs. Accordingly, they propose a family of triggering functions involving hyperparameters to be optimized for each tradeoff. 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 [19] which compares several techniques for early classification of time series.

¹ A non-myopic approach predicts the best future decision time.

In [8], 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 to estimate the expected costs for an incoming subsequence \mathbf{x}_t at all future time steps allowing non-myopic decisions.

It is apparent that approaches that do not explicitly consider costs are ill-equipped to deal with the possibility of *revocable* decisions. At best, they could base such revisions on observing that the confidence level falls below the pre-set threshold, and possibly exceeds it again, but this would not allow for the associated costs: of decision change and of delay.

III. A COST-BASED NON-MYOPIC FRAMEWORK

In this section, we introduce a cost-based non-myopic framework that was designed for the irrevocable regime [1]. The purpose of the following sections is to show how it can be adapted to the revocable regime. Notice that this framework leads to the best performance observed to date, as empirically demonstrated in [1].

We suppose that a training set $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1 \leq i \leq m}$ of complete time series, each with its associated labels, exists.

I- For each time step, $t \in \{1, \dots, T\}$, and using the training set, a classifier h_t can be learned $h_t : \mathcal{X}^t \rightarrow \mathcal{Y}$.

II- Using these classifiers and the knowledge that can be extracted from \mathcal{S} when estimating the likely future of an incoming time series \mathbf{x}_t , it is possible to estimate the optimal instant for deciding a prediction about its class.

More precisely, given the *misclassification cost* function $C_m(\hat{y}|y) : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ and the *delay cost* function $C_d(t) : \mathbb{R} \rightarrow \mathbb{R}$, the expectancy of the cost of taking a decision at time t given the incoming time series \mathbf{x}_t is:

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

where $\mathbb{E}_{(\hat{y}, y) \in \mathcal{Y}^2}^t [\]$ is the expectancy at time t , over the variables y and \hat{y} . $P_t(y|\mathbf{x}_t)$ is the probability of the class y given a time series that starts as \mathbf{x}_t , and $P_t(\hat{y}|y, \mathbf{x}_t)$ is the probability that the classifier h_t makes the prediction \hat{y} given \mathbf{x}_t as input and when y would be its true label. In this non-myopic setting, the idea is that the decision of making a prediction is made at the current time t only

insofar that it is not expected that a lower cost could be achieved at a later time. This could happen if the expected misclassification cost would drop sufficiently to offset the increase of $C_d(t)$.

For any time in the future $t + \tau$ ($1 \leq \tau \leq T - t$), the expected cost of making a prediction can be estimated as:

$$\begin{aligned} f_\tau(\mathbf{x}_t) &= \mathbb{E}_{(\hat{y}, y) \in \mathcal{Y}^2}^{t+\tau} [C_m(\hat{y}|y)] + C_d(t + \tau) \\ &= \sum_{y \in \mathcal{Y}} P_{t+\tau}(y|\mathbf{x}_t) \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|y, \mathbf{x}_t) C_m(\hat{y}|y) \\ &\quad + C_d(t + \tau) \end{aligned} \quad (2)$$

This equation naively expands the expectation $\mathbb{E}_{(\hat{y}, y) \in \mathcal{Y}^2}^{t+\tau} [C_m(\hat{y}|y)]$ and implies a sum over the possible values of the true label y and another sum over the predicted label \hat{y} . Then the optimal decision time, at time t , is expected to be:

$$\tau^* = \underset{\tau \in \{0, \dots, T-t\}}{\text{ArgMin}} f_\tau(\mathbf{x}_t) \quad (3)$$

As the reader may notice Equation 2 is not tractable since the true label y is not available when a new input time series is observed. The authors in [1] propose to partition the training set in order to identify groups \mathcal{G} of time series that share similar patterns³, or share the same confidence levels with respect to the output of the classifiers⁴. This makes it possible to estimate the terms $P_t(y|\mathbf{x}_t)$ $P_t(\hat{y}|y, \mathbf{x}_t)$ of Equation 2 by replacing them with $P_t(y|\mathbf{g}_k)$ and $P_{t+\tau}(\hat{y}|y, \mathbf{g}_k)$ which can be easily estimated. Thus, Equation 2 becomes:

$$\begin{aligned} f_\tau(\mathbf{x}_t) &= \left(\sum_{\mathbf{g}_k \in \mathcal{G}} P_t(\mathbf{g}_k|\mathbf{x}_t) \sum_{y \in \mathcal{Y}} P_t(y|\mathbf{g}_k) \right. \\ &\quad \left. \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|y, \mathbf{g}_k) C_m(\hat{y}|y) \right) + C_d(t + \tau) \end{aligned} \quad (4)$$

The idea is to estimate the cost of a decision at all future time steps, up until $t = T$, based on the current knowledge about the incoming time series, and to postpone the decision to the time step that appears to be the best.

