Early and Revocable Time Series Classification

Youssef Achenchabe   YOUSSEF.ACHENCHABE@UNIVERSITE-PARIS-SACLAY.FR
Paris-Saclay University
Orange Labs

Alexis Bondu   ALEXIS.BONDU@ORANGE.COM
Orange Labs

Antoine Cornuéjols   ANTOINE.CORNUEJOLS@AGROPARISTECH.FR
Paris-Saclay University

Vincent Lemaire   VINCENT.LEMAIRE@ORANGE.COM
Orange Labs

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.

Keywords: revocable decisions, cost estimation, online decision making

1. Introduction

Consider Eisenhower, in June 1944, having to decide when to launch the landing on the French coast (Eisenhower (1944)). 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.

Here again, we are faced with an an early classification problem: having to decide to make a prediction about a situation before all information becomes available because it is costly to wait, but now having the opportunity to change the prediction made if needed, potentially several times.

We can liken our work to what is studied in control theory (Bennett (1993)). Control theory is concerned as well with online decision making. For instance, when firing a rocket, the engines are controlled in strength and direction such as to maintain the rocket on the correct course. While the control here is instantaneous, other examples involve some prediction at what is likely to happen. This is the case with anti-aircraft guns that must point ahead of the plane in order to hit it. And the farther away is the plane from the gun, the more ahead of the plane should the pointing be. Notice that a whole lot of standard rules in control theory use as well integration of measurements over some time interval, thus taking stock of the past to take decisions.

Control theory is the province of smart engineers who know how to design mathematical formula that capture all of the knowledge pertinent for the problems to be solved. In numerous fields however, it is difficult or impossible to formulate such mathematical rules, either because the field is ripped with intricate and numerous factors, like in biology, economy or sociology, or because the environment changes in such a way that it is not worth trying to find mathematical formulas that would soon be useless. In these cases, one way to circumvent this difficulty is to rely on heuristic rules learned from data representative of the environment using machine learning methods.

Early classification of time series goes one step further and tries to learn second order knowledge. The idea is to learn *when* there will be enough information to decide in such a way as to optimize both the expectation of the misclassification cost and the cost associated with delaying the decision.

Back to control theory, tremendous progress has been achieved when feedback signals have been taken into account. Then, the system, by observing the discrepancy between the prediction and further measurements, can correct its earlier decisions. Again, this is what is done in anti-aircraft systems. And, here also, these feedback loops which control by how much to correct, after which time interval, and so on, have to be designed by engineers using their knowledge of the domain.

In the early classification of time series context, this translates into revocable early classification of time series. There, the system is allowed to estimate the best time to make a decision, but also to estimate when to revise prior decisions, if the need seems warranted. There also, the rules can be learned from the available data. For instance, it can be anticipated that large amounts of such data are and will become available in the domain of autonomous vehicles.

To illustrate in a more concrete way the usefulness of extending Early Classification to revocable decisions, let us consider the example of an emergency stop system implemented in an autonomous car. Let us assume that this system is equipped with several sensors, consisting of radars and cameras, which scan the road in order to detect a possible obstacle. The reliability of these measurements decreases when the distance between the car and the observed point increases. The car is driven at 120 km/h on the highway and a dark shape

located 300 meters ahead is detected on the camera image. At this moment, there is a doubt on the nature of this shape which could be an obstacle on the road with a probability of 0.008. Thus, the system decides to brake. While approaching, the image of the camera becomes more precise and the radars which have a range limited to 100 meters can now be used: it does not detect an obstacle. The system decides to release the brake because it recognizes a dark spot on the road which is at the origin of the shape initially observed.

Some 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 $\mathbf{x}_T = \langle x_1, \ldots, x_T \rangle$ 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: $\mathbf{x}_t = \langle x_1, \ldots, x_t \rangle$ 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 $\mathcal{D}_k$, the sequence of the k predictions $\langle \hat{y}_{t_1}, \ldots, \hat{y}_{t_k} \rangle$ made at times $t_1, \ldots, t_k$ in the time interval $\{1, \ldots, T\}$. Each decision change from $\hat{y}_{t_i}$ to $\hat{y}_{t_{i+1}}$ entails a cost $C_{cd}(\hat{y}_{t_{i+1}}|\hat{y}_{t_i})$. Then, the cost endured by the decision maker at time T is given by Equation 3 which sums all previously introduced costs.

Interestingly, 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. One reason could be that this does not seem worthwhile. Are there so many situations, after all, where changing decisions could reduce the overall cost of the expected misclassification cost and of the cost associated with further delays? A second reason is that the problem seems difficult. Our work shows in effect that it is not straightforward to identify the best instants to revoke a decision. But yet, it shows also that such revocations can yield significant gains. This is demonstrated by the statistics presented at the beginning of the experimental section: indeed, there are very few useful revocations for sure, but the system presented is able to identify them. And, as shown by comparison between our approach and the best known irrevocable method, to our knowledge, the gain in performance is statistically significant.

