

# Early Classification of Time Series in Non-Stationary Cost Regimes

Aurélien Renault  
Orange Research, France  
Email: aurelien.renault@orange.com

Antoine Cornuéjols  
AgroParisTech, France  
Email: antoine.cornuejols@agroparistech.fr

Alexis Bondu  
Orange Research, France  
Email: alexis.bondu@orange.com

Vincent Lemaire  
Orange Research, France  
Email: vincent.lemaire@orange.com

**Abstract**—Early Classification of Time Series (ECTS) addresses decision-making problems in which predictions must be made as early as possible while maintaining high accuracy. Most existing ECTS methods assume that the time-dependent decision costs governing the learning objective are known, fixed, and correctly specified. In practice, however, these costs are often uncertain and may change over time, leading to mismatches between training-time and deployment-time objectives. In this paper, we study ECTS under two practically relevant forms of cost non-stationarity: drift in the balance between misclassification and decision delay costs, and stochastic realizations of decision costs that deviate from the nominal training-time model. To address these challenges, we revisit representative ECTS approaches and adapt them to an online learning setting. Focusing on separable methods, we update only the triggering model during deployment, while keeping the classifier fixed. We propose several online adaptations and baselines, including bandit-based and RL-based approaches, and conduct controlled experiments on synthetic data to systematically evaluate robustness under cost non-stationarity. Our results demonstrate that online learning can effectively improve the robustness of ECTS methods to cost drift, with RL-based strategies exhibiting strong and stable performance across varying cost regimes.

## I. MOTIVATION

Many real-world decision-making systems operate under *time pressure*, where delaying a prediction can be as costly as making an incorrect one. Representative examples include real-time monitoring, predictive maintenance, medical triage, and adaptive human-machine interaction. In such settings, the objective is not only to achieve high predictive accuracy, but to do so as early as possible. This inherent trade-off between classification accuracy and decision earliness lies at the core of *Early Classification of Time Series* (ECTS) and has been extensively studied in this field [1], [2], [3]. ECTS belongs to the broader class of early decision-making [4] and anytime prediction problems. More specifically, ECTS can be naturally formulated as an empirical risk minimization problem with a time-dependent loss function that jointly penalizes misclassification errors and decision delays (see Section II).

Most existing ECTS approaches, however, rely on the assumption that this time-dependent loss and the decision costs it encodes are known, fixed, and accurately specified

during training. In practice, these costs must be elicited from domain experts and are often subject to substantial uncertainty. At deployment time, the costs actually incurred may deviate from those assumed during training due to imperfect cost specification, evolving operational conditions, or intrinsic randomness across decision instances. Although cost mismatches are inevitable in practice, their impact on ECTS systems remains largely unexplored, making cost non-stationarity a fundamental challenge for ECTS and, more broadly, for online and continual neural learning systems. Concrete applications illustrate how decision costs may be non-stationary over time:

**I** - In healthcare monitoring [5], for instance, the *balance* between misclassification and decision delay costs evolves as clinical protocols, hospital workload, or available treatments change. Early detection of patient deterioration enables timely intervention; yet, premature alarms may lead to unnecessary treatments, inefficient resource allocation. Thus, the relative importance of accuracy versus earliness may drift over time, inducing a deployment-time objective that differs from the one assumed during training.

**II** - In addition to cost drift, *stochastic* decision costs naturally arise in industrial early defect detection. Consider a production line manufacturing composite mechanical parts, such as carbon-fiber-reinforced components used in aerospace or automotive industries. Defective parts may be ejected from the line as scrap, but the economic cost of rejection increases as manufacturing progresses: early rejection allows partial material recovery, while late rejection often makes recycling infeasible once curing or assembly steps are completed. Crucially, the realized cost of a late decision depends on the type of defect—unknown at prediction time—which affects the ability to separate and reuse constituent materials. Thus, the actual decision cost is inherently stochastic and may deviate from the nominal cost assumed during training.