If $\tau^* = 0$, then the best time for prediction seems to be now, the prediction $h_t(\mathbf{x}_t)$ is returned and the classification process is terminated. Otherwise the decision

² Notice that $f_0(\mathbf{x}_t) = \mathbb{E}_{y \in \mathcal{Y}}^t [C_m(\hat{y}|y)] + C_d(t)$ since we have access to predictions at current time.

³ Approach called Economy_K in [1].

⁴ Approach called Economy_γ in [1].

is postponed to the next time step, and Equation 3 is computed again, this time with \mathbf{x}_{t+1} . The process goes on until a decision is made or $t = T$ at which point a prediction is forced.

The question then arises as to how best adapt the *irrevocable* non-myopic strategy just described to the *revocable* regime where changes of decisions are allowed until T , but at the expense of incurring the additional costs $C_{cd}(\cdot)$ associated with these changes.

IV. A NEW FRAMEWORK FOR REVOCABLE DECISIONS

Suppose that while the measurements x_t about time series \mathbf{x}_T unfold from time $t = 1$ to $t = T$, the decision-maker can change its mind as many times as it sees fit and ends up triggering a sequence of predictions $\mathcal{D}_\ell = \langle \hat{y}_{t_1}, \dots, \hat{y}_{t_\ell} \rangle$ about the class of the input time series. The *final* cost incurred will be:

$$g(\mathcal{D}_\ell | \mathbf{x}_T, y) = C_m(\hat{y}_{t_\ell} | y) + C_d(t_\ell) + \sum_{\substack{i=1 \\ \hat{y}_{t_i}, \hat{y}_{t_{i+1}} \in \mathcal{D}_\ell}}^{\ell-1} C_{cd}(\hat{y}_{t_{i+1}} | \hat{y}_{t_i}) \quad (5)$$

where t_ℓ is the time of the last change of decision yielding the prediction $\hat{y}_{t_\ell} = h_{t_\ell}(\mathbf{x}_{t_\ell})$.

Formally, the problem is now to find a sequence of decisions $\mathcal{D}^* \in \mathbb{D}_T$ that minimizes Equation 5:

$$\mathcal{D}^* = \underset{\mathcal{D} \in \mathbb{D}_T}{\text{ArgMin}} g(\mathcal{D} | \mathbf{x}_T, y) \quad (6)$$

where \mathbb{D}_T is the set of all possible sequences of maximum length T .

Non-myopia of second order: When, at time t , only a partial knowledge \mathbf{x}_t is available about the incoming time series, Equation 5 cannot be computed. A sequence of decisions $\mathcal{D}_k = \langle \hat{y}_{t_1}, \dots, \hat{y}_{t_k} \rangle$ has been taken so far, and the question is to see if changing the last decision \hat{y}_{t_k} now, at time t , is favorable, because it would bring a better expected cost, and it would not seem better to postpone such a possible change to a later time $t + \tau$. Note that this is where *second order* considerations enter the optimization problem. In order to decide if now is a good time to change decision, one has to look if another change of decision is likely to happen in the future, at any time $t + \tau$ (see the term $P_{t+\tau}(\hat{y} | \hat{y}_{t_k}, \mathbf{x}_t)$ of Equation 8).