The contribution of this paper is threefold. First, it formalizes the optimization problem associated with the revocable regime for the early classification problem. Second, it proposes two approaches to tackle this problem. Both approaches are *non-myopic* in that they take into account expectancies about likely futures to take their decisions. The first approach does not consider explicitly the cost of changing decision, while the second one does. Third, extensive experiments are presented that both allow the comparison of the two approaches and show that it is actually better to be able to revise decisions than to implement an irrevocable decision strategy.

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). The framework and approaches presented in this paper can be adapted to multivariate time series in a trivial way.

The rest of this paper is organized as follows. Section 2 provides an overview of classical *early classification* approaches, all of which deal with the irrevocable regime. Section 3 focuses on a non-myopic framework which is designed for the irrevocable regime. The *early and revocable classification* problem is defined in Section 4. Then, two new approaches are proposed, which are evaluated through extensive experiments in Section 5. Perspectives and future work are discussed in Section 6.

2. State of the art on early classification

All works that we are aware of deal with the early classification of time series problem in the irrevocable regime. Most of these works do not take into account explicitly the costs associated with a decision: the cost of misclassification and the cost of delaying decision. Instead, they generally base their decision on some form of confidence criterion and wait until a predefined threshold is reached before making a decision. For instance, in (Xing et al. (2009)), the best time step to trigger the decision is estimated by determining the earliest time step for which the predicted label does not change, based on a 1NN classifier. Similarly, (Mori et al. (2017b)) proposes a method where the accuracy of a set of probabilistic classifiers is monitored over time, which allows the identification of time steps from whence it seems safe to make predictions. In (Parrish et al. (2013); Hatami and Chira (2013); Ghalwash et al. (2012)), a classifier $h_t(\cdot)$ is learned for each time step and various stopping rules are defined (e.g. threshold on the confidence level).

With the advent of deep learning, researchers tried to revisit the problem of early classification in the light of these modern techniques. (Rußwurm et al. (2019)) proposed a trainable framework for early classification of time series that can be fine-tuned end-to-end using standard gradient back-propagation. (Martinez et al. (2018)) and (Hartvigsen et al. (2019)) approach early classification as a reinforcement learning problem, a perspective extended to multi-label classification in (Hartvigsen et al. (2020)).

Few works do explicitly take costs into account. A notable example is (Mori et al. (2019)) where the conflict between earliness and accuracy is explicitly addressed. Moreover, instead of setting the trade-off in a single objective optimization criterion as in (Mori et al. (2017a)), the authors keep it as a multi-objective criterion and to explore the Pareto front of the multiple dominating trade-offs.

It is however in (Dachraoui et al. (2015)), that the early classification problem is for the first time 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. Importantly, besides the fact that this optimization criterion is well-founded, it permits also the expected costs for an incoming subsequence $\mathbf{x}_t$ to be estimated for future time steps. A non-myopic decision procedure can thus be used. These expectations about the foreseeable future of an incoming time series can be learned from the training set of m full-length time series $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1 \leq i \leq m}$.

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.

In the following, we therefore focus our attention on a setting where the costs are explicit factors entering the optimization problem.

3. A cost-based non-myopic framework

In this section, we introduce a cost-based non-myopic framework that was designed for the irrevocable regime (Achenchabe et al. (2021)). The following sections will show how it can be adapted to the revocable regime.

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

I- For each time step, $t \in \{1, \ldots, T\}$, a classifier h_t can be learned $h_t : \mathcal{X}^t \to \mathcal{Y}$. Note that this contrasts with learning from data stream where the world can be non-stationary and the classifiers might have to evolve over time.

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[1].

More precisely, given the *misclassification cost* function $C_m(\hat{y}|y) : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ that expresses the cost of predicting $\hat{y}$ when the true class is y and the *delay cost* function $C_d(t) : \mathbb{R} \to \mathbb{R}$ which is assumed to be an increasing function of time, the expectancy of the cost of taking a decision at time t given the incoming time series $\mathbf{x}_t$ is:

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

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:

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

---

1. This can be seen as an instance of 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 $\mathcal{S} = \{(\mathbf{x}_T^i, y_i)\}_{1\leq i\leq m}$, while at testing time, only a subsequence $\mathbf{x}_t$ $(t < T)$ is known.

and $f_0(\mathbf{x}_t) = \mathbb{E}^{t}_{y \in \mathcal{Y}}[C_m(\hat{y}_t|y)] + C_d(t)$ since we have access to predictions at current time. Then the optimal decision time, at time t, is expected to be: $\tau^* = \text{ArgMin}_{\tau \in \{0,\ldots,T-t\}} f_\tau(\mathbf{x}_t)$

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 is postponed to the next time step, and Eq. 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.