Motivated by these observations, this paper provides a first systematic study of ECTS under two practically relevant forms of cost non-stationarity, which are explicitly addressed in our experimental evaluation: (i) drift in the balance between misclassification costs and decision delay costs, and (ii) stochastic

realizations of decision costs that deviate from the nominal training-time regime. To tackle these challenges, we revisit the ECTS literature with the goal of adapting widely used approaches to an online learning setting, a direction that has not been investigated to our knowledge. We focus on *separable* ECTS methods, which first learn a classifier and subsequently learn a triggering model, as this paradigm encompasses the vast majority of existing work [1], [2], [3]. In our setting, only the triggering model is updated during deployment, while the classifier remains fixed, reflecting that cost drift primarily affects decision timing rather than class discrimination. Our contributions are threefold:

- **Online adaptation of representative ECTS methods.** We adapt several widely used ECTS approaches to the online setting, including an online extension of *Calimera* [6] based on a deep learning architecture. We also introduce (i) four new baselines: two based on multi-armed bandits and two adapting the probability-thresholding paradigm commonly used in classical ECTS and (ii) propose lightweight online adaptations of *Alert* [7] and *Economy* [8], which were originally designed with online decision-making in mind but had not previously been evaluated in this context.
- **Controlled experiments on cost non-stationarity.** We present the first experimental evaluation of ECTS under both drift in the balance between misclassification and delay costs and stochastic cost realizations with controlled distributions, using synthetic data that enables precise control over deployment conditions. This evaluation framework readily extends to real-world applications.
- **Proof of concept for cost-robust online ECTS.** Our results demonstrate that ECTS methods can be effectively adapted to online learning to improve robustness under non-stationary costs. In particular, we observe that RL-based approaches exhibit strong and stable behavior across varying cost regimes, paving the way for future real-world applications.

The remainder of this paper is organized as follows. Section II provides background on ECTS. Section III reviews closely related research areas. Section IV details the proposed online formulation of ECTS and the methods introduced in this paper, including a categorization of the proposed approaches according to different update timings, corresponding to increasing levels of information available about a new instance at update time. Section V describes the experimental protocol, including the synthetic datasets, the controlled cost drift scenarios, and the evaluation metrics. Finally, Section VI presents and discusses the experimental results for the two forms of cost non-stationarity considered in this work.

## II. BACKGROUND

### A. Problem statement of standard ECTS

In the ECTS problem, an input time series of size  $T$  is observed progressively. At time  $t \leq T$ , only a prefix  $\mathbf{x}_t = \langle x_1, \dots, x_t \rangle$  is available, where  $x_{i(1 \leq i \leq t)}$  denotes

the measurements (possibly multivariate). The full series  $\mathbf{x}_T$  belongs to an unknown class  $y \in \mathcal{Y}$ . The task is to output a prediction  $\hat{y} \in \mathcal{Y}$  for the class of the incoming time series at some time  $\hat{t} \in [1, T]$  before the deadline  $T$ .

ECTS explicitly addresses a trade-off between predictive accuracy and earliness. The correctness of the prediction is measured by a misclassification cost  $C_m(\hat{y}|y)$ , and time pressure is modeled by a delay cost  $C_d(t)$ , assumed positive and typically non-decreasing over time. Formally, we consider:  $C_m : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ ,  $C_d : \mathbb{R}^+ \rightarrow \mathbb{R}$ .

An ECTS system involves a predictor  $\hat{y}(\mathbf{x}_t)$  that outputs a class from any prefix  $\mathbf{x}_t$ . The cost incurred when a prediction is triggered at time  $t$  is given by a time-dependent loss  $\mathcal{L}(\hat{y}(\mathbf{x}_t), y, t) = C_m(\hat{y}(\mathbf{x}_t)|y) + C_d(t)$ .