The cost of adding a new decision at time $t + \tau$, can be estimated as:

$$f_\tau^{\text{rev}}(\mathcal{D}_k, t + \tau | \mathbf{x}_t) = \mathbb{E}_{(\hat{y}, y) \in \mathcal{Y}^2}^{t+\tau} [C_m(\hat{y} | y) | \mathbf{x}_t] + \underbrace{\sum_{i=1}^{k-1} C_{cd}(\hat{y}_{t_{i+1}} | \hat{y}_{t_i})}_{\text{known values of past changes}} + \underbrace{\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y} | \hat{y}_{t_k}) | \mathbf{x}_t]}_{\text{expected value at } t + \tau} + C_d(t + \tau) \quad (7)$$

The expected cost of changing decision is defined as follows for $(1 \leq \tau \leq T - t)$:

$$\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y} | \hat{y}_{t_k}) | \mathbf{x}_t] = \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y} | \hat{y}_{t_k}, \mathbf{x}_t) C_{cd}(\hat{y} | \hat{y}_{t_k}) \quad (8)$$

Given that the notation \mathcal{D}_{k+1} is used to denote the sequence of decisions $\langle \hat{y}_{t_1}, \dots, \hat{y}_{t_k}, \hat{y}_t \rangle$, with $\hat{y}_t = h(\mathbf{x}_t)$, the criterion *cd* for changing decision at time t becomes:

$$cd = \begin{cases} \hat{y}_t \neq \hat{y}_{t_k} \\ \text{and } \underset{\tau \in \{0, \dots, T-t\}}{\text{ArgMin}} f_\tau^{\text{rev}}(\mathcal{D}_k, t + \tau | \mathbf{x}_t) = 0 \\ \text{and } f_{\tau=0}^{\text{rev}}(\mathcal{D}_{k+1}, t | \mathbf{x}_t) < f_{\tau=0}^{\text{rev}}(\mathcal{D}_k, t_k | \mathbf{x}_t) \end{cases} \quad (9)$$

A decision is thus taken at time t only if (i) the current prediction \hat{y}_t would differ from the last one \hat{y}_{t_k} , (ii) if it seems that now is the best time to make a new decision, and (iii) if the estimated cost with the new prediction would be less than the engaged one with the previous decision.

An interesting case occurs *when changing decision is costless*: $\forall y, y' \in \mathcal{Y} \times \mathcal{Y}, C_{cd}(y | y') = 0$. Equation 7 becomes:

$$f_\tau^{\text{rev}}(\mathcal{D}_k, \tilde{t} | \mathbf{x}_t) = \mathbb{E}_{(\hat{y}, y) \in \mathcal{Y}^2}^{t+\tau} [C_m(\hat{y} | y) | \mathbf{x}_t] + C_d(\tilde{t}) \quad (10)$$

which is Equation 2. Then, the strategy is to change decision when the gain in the expected misclassification cost with a new decision offsets the increased delay cost.

Now a question is: what would be the **optimal sequence of decisions**?

Theorem IV.1 (Optimal sequence of decisions). *Let us assume that $\forall (y, y') \in \mathcal{Y}^2, C_{cd}(y | y') > 0$. Then, for any time series \mathbf{x}_T of class y , the optimal sequence of decision is reduced to a **one decision** sequence where the optimal time⁵ t^* is defined by: $t^* = \underset{1 \leq t \leq T}{\text{ArgMin}} \{C_m(\hat{y}_t | y) + C_d(t)\}$.*

⁵Actually several optimum may occur at different times, and then any one of them can be chosen.

Proof. Let $\mathcal{D}_k = \langle \hat{y}_{t_1}, \dots, \hat{y}_{t_k} \rangle$ be a sequence of decisions taken at times $\{t_1, \dots, t_k\}$. Then the cost paid at time T is: $\sum_{i=1}^{k-1} C_{cd}(\hat{y}_{i+1}|\hat{y}_i) + C_m(\hat{y}_{t_k}|y) + C_d(t_k)$ which cannot be less than: $C_m(\hat{y}_{t^*}|y) + C_d(t^*)$. \square

Theorem IV.1 shows that it is better to make the optimal decision at the right time rather than revoking a decision since this can only lead to sub-optimal sequences of decisions. However, in practice, the ground truth y is unknown, and it may be unavoidable to make a first decision, because it seems the optimal time to do so, only to find later that it should be changed.

It must be noted that the *criterion* (Eq. 9) does not specify how and when to make **the first prediction** \hat{y}_{t_1} . Since a decision is mandatory in the framework of decision making, we assume that a “no decision” is associated with an infinite cost: $f_\tau^{\text{rev}}(\emptyset | \mathbf{x}_t) = +\infty$, forcing a decision before T , according to the non-myopic strategy defined by $f_\tau(\mathbf{x}_t)$ in its irrevocable regime (see Equation 1).