While this irrevocable decision process is well-founded and has proven to be quite efficient in extensive experiments (Achenchabe et al. (2021)), it can nonetheless lead to non optimal decisions when the estimated expected future cost of decision reveals itself to be erroneous. Figure 1 provides an example where at time t, it is expected that all future instants will lead to worse costs, when actually a better decision time occurs later, that could even have been anticipated given just a few additional measurements.
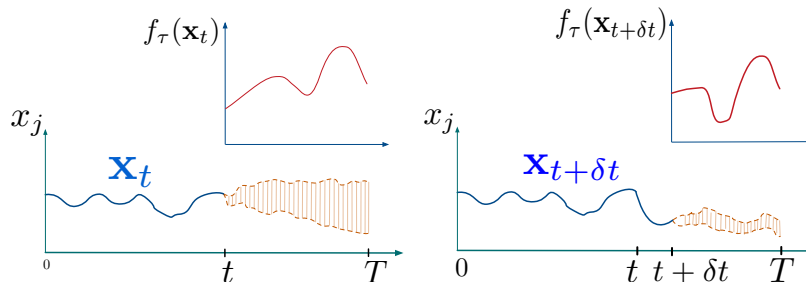

Figure 1: (Left) At time t, the system foresees that this is the best time to make a decision, $f_\tau(\mathbf{x}_t)$ is at its lowest. (Right) At some ulterior time $t + \delta t$, it appears that there should be a better time for decision in the future.

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 associated with these changes.

4. 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 $\mathcal{D}_\ell = \langle \hat{y}_{t_1}, \ldots, \hat{y}_{t_\ell} \rangle$ about the class of the input time series. The final cost incurred will be:

$$g(\mathcal{D}_\ell|\mathbf{x}_T, y) = C_m(\hat{y}_{t_\ell}|y) + C_d(t_\ell) + \sum_{\substack{i=1 \\ \hat{y}_{t_i}, \hat{y}_{t_{i+1}} \in \mathcal{D}_\ell}}^{\ell-1} C_{cd}(\hat{y}_{t_{i+1}}|\hat{y}_{t_i}) \tag{3}$$

6

where t_ℓ is the time of the last change of decision yielding the prediction $\hat{y}_{t_\ell} = h_{t_\ell}(\mathbf{x}_{t_\ell})$, when the true class of $\mathbf{x}_T$ is y. Moreover, the cost of changing decision is defined as $\mathrm{C}_{cd}(\hat{y}_1|\hat{y}_2) : \mathcal{Y} \times \mathcal{Y} \to \mathbb{R}$ with $\mathrm{C}_{cd}(\hat{y}_1|\hat{y}_2) = 0$ if $\hat{y}_1 = \hat{y}_2$.

Formally, the problem is to find a sequence of decisions $\mathcal{D}^\star \in \mathbb{D}_T$ that minimizes Equation 3:

$$\mathcal{D}^\star = \underset{\mathcal{D} \in \mathbb{D}_T}{\mathrm{ArgMin}}\, g(\mathcal{D}|\mathbf{x}_T, y) \tag{4}$$

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

When, at time t, only a partial knowledge $\mathbf{x}_t$ is available about the incoming time series, Equation 3 cannot be computed. A sequence of decisions $\mathcal{D}_k = \langle \hat{y}_{t_1}, \ldots, \hat{y}_{t_k} \rangle$ 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 $t + \tau$.

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

$$f_\tau^{\mathrm{rev}}(\mathcal{D}_k, \tilde{t} \,|\, \mathbf{x}_t) = \mathbb{E}_{(\hat{y},y) \in \mathcal{Y}^2}^{t+\tau} [\mathrm{C}_m(\hat{y}|y)|\, \mathbf{x}_t] + \sum_{\substack{i=1 \\ \hat{y}_{t_i}, \hat{y}_{t_{i+1}} \in \mathcal{D}_k}}^{k-1} \mathrm{C}_{cd}(\hat{y}_{t_{i+1}}|\hat{y}_{t_i}) +$$
$$\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [\mathrm{C}_{cd}(\hat{y}|\hat{y}_{t_k})|\, \mathbf{x}_t] + \mathrm{C}_d(\tilde{t}) \tag{5}$$