If the full time series  $\mathbf{x}_T$  and its label  $y$  were known in advance, the optimal decision time would satisfy:

$$t^* = \arg \min_{t \in [1, T]} \mathcal{L}(\hat{y}(\mathbf{x}_t), y, t). \quad (1)$$

Let  $s^* \in \mathcal{S}$  be an optimal ECTS function belonging to a class of functions  $\mathcal{S}$ , whose output at time  $t$  when receiving  $\mathbf{x}_t$  is:

$$s^*(\mathbf{x}_t) = \begin{cases} \emptyset & \text{if extra measures are queried;} \\ y^* = \hat{y}(x_{t^*}) & \text{if prediction is triggered at } t^*; \end{cases} \quad (2)$$

In practice, ECTS is an online optimization problem: at each time  $t$ , the function  $s(\mathbf{x}_t)$  must decide whether to predict or wait for more data, based only on the prefix  $\mathbf{x}_t$ . Equation (1) is therefore not directly operational, as it assumes access to the full series. Instead,  $s$  induces a stopping time  $\hat{t}$  on each incoming series, and the goal is to choose  $\hat{t}$  such that the incurred loss  $\mathcal{L}(\hat{y}(\mathbf{x}_{\hat{t}}), y, \hat{t})$  is as small as possible.

From a machine learning perspective, the goal is to find a function  $s \in \mathcal{S}$  that best optimizes the time-dependent loss function  $\mathcal{L}$ , minimizing the true risk over the data-generating process  $\mathbb{P}(\mathcal{X} \times \mathcal{Y})$  [3]:

$$\arg \min_{s \in \mathcal{S}} \mathbb{E}_{(\mathbf{x}, y) \sim \mathbb{P}(\mathcal{X} \times \mathcal{Y})} [\mathcal{L}(\hat{y}(\mathbf{x}_{\hat{t}}), y, \hat{t})] \quad (3)$$

In order to solve the ECTS problem, a training set is made available that is composed of  $N$  labeled time series, denoted by  $(\mathbf{x}_T^i, y^i)_{i \in [0, N]} \in (\mathcal{X} \times \mathcal{Y})$ , where each series  $\mathbf{x}_T = \langle x_1, \dots, x_T \rangle$  is complete and of the same size  $T$ , and associated with its label  $y \in \mathcal{Y}$ .

### B. Literature on standard ECTS

The ECTS literature can be broadly divided into *end-to-end* and *separable* approaches. End-to-end methods jointly learn a classifier and a triggering mechanism. While flexible, they are classifier-dependent and often difficult to compare consistently across studies. In contrast, *separable* approaches decouple classification and decision triggering. This modular design dominates the literature and has been extensively evaluated in previous works [1], [2], [3], which motivates the focus on separable approaches in this paper.

A widely used separable baseline is *Proba-Threshold*, which triggers a prediction when the maximum posterior probability exceeds a learned threshold. Despite its simplicity, this myopic confidence-based approach has been shown to achieve competitive performance across many datasets and cost settings, making it a natural reference method [3].

Beyond myopic strategies, non-myopic separable approaches explicitly anticipate future observations. The *Economy* family [8], [9] triggers a decision when the expected cost at the current time step is minimal among all anticipated future ones, relying on estimates of future costs. *Calimera* [6] follows a similar anticipation-based principle and uses regression models learned by backward induction to predict future costs. *Alert* [7], whose triggering mechanism is based on reinforcement learning, is inherently suited to online adaptation. These methods consistently outperform purely myopic approaches while preserving a separable design [3].

To the best of our knowledge, none of these ECTS approaches have been evaluated under deployment-time cost non-stationarity. In this work, *Proba-Threshold*, *Economy* and *Calimera* are adapted to the online learning setting, as these approaches provide strong performance and are widely used in the standard ECTS regime [3], [6], [7].