One goal of our research is to evaluate the added value of explicitly taking into account the cost of the changes of decision with respect to a revocable strategy which would not. Accordingly, we implemented two algorithms, based on the ECONOMY- γ algorithm [1].

- 1) The first one is named ECO-REV-CU for cost unaware (as in Equation 10).
- 2) The second is named ECO-REV-CA for cost aware (as in Equation 7).

A generic algorithmic implementation of the revocable decision-making criterion as defined in Equation 9 is presented in Algorithm 1.

Complexity Analysis

We present here the time complexity of the two proposed algorithms. First, let us define some notations:

- *Learn(m)*: time complexity for learning a single classifier;
- *Predict*: time complexity of inference phase of a classifier on a time series;
- *Partitioning*: time complexity for partitioning a set of time series;
- *K*: number of groups in the data partition;
- *m*: number of time series within the dataset.

The *training stage* consists of multiple steps: (i) learning a classifier for each timestamp, in a $\mathcal{O}(T \cdot \text{Learn}(m))$ complexity; (ii) partitioning the training set, in $\mathcal{O}(\text{Partitioning})$; (iii) computing predictions for all examples in the training set at each timestamp in order to compute confusion matrices, $\mathcal{O}(T \cdot m \cdot \text{Predict})$;

Algorithm 1: GENERIC REVOCABLE REGIME ALGORITHM

Require: K : number of groups

- 1: $decisions \leftarrow \langle \hat{y}_{t_1} \rangle$
 - 2: $t_{prev} \leftarrow t_1$
 - 3: **for all** $t = t_1 + 1 \dots T$ **do**
 - 4: $\tau^* \leftarrow \text{ArgMin}_{\tau \in \{0 \dots T-t\}} f_\tau^{\text{rev}}(decisions, t + \tau | \mathbf{x}_t)$
 - 5: $cost_{new} \leftarrow f_{\tau=0}^{\text{rev}}(decisions, t + \tau^* | \mathbf{x}_t)$
 - 6: $cost_{prev} \leftarrow f_{\tau=0}^{\text{rev}}(decisions, t_{prev} | \mathbf{x}_t)$
 - 7: **if** $\hat{y}_t \neq \hat{y}_{t_{prev}}$ **and** $\tau^* = 0$ **and**
 $cost_{new} < cost_{prev}$ **then**
 - 8: $t_{prev} = t$
 - 9: $decisions \leftarrow decisions \cup \hat{y}_t$
 - 10: **end if**
 - 11: **end for**
 - 12: **return** $decisions$
-

(iv) the prior of each class in each group must be computed in $\mathcal{O}(|Y| \cdot K \cdot m)$; (v) The expected cost of decision changes is computed for all future timestamps at each timestamp (T^2 in complexity) according to Eq.7 which results in a complexity of $\mathcal{O}(|Y|^2 \cdot T^2 \cdot K)$. Finally, the overall time complexity of ECO-REV-CU is $\mathcal{O}(T \cdot \text{Learn}(m) + \text{Partitioning} + T \cdot m \cdot \text{Predict} + |Y| \cdot K \cdot m)$ and ECO-REV-CA is $\mathcal{O}(T \cdot \text{Learn}(m) + \text{Partitioning} + T \cdot m \cdot \text{Predict} + |Y| \cdot K \cdot m + |Y|^2 \cdot T^2 \cdot K)$.

For the *testing part*, the time complexity of estimating the cost expectancy of future time step is similar to the *irrevocable* regime which is $\mathcal{O}(T^2 \cdot |Y|^2 \cdot K)$ as presented in [1]. Taking into account the final decision and the intermediate predictions of the classifiers, this complexity becomes $\mathcal{O}(\text{Predict} \cdot T^2 \cdot |Y|^2 \cdot K)$. \square

V. EXPERIMENTS

The experiments aim at measuring the true added value of a revocable strategy. Specifically, the question is twofold. *First*, does such a strategy recognize useful changes of decisions: those that increase the performance? *Second*, does it pay off to implement a revocable strategy that takes into account the costs of changing decisions by comparison to a naive one that would not consider these costs? In the following, we report results obtained on 34 datasets (see Section V-B) for a whole range of values for the delay cost C_d and the cost incurred if changing decision C_{cd} .