and for $\tau = 0$,

$$f_{\tau=0}^{\mathrm{rev}}(\mathcal{D}_k, \tilde{t} \,|\, \mathbf{x}_t) = \mathbb{E}_{y \in \mathcal{Y}}^{t} [\mathrm{C}_m(\hat{y}_t|y)|\, \mathbf{x}_t] + \sum_{\substack{i=1 \\ \hat{y}_{t_i}, \hat{y}_{t_{i+1}} \in \mathcal{D}_k}}^{k-1} \mathrm{C}_{cd}(\hat{y}_{t_{i+1}}|\hat{y}_{t_i}) + \mathrm{C}_d(\tilde{t})$$

where $\tilde{t}$ is introduced as input of the function f_τ^{rev} in order to control the delay cost paid by the user. In addition, the expected cost of changing decision is defined as follows for $(1 \leq \tau \leq T - t)$:

$$\mathbb{E}_{\hat{y} \in \mathcal{Y}}^{t+\tau} [\mathrm{C}_{cd}(\hat{y}|\hat{y}_{t_k})|\, \mathbf{x}_t] = \sum_{\hat{y} \in \mathcal{Y}} P_{t+\tau}(\hat{y}|\hat{y}_{t_k}, \mathbf{x}_t) \mathrm{C}_{cd}(\hat{y}|\hat{y}_{t_k}) \tag{6}$$

and for $\tau = 0$, this term equals zero by convention because we have full knowledge of the prediction at the current time step t. Then, given that the notation $\mathcal{D}_{k+1}$ is used to denote the sequence of decisions $\langle \hat{y}_{t_1}, \ldots, \hat{y}_{t_k}, \hat{y}_t \rangle$, the criterion for changing decision at time t becomes:

$$criterion = \begin{cases} & \hat{y}_t \neq \hat{y}_{t_k} \\ \mathrm{and} & \underset{\tau \in \{0, \ldots, T-t\}}{\mathrm{ArgMin}}\, f_\tau^{\mathrm{rev}}(\mathcal{D}_k, t+\tau \,|\, \mathbf{x}_t) = 0 \\ \mathrm{and} & f_{\tau=0}^{\mathrm{rev}}(\mathcal{D}_{k+1}, t \,|\, \mathbf{x}_t) < f_{\tau=0}^{\mathrm{rev}}(\mathcal{D}_k, t_k \,|\, \mathbf{x}_t) \end{cases} \tag{7}$$

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 5 becomes:

$$f_\tau^{\text{rev}}(\mathcal{D}_k, \tilde{t} \,|\, \mathbf{x}_t) \;=\; \mathbb{E}_{(\hat{\mathbf{y}},\mathbf{y}) \in \mathcal{Y}^2}^{t+\tau} \left[C_m(\hat{y}|y) | \mathbf{x}_t \right] \;+\; C_d(\tilde{t}) \tag{8}$$

which is Equation 2. Then, the strategy is to change decision each time 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** if the decision maker had access to the true class y of the incoming time series but could only use the classifiers $h_t(t \in \{1, \ldots, T\})$ to make its prediction? (Thus, it could not output y before the first time t st. $h_t(\mathbf{x}_t) = \hat{y}_t = y$).

Theorem 1 (Optimal sequence of decisions when $\forall (y, y') \in \mathcal{Y}^2$, $C_{cd}(y|y') > 0$) *For any time series $\mathbf{x}_T$ of class y, the optimal sequence of decision is reduced to a one decision sequence where the (or one of possibly several) optimal time(s) $t^\star$ is defined by: $t^\star = \text{ArgMin}_{1 \leq t \leq T}\{C_m(\hat{y}_t|y) + C_d(t)\}$.*

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

Theorem 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, without having access to the ground truth y, it may be beneficial to make a first guess and to change it later on.

In a consistent way, f_τ^{rev} envisions a single future decision $\hat{y}_{t+\tau}$, which entails a minimal cost compared to longer sequences involving several decision changes between t and $t + \tau$. Indeed, the costs associated with the successive decision changes would be added to the Equation 5, and would necessarily lead to a higher total cost.

It must be noted that the *criterion* (Equation 7) 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: $f_\tau^{\text{rev}}(\emptyset \,|\, \mathbf{x}_t) = +\infty$. By contrast, when a first decision is taken, its expected cost is:

$$f_{\tau=0}^{\text{rev}}(\langle \hat{y}_1 \rangle, t, \,|\, \mathbf{x}_t) = \mathbb{E}_{y \in \mathcal{Y}^2}^{t} \left[C_m(\hat{y}_t|y)|\mathbf{x}_t \right] + 0 + C_d(t)$$

Consequently, $f_{\tau=0}^{\text{rev}}(\langle \hat{y}_1 \rangle \,|\, \mathbf{x}_t) < f_{\tau=0}^{\text{rev}}(\emptyset \,|\, \mathbf{x}_t)$ and the first decision is made accordingly to the non-myopic strategy defined by $f_\tau(\mathbf{x}_t)$ in its irrevocable regime (see Equation 1).