### C. Online ECTS with deployment-time cost mismatch

The standard ECTS formulation assumes that the misclassification and delay costs used for training,  $C_m$  and  $C_d$ , accurately reflect the costs incurred at deployment. In practice, this assumption is often violated: the *actual* costs observed in production may differ from the nominal ones and can evolve over time. Such cost variations may occur either within the duration of a single time series (i.e. along the time index  $t \in [0, T]$ ) or across longer time scales, for instance due to seasonal or operational changes.

To clearly distinguish these two temporal dimensions, we explicitly separate the time index of the time series from the deployment timeline. We denote by  $C_m^{(u)}$  and  $C_d^{(u)}$  the *deployment-time* misclassification and delay costs at deployment step  $u \in \mathbb{N}^+$ , where  $u$  indexes successive processed incoming-time series. We define the corresponding loss as follows:  $\mathcal{L}^{(u)}(\hat{y}(\mathbf{x}_t), y, t) = C_m^{(u)}(\hat{y}(\mathbf{x}_t) | y) + C_d^{(u)}(t)$ . In the general case,  $C_m^{(u)}$  and  $C_d^{(u)}$  are (i) *unknown* to the learner at deployment; (ii) may differ from the training-time costs; (iii) and can be *non-stationary* (drifting over time) and/or *stochastic* (varying across instances).

At deployment, an ECTS system repeatedly processes incoming time series. For each new series, it chooses a prediction time  $\hat{t}$  based on the observed prefix  $\mathbf{x}_{\hat{t}}$ , outputs a prediction  $\hat{y}(\mathbf{x}_{\hat{t}})$ , and then incurs a deployment-time loss  $\mathcal{L}^{(u)}(\hat{y}(\mathbf{x}_{\hat{t}}), y, \hat{t})$  once the true label  $y$  and the realized costs are revealed. Its deployment performance is measured by

$$\mathbb{E}[\mathcal{L}^{(u)}(\hat{y}(\mathbf{x}_{\hat{t}}), y, \hat{t})], \quad (4)$$

where the expectation is taken over the data-generating process and, when applicable, over the randomness in the costs.

The purpose of this paper is to allow ECTS systems to *adapt online* to these unknown deployment-time costs, while the classifier has been trained only once under nominal costs  $C_m, C_d$ . Concretely:

- *Offline*, a classifier is trained using complete labeled time series and nominal costs.
- *Online*, during deployment, only the triggering model is updated based on the observed feedback  $\mathcal{L}^{(u)}$ , with the objective of minimizing the expected loss in Equation 4 under cost drift and stochastic costs.

The remainder of the paper instantiates various online adaptation strategies for the most popular ECTS approaches and evaluates their robustness under controlled forms of cost non-stationarity.

## III. RELATED DOMAINS

This section reviews work related to the proposed setting of ECTS under deployment-time cost non-stationarity.

### A. Cost-Sensitive Learning

Learning algorithms that optimize decisions under user-defined cost functions have been extensively studied in cost-sensitive learning [10], [11], [12], typically by incorporating misclassification cost matrices into standard classifiers to minimize expected cost rather than the error rate (e.g., via threshold adjustment, cost-sensitive losses, or tailored ensembles).

In an online setting, cost-sensitive learning has also been explored, e.g., [13] integrate fixed misclassification costs into streaming classifiers; however, they do not address changes in the cost function over time. In contrast, the ECTS setting studied here optimizes an online time-dependent loss, combining misclassification and delay costs, under non-stationary deployment-time conditions.

### B. Concept Drift in Data Streams

Concept drift refers to changes in the joint distribution of inputs and labels over time, and is a central topic in data stream mining [14], [15], [16]. Many works have studied how to detect, characterize, and adapt to drift in streaming environments. Existing approaches include drift detectors based on statistical tests or resampling [14], adaptive ensembles that reweight or replace base learners [17], and prequential evaluation protocols that interleave prediction and update [15].

The scope of the present work is more specific. Rather than addressing changes in the data or label distribution, the focus is