A. Implementation choices

In our experiments, the ECO-REV-CU algorithm is simply the ECONOMY- γ algorithm allowed to be reit-

erated after each decision. It thus does not take into account the costs associated with changing decisions, whereas ECO-REV-CA does. More technically, ECO-REV-CA approximates $\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y}|\hat{y}_{t_k}) | \mathbf{x}_t]$ in Equation 8 by using the groups of time series, denoted by \mathcal{G} :

$$\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y}|\hat{y}_{t_k}) | \mathbf{x}_t] \approx \sum_{\mathbf{g}_k \in \mathcal{G}} P(\mathbf{g}_k | \mathbf{x}_t) \mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y}|\hat{y}_{t_k}) | \mathbf{g}_k] \quad (11)$$

Then, the probability $P_{t+\tau}(\hat{y}|\hat{y}_{t_k}, \mathbf{g}_k)$ entering in the term $\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [C_{cd}(\hat{y}|\hat{y}_{t_k}) | \mathbf{g}_k]$ is estimated in a frequentist way as the proportion of time series predicted to belong to \hat{y}_{t_k} at time t_k , and for which the classifier changed its decision at time $t + \tau$ by predicting the class \hat{y} . For full reproducibility of the experiments presented in this paper, an open-source code is available in <https://github.com/YoussefAch/rev-economy>.

B. Data and feature extraction

Because ECONOMY- γ is restricted to binary classification problems, and in order to be able to directly compare our results with those reported in [1], we chose to use the same 34 datasets that are taken from the UEA & UCR Time Series Classification Repository⁶ [3]. It is important to note that the revocable framework presented here could as well accommodate multi-class classification problems. Additional experiments on multi-class problems are reported in [9].

Each training set is built with 70% of the examples randomly uniformly selected, while the remaining 30% are used as test set (note that in each dataset, all time series have the same length). In addition, each training set is divided into three disjoint subsets: (i) 40% for training the Xgboost [6] classifiers $\{h_t\}_{t \in \{1, \dots, T\}}$ that are the base classifiers used in the ECONOMY- γ method, which offer a good trade-off between computing time and accuracy; (ii) 40% for estimating the probabilities in f_{τ}^{rev} and f_{τ} ; and (iii) the remaining 20% for optimizing the number of groups $|\mathcal{G}|$ in ECONOMY- γ which is its only hyper-parameter.

In order to give equal weight to all data sets in the comparison, it is important that they offer the same number of opportunities for decision changes. This is why the instants for potential changes are sampled every $n\%$ of the length of the times series in each data set (in our case, $n = 5\%$). For each possible length, 60 features on the statistical, temporal and spectral domains are extracted using the Time Series Feature Extraction Library [4], and are used for training the classifiers $\{h_t\}_{t \in \{1, \dots, T\}}$.

⁶Available at: <http://www.timeseriesclassification.com>

C. The evaluation criterion

The cost incurred using an early classification system on a time series \mathbf{x}_T is the sum of three costs, the cost of misclassification, the delay cost incurred at the time of the last decision, and the sum of the costs associated with all changes of decision if any:

$$\text{Cost}(\mathbf{x}_T) = C_m(h_{t_l}(\mathbf{x}_{t_l})|y) + C_d(t_l) + \sum_{i=1}^{|\mathcal{D}_l|-1} C_{cd}(\hat{y}_{i+1}|\hat{y}_i) \quad (12)$$

In order to evaluate a method, we compute its mean performance on the test set \mathcal{T} :

$$\text{AvgCost}(\mathcal{T}) = \frac{1}{|\mathcal{T}|} \sum_{i=1}^{|\mathcal{T}|} \text{Cost}(\mathbf{x}_T^i) \quad (13)$$

D. Description of the experiments

In our experiments, we compared three algorithms: ECONOMY- γ which is an irrevocable decision-maker, ECO-REV-CU which is the revocable version of ECONOMY- γ but unaware of the costs of changing decision, and ECO-REV-CA which is aware of these changing costs. We are thus able to measure the added-value of the revocable strategy (ECO-REV-CU vs. ECONOMY- γ) and the added-value of being aware of the costs of changing decision (ECO-REV-CA vs. ECO-REV-CU).