A generic algorithmic implementation of the revocable decision-making criterion as defined in Equation 7 is presented in Algorithm 1. The sequence of decisions is initialized with $\langle \hat{y}_{t_1} \rangle$, where $\hat{y}_{t_1}$ is provided using the irrevocable strategy based on $f_\tau(\mathbf{x}_t)$ (see Equation 1) and t_1 is the corresponding triggering time. The cost of the previous decision $cost_{prev}$ is initialized with the cost of the first decision $f_{\tau=0}(\mathbf{x}_{t_1})$. A new decision is triggered if the conditions of Equation 7 are satisfied.

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 coded two algorithms.

1. The first one is named ECO-REV-CU for cost unaware (as in Equation 8).

2. The second is named ECO-REV-CA for cost aware (as in Equation 5).

Algorithm 1 Generic revocable regime algorithm

$decisions \leftarrow \langle \hat{y}_{t_1} \rangle$
$t_{prev} \leftarrow t_1$
for $t = t_1 + 1 \ldots T$ **do**
    $\tau^\star \leftarrow ArgMin_{\tau \in \{0\ldots T-t\}} f_\tau^{\text{rev}}(decisions, t + \tau | \mathbf{x}_t)$
    $cost_{new} \leftarrow f_{\tau=0}^{\text{rev}}(decisions, t + \tau^* | \mathbf{x}_t)$
    $cost_{prev} \leftarrow f_{\tau=0}^{\text{rev}}(decisions, t_{prev} | \mathbf{x}_t)$
    if $\hat{y}_t \neq \hat{y}_{t_{prev}}$ **and** $\tau^* = 0$ **and** $cost_{new} < cost_{prev}$ **then**
        $t_{prev} = t$
        $decisions \leftarrow decisions \cup \hat{y}_t$
end
$return decisions$

---

5. Experiments

The criterion defined by Equation 7 and its implementation in Algorithm 1 aim at identifying the advantageous changes of decision, those that take advantage of the knowledge gained with additional measurements of the incoming time series. With the experiments, we want to measure the true added value of this strategy. Specifically, the question is twofold. First, does it 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 one that would not consider these costs? In the following, we report results obtained on 34 datasets (see Section 5.2) for a whole range of values for the delay cost C_d and the cost incurred if changing decision C_{cd}.

5.1. Implementation choices

Equation 2 for the irrevocable strategy has been proposed and has given way to several different algorithmic versions generically called ECONOMY as described in (Achenchabe et al. (2021)). They differ in the way they group time series in order to estimate the expectation $\mathbb{E}_{(\hat{y},y) \in \mathcal{Y}^2}^{t+\tau} [C_m(\hat{y}|y)]$. Of all these methods, ECONOMY-γ is the one that stands out, both because its refined way of predicting the likely future of an incoming time series and its significantly better performances demonstrated in extensive experiments over the other ECONOMY versions as well as with the method of (Mori et al. (2017a)). This is why it is the method used in our experiments, both in a revocable version that takes the cost of changing decisions into account, and one that does not.

Implementation of the two proposed approaches: in our experiments, the ECO-REV-CU algorithm is simply the ECONOMY-γ algorithm allowed to be reiterated 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 $\mathbb{E}_{\hat{y}\in\mathcal{Y}}^{\,t+\tau}\left[C_{cd}(\hat{y}|\hat{y}_{t_k})|\,\mathbf{x}_t\right]$ in Equation 6 by using the groups of time series, denoted by $\mathcal{G}$:

$$\mathbb{E}_{\hat{y}\in\mathcal{Y}}^{\,t+\tau}\left[C_{cd}(\hat{y}|\hat{y}_{t_k})|\,\mathbf{x}_t\right] \approx \sum_{\mathfrak{g}_k\in\mathcal{G}} P(\mathfrak{g}_k|\mathbf{x}_t)\,\mathbb{E}_{\hat{y}\in\mathcal{Y}}^{\,t+\tau}\left[C_{cd}(\hat{y}|\hat{y}_{t_k})|\,\mathfrak{g}_k\right] \tag{9}$$

Then, the probability $P_{t+\tau}(\hat{y}|\hat{y}_{t_k}, \mathfrak{g}_k)$ 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 the supplementary material.

Overview of the Economy-γ approach: more precisely, the groups $\mathcal{G}$ are obtained by stratifying the time series by confidence levels[2] of h_t. 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 $\{I_t^1, \ldots, I_t^K\}$. Then, the future expectancy $\mathbb{E}_{y\in\mathcal{Y}^2}^{\,t+\tau}\left[C_m(\hat{y}_t|y)|\mathbf{x}_t\right]$ is estimated by modeling the term $P_{t+\tau}(\hat{y}|y, \mathbf{x}_t \in \mathfrak{I}_k^t)$ as a projection into the future of the probability distribution over $\mathfrak{I}_k^{t+\tau}$, the confidence intervals of $h_{t+\tau}$. A Markov-chain model is used for this purpose. A fully detailed description of the ECONOMY-γ is provided in (Achenchabe et al. (2021)).

5.2. 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 (Achenchabe et al. (2021)), we chose to use the same 34 datasets that are taken from the UEA & UCR Time Series Classification Repository[3] (Bagnall et al. (2017)). However, it is important to note that the revocable framework presented here could as well accommodate multi-class classification problems. Additional experiments on multi-class problem are reported in the supplementary material

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 (Chen and Guestrin (2016)) classifiers $\{h_t\}_{t\in\{1,\ldots,T\}}$ 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 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

---

2. This restricts these methods to binary classification problems.

3. Available at: http://www.timeseriesclassification.com

spectral domains are extracted using the Time Series Feature Extraction Library (Barandas et al. (2020)), and are used for training the classifiers $\{h_t\}_{t \in \{1,...,T\}}$.

5.3. 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:

$$\text{Cost}(\mathbf{x}_T) \;=\; \text{C}_m(h_{t_l}(\mathbf{x}_{t_l})|y) + \text{C}_d(t_l) + \sum_{i=1}^{|\mathcal{D}_l|-1} C_{cd}(\hat{y}_{i+1}|\hat{y}_i) \tag{10}$$

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

$$AvgCost(\mathcal{T}) \;=\; \frac{1}{|\mathcal{T}|} \sum_{i=1}^{|\mathcal{T}|} \text{Cost}(\mathbf{x}_T^i) \tag{11}$$

5.4. 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, finally, ECO-REV-CA the revocable decision-maker that 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, it is expected that the various costs, relative to misclassifications, delays and changes of decision, will be provided by the domain expert. For our experiments, we explored the performance of the three methods on a wide range of cost values:

- The *misclassification cost* was set to $\text{C}_m(\hat{y}|y) = 1$ 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 $\text{C}_{cd}(\hat{y_1}|\hat{y_2}) = \beta$ if $\hat{y_1} \neq \hat{y_2}$, and $= 0$ otherwise. The parameter β being taken in the same set of values as α[4]

The AvgCost criterion defined in Equation 11 was evaluated on the 34 test sets for all cost values, and the Wilcoxon signed-rank test (Wilcoxon (1945)) was performed for all the range of cost values, in order to assess whether the observed performance gap between methods is significant *("+" and "-")* or not *("∘")*.

---

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

5.5. 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, it is not clear that a significant added-value can be found in favor of the revocable strategies. That such an advantage is nonetheless what the experiments bring out is thus remarkable (see Supplementary Material for a more detailed analysis).

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 3(a) and 3(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.
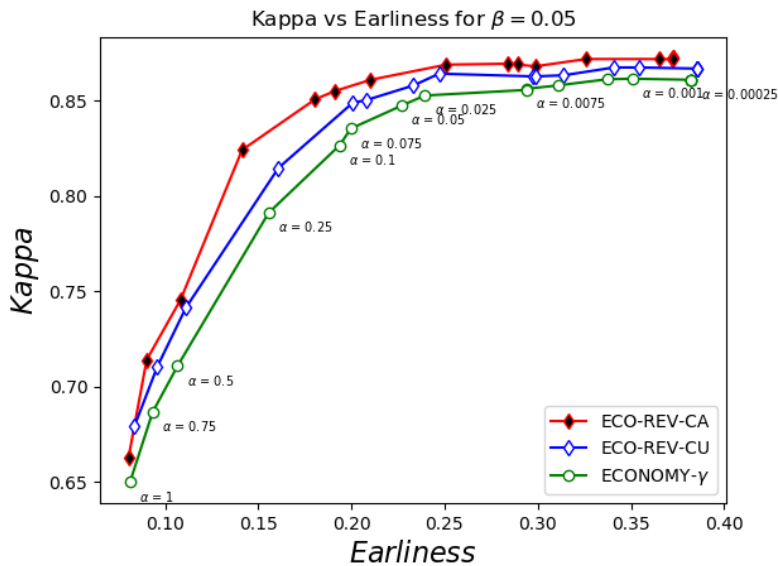

Figure 2: Average *Earliness* vs. Average *Kappa* score obtained over all the 34 datasets for $\beta = 0.05$ and by varying the slope α of the delay cost. The reader may find the same behavior for other β values in the supplementary material.