For a given application, the various costs, relative to misclassifications, delays and changes of decision, must be provided by the domain expert⁷. For our experiments, we explored the performance of the three methods on a wide range of cost values:

- The *misclassification cost* was set to $C_m(\hat{y}|y) = 1$ if $\hat{y} \neq y$, and $= 0$ if not.
- The *delay cost* was assumed to be linear with a positive slope: $C_d = \alpha \times \frac{t}{T}$ starting from very low $\alpha = \{0.0001, 0.00025, 0.0005, 0.00075\}$, to low $\alpha = \{0.001, 0.0025, 0.005, 0.0075\}$, to medium values $\alpha = \{0.01, 0.025, 0.05, 0.075\}$ and to high values $\alpha = \{0.1, 0.25, 0.5, 0.75, 1\}$.
- The *cost of changing decision* was set to $C_{cd}(\hat{y}_1|\hat{y}_2) = \beta$ if $\hat{y}_1 \neq \hat{y}_2$, and $= 0$ otherwise. The parameter β being taken in the same set of values as α ⁸.

The AvgCost criterion defined in Equation 13 was evaluated on the 34 test sets for all cost values, and the Wilcoxon signed-rank test [20] was performed for all the

⁷ For instance, for condition of septic shock, every hour delay in antibiotic treatment leads to 8% increase in the risk of mortality [13]

⁸ α and β were chosen in a very large spectrum of values so as not biasing the results.

range of cost values, in order to assess whether the observed performance gap between methods is significant (“+” and “-”) or not (“o”).

E. Results and analysis

Before comparing the methods, it is important to measure the proportion of time series that offer useful opportunities for revocable decisions. Those are the ones where the first decision taken by an irrevocable strategy, here ECONOMY- γ , turns out not to be optimal. For the 34 datasets under study, it turns out that (i) for a low delay cost $C_d = 0.0025 \times \frac{t}{T}$ only 3% of the first decisions can be usefully revoked; (ii) for a medium delay cost $C_d = 0.025 \times \frac{t}{T}$ this percentage rises to 3.6%; and (iii) for a high delay cost $C_d = 0.5 \times \frac{t}{T}$ this percentage reaches 8%. These figures show that, for these datasets and this range of cost values, opportunities for a revocable strategy to overcome an irrevocable one seem seldom. (see [9] for a complete detailed analysis over all the 34 datasets and all couples of values (C_{cd}, C_d)).

However, the *first lesson* is that both revocable methods ECO-REV-CU and ECO-REV-CA get significantly better results than the irrevocable method ECONOMY- γ on a wide range of delay cost C_d and decision change cost values β (see Figures 2(a) and 2(b)). The *second lesson* is that it pays off to use a strategy which takes into account the costs of changing decision. Indeed, ECO-REV-CA beats ECONOMY- γ on a wider range of conditions than ECO-REV-CU.

Both revocable strategies fail to overcome the irrevocable one, ECONOMY- γ , when β is large (i.e. more than 0.1), and then ECO-REV-CU fails more often than ECO-REV-CA. This behavior is not surprising since, when it is very costly to delay a decision, the best strategy is generally to make a very early decision and not to revise it afterwards.

Figure 2(c) shows the results of the Wilcoxon signed-rank test between the two revocable strategies. It appears that the cost aware approach ECO-REV-CA performs significantly better than the cost unaware approach ECO-REV-CU, for almost one third of the pairs of values (α, β) . As the slope of the delay cost α grows, ECO-REV-CA becomes significantly better than ECO-REV-CU for an increasing larger range of values for β . This means that when the delay cost is rather high, it pays off to use a revocable strategy that takes into account the cost of changing decision. In addition, the Friedman test [17] shows that ECO-REV-CA is on average better ranked than ECO-REV-CU in 96% of pairs (α, β) . (Further details are available in [9]).

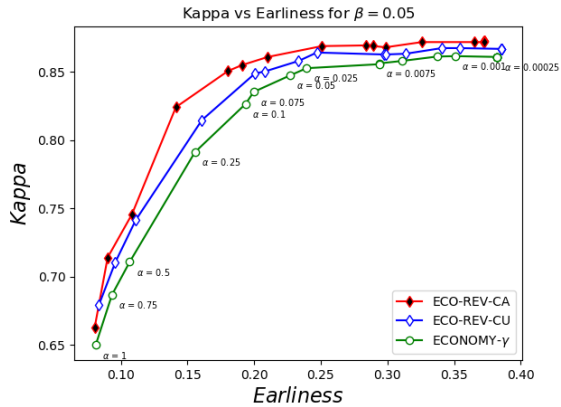


Fig. 1: Average *Earliness* vs. Average *Kappa* score obtained over all the 34 datasets for $\beta = 0.05$ and by varying the slope α of the delay cost. The reader may find the same behavior for other β values in the [9]

In order to get a global view of the merits of each method, we have drawn *Pareto curves* (see Figure 1) with respect to the average Cohen’s *kappa* score [7] and the average *earliness*, which is defined as the mean of the last triggering times normalized by the length of the series $earliness = Avg\{t_\ell/T\}$. These two quantities are averaged over the 34 datasets by varying α in the range of values defined in Section V-D, and for $\beta = 0.05$. The Pareto curves confirm that (i) the baseline irrevocable ECONOMY- γ method is dominated by the two revocable strategies; (ii) and ECO-REV-CA dominates ECO-REV-CU. More finely, it is apparent that, as the slope α of the delay cost increases, from 0.00025 to 1, all methods first maintain a high kappa, before being unable to maintain it as they are forced to make decisions too early. Still, the ECO-REV-CA algorithm is the one that best resists.

Overall, our experiments show the interest of using revocable strategies for the early classification of time series in a wide range of delay and change of decision costs.

VI. CONCLUSION

Until now, the problem of early classification of time series was addressed by triggering irrevocable decisions. For the first time, this paper defines the revocable version of this problem and introduces: i) the notion of *second order non-myopia*; ii) the corresponding optimization problem. Two versions of an algorithm have been implemented, one which takes into account the cost of changing the decision and the second which does not. Extensive experiments have shown that the algorithm which

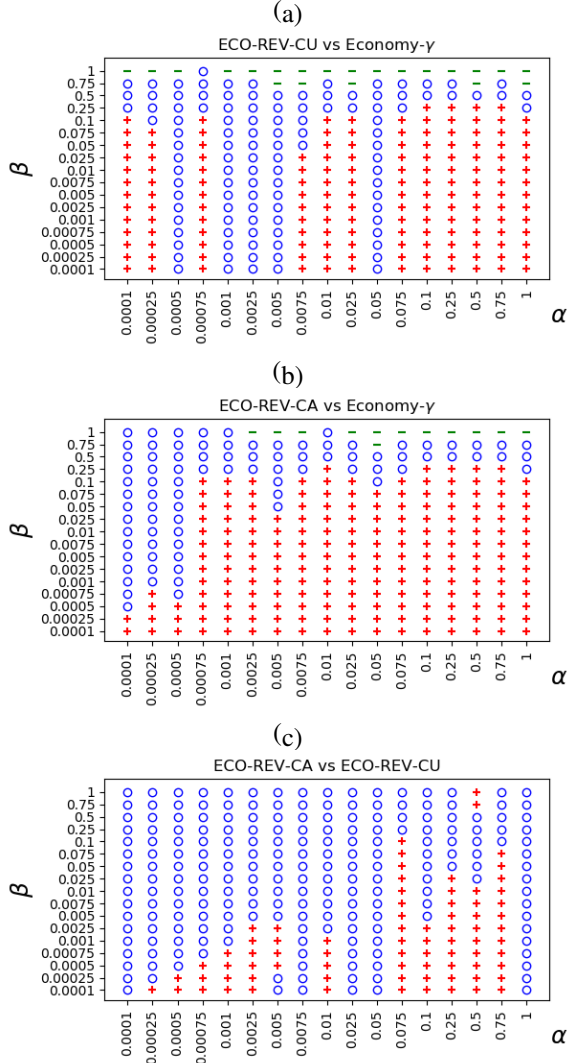


Fig. 2: (a) ECO-REV-CU vs. ECONOMY- γ ; (b) ECO-REV-CA vs. ECONOMY- γ ; (c) ECO-REV-CA vs. ECO-REV-CU. Wilcoxon signed-rank test applied on the *AvgCost* criterion over the 34 test sets, for a range of couples of values α and β , with “+” indicating a significant success of the first approach, “o” an insignificant difference and “-” indicating a significant failure of the first approach.

explicitly takes into account the cost of changing decisions, significantly overcomes the algorithm that does not. In addition, both proposed algorithms outperform the irrevocable scheme. The potential impact of these results on applications such as predictive maintenance or intensive care units, to name a few, is noteworthy.