Figure 3(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 (Nemenyi (1962)) 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 the the supplementary material).

In order to get a global view of the merits of each method, we have drawn *Pareto curves* (see Figure 2) with respect to the average Cohen's *kappa* score (Cohen (1960)) 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 5.4, and for $\beta = 0.05$. The Pareto curves that are obtained show that (i) the baseline irrevocable ECONOMY-γ method is dominated by the two revocable strategies; (ii) and ECO-REV-CA dominates ECO-REV-CU. These results are consistent with the ones cited above and hold for other values of β a well (see the supplementary material for more details). 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. They demonstrate also the relevance of the formal criterion that we propose for these strategies.

6. Conclusion

Applications abound, in which incoming time series must be labeled as early and as accurately as possible, before all measurements are available. 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 an associated optimization problem. Optimization is performed over the space of all possible sequences of decisions given an incoming time series. Thanks to a non-myopic criterion given in the paper its exploration is simplified and an algorithm follows naturally. This algorithm has been implemented in two versions, the first one takes into account the cost of changing the decision and the second one does not. Extensive experiments have shown that this algorithm, which explicitly takes into account the cost of changing decisions, has significantly better performances than the same algorithm that does not. In addition, comparison with the irrevocable regime (ECONOMY-γ) shows that the two proposed algorithms make useful revocations, because they both outperform the irrevocable regime.

As future work, we plan to adapt the proposed framework to the online setting, where decisions should be taken based on a data stream. We will also improve the multi-class approaches proposed in the supplementary material.
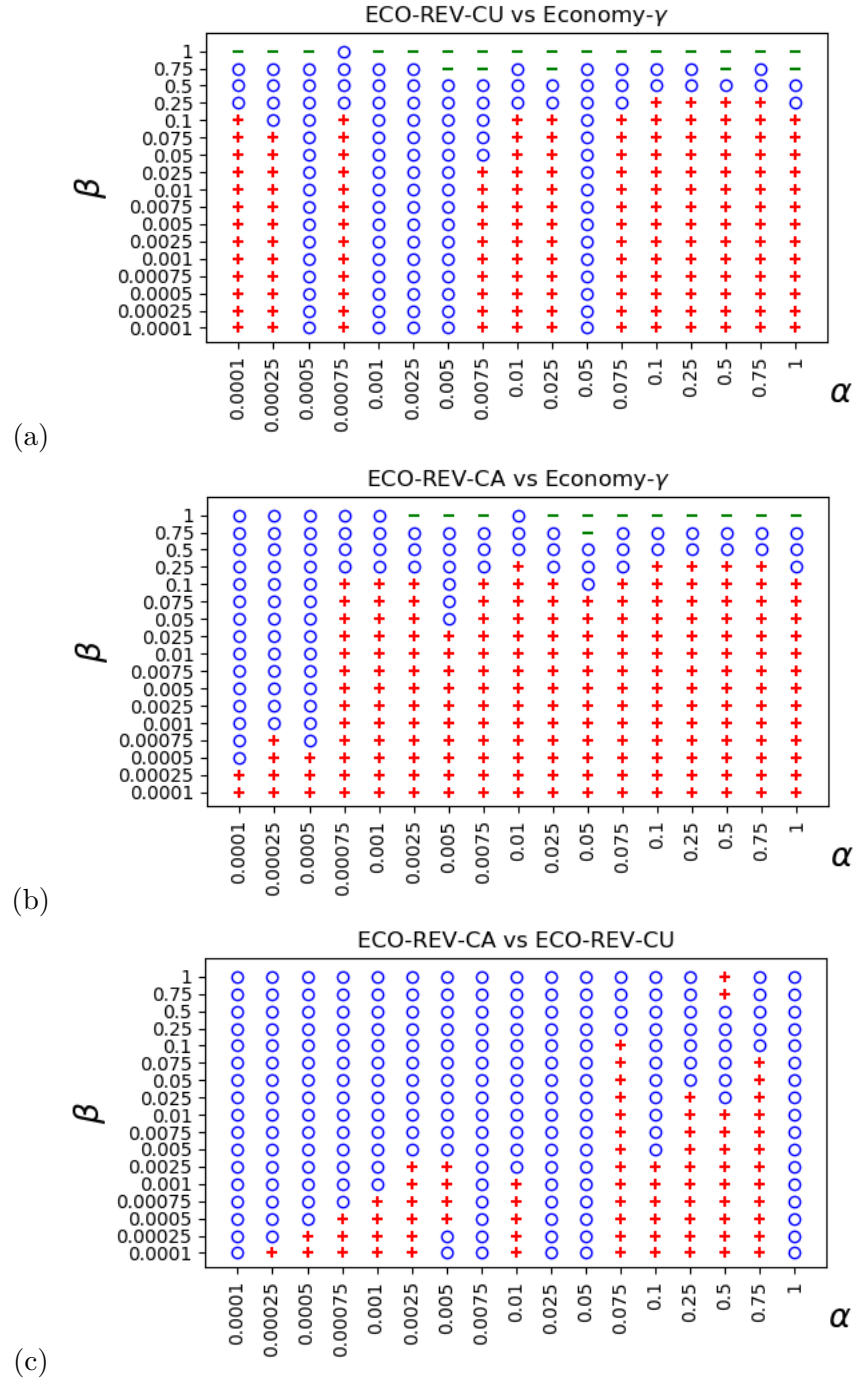
(a)