REFERENCES

[1] Achenchabe, Y., Bondu, A., Cornuéjols, A., Dachraoui, A.: Early classification of time series. *Machine Learning* pp. 1–24 (2021)

[2] Anderson, H.S., Parrish, N., Tsukida, K., Gupta, M.: Early time-series classification with reliability guarantee. *Sandria Report* (2012)

[3] Bagnall, A., Lines, J., Bostrom, A., Large, J., Keogh, E.: The great time series classification bake off: a review and experimental evaluation of recent algorithmic advances. *Data Mining and Knowledge Discovery* **31**, 606–660 (2017)

[4] Barandas, M., Folgado, D., Fernandes, L., Santos, S., Abreu, M., Bota, P., Liu, H., Schultz, T., Gamboa, H.: Tsfel: Time series feature extraction library. *SoftwareX* **11**, 100456 (2020). <https://github.com/fraunhoferportugal/tsfel>

[5] Bondu, A., Achenchabe, Y., Bifet, A., Clérot, F., Cornuéjols, A., Gama, J., Hébrail, G., Lemaire, V., Marteau, P.F.: Open challenges for machine learning based early decision-making research. *arXiv preprint arXiv:2204.13111* (2022)

[6] Chen, T., Guestrin, C.: XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* pp. 785–794 (2016)

[7] Cohen, J.: A coefficient of agreement for nominal scales. *Educational and psychological measurement* **20**(1), 37–46 (1960)

[8] Dachraoui, A., Bondu, A., Cornuéjols, A.: 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 (2015)

[9] Eco-rev: Early and Revocable classification library. <https://github.com/eco-rev/rev-algo> (2021). Online; January 2021

[10] Eisenhower: (1944). <https://www.eisenhowerlibrary.gov/sites/default/files/research/online-documents/d-day/order-of-the-day.pdf>

[11] Ghalwash, M.F., Ramljak, D., Obradović, Z.: 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 (2012)

[12] Hatami, N., Chira, C.: Classifiers with a reject option for early time-series classification. In: *Computational Intelligence and Ensemble Learning (CIEL), 2013 IEEE Symposium on*, pp. 9–16. IEEE (2013)

[13] Khoshnevisan, F., Chi, M.: Unifying domain adaptation and domain generalization for robust prediction across minority racial groups. In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 521–537. Springer (2021)

[14] Mori, U., Mendiburu, A., Dasgupta, S., Lozano, J.: Early classification of time series from a cost minimization point of view. In: *Proceedings of the NIPS Time Series Workshop* (2015)

[15] Mori, U., Mendiburu, A., Dasgupta, S., Lozano, J.A.: Early classification of time series by simultaneously optimizing the accuracy and earliness. *IEEE transactions on neural networks and learning systems* **29**(10), 4569–4578 (2017)

[16] Mori, U., Mendiburu, A., Miranda, I.M., Lozano, J.A.: Early classification of time series using multi-objective optimization techniques. *Information Sciences* **492**, 204–218 (2019)

[17] Nemenyi, P.: Distribution-free multiple comparisons. *Biometrics* **18**(2), 263 (1962)

[18] Parrish, N., Anderson, H.S., Gupta, M.R., Hsiao, D.Y.: Classifying with confidence from incomplete information. *J. of Mach. Learning Research* **14**(1), 3561–3589 (2013)

[19] Rußwurm, M., Lefevre, S., Courty, N., Emonet, R., Körner, M., Tavenard, R.: End-to-end learning for early classification of time series. *arXiv preprint arXiv:1901.10681* (2019)

[20] Wilcoxon, F.: Individual comparisons by ranking methods. *Biometrics Bulletin* **1**(6), 80–83 (1945). URL <http://www.jstor.org/stable/3001968>