(b)

(c)

Figure 3: (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, "∘" an insignificant difference and "−" indicating a significant failure of the first approach.

References

Youssef Achenchabe, Alexis Bondu, Antoine Cornuéjols, and Asma Dachraoui. Early classification of time series. *Machine Learning*, pages 1–24, 2021.

A. Bagnall, J. Lines, A. Bostrom, J. Large, and E. Keogh. 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.

Marília Barandas, Duarte Folgado, Letícia Fernandes, Sara Santos, Mariana Abreu, Patrícia Bota, Hui Liu, Tanja Schultz, and Hugo Gamboa. Tsfel: Time series feature extraction library. *SoftwareX*, 11:100456, 2020. https://github.com/fraunhoferportugal/tsfel.

Stuart Bennett. *A history of control engineering, 1930-1955.* Number 47. IET, 1993.

Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

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

Asma Dachraoui, Alexis Bondu, and Antoine Cornuéjols. 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*, pages 433–447. Springer, 2015.

Eisenhower, 1944. https://www.eisenhowerlibrary.gov/sites/default/files/research/online-documents/d-day/order-of-the-day.pdf.

Mohamed F Ghalwash, Dušan Ramljak, and Zoran Obradović. Early classification of multivariate time series using a hybrid hmm/svm model. In *2012 IEEE International Conference on Bioinformatics and Biomedicine*, pages 1–6. IEEE, 2012.

Thomas Hartvigsen, Cansu Sen, Xiangnan Kong, and Elke Rundensteiner. Adaptive-halting policy network for early classification. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 101–110, 2019.

Thomas Hartvigsen, Cansu Sen, Xiangnan Kong, and Elke Rundensteiner. Recurrent halting chain for early multi-label classification. 2020.

Nima Hatami and Camelia Chira. Classifiers with a reject option for early time-series classification. In *Computational Intelligence and Ensemble Learning (CIEL), 2013 IEEE Symposium on*, pages 9–16. IEEE, 2013.

Coralie Martinez, Guillaume Perrin, Emmanuel Ramasso, and Michèle Rombaut. A deep reinforcement learning approach for early classification of time series. In *2018 26th European Signal Processing Conference (EUSIPCO)*, pages 2030–2034. IEEE, 2018.

Usue Mori, Alexander Mendiburu, Sanjoy Dasgupta, and Jose A Lozano. Early classification of time series by simultaneously optimizing the accuracy and earliness. *IEEE transactions on neural networks and learning systems*, 29(10):4569–4578, 2017a.

Usue Mori, Alexander Mendiburu, Eamonn Keogh, and Jose A Lozano. Reliable early classification of time series based on discriminating the classes over time. *Data mining and knowledge discovery*, 31(1):233–263, 2017b.

Usue Mori, Alexander Mendiburu, Isabel Marta Miranda, and José Antonio Lozano. Early classification of time series using multi-objective optimization techniques. *Information Sciences*, 492:204–218, 2019.

Peter Nemenyi. Distribution-free multiple comparisons. *Biometrics*, 18(2):263, 1962.

Nathan Parrish, Hyrum S Anderson, Maya R Gupta, and Dun Yu Hsiao. Classifying with confidence from incomplete information. *J. of Mach. Learning Research*, 14(1):3561–3589, 2013.

Marc Rußwurm, Sébastien Lefevre, Nicolas Courty, Rémi Emonet, Marco Körner, and Romain Tavenard. End-to-end learning for early classification of time series. *arXiv preprint arXiv:1901.10681*, 2019.

Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009.

Frank Wilcoxon. Individual comparisons by ranking methods. *Biometrics Bulletin*, 1(6): 80–83, 1945. ISSN 00994987. URL http://www.jstor.org/stable/3001968.

Zhengzheng Xing, Jian Pei, and S Yu Philip. Early prediction on time series: A nearest neighbor approach. In *IJCAI*, pages 1297–1302. Citeseer, 2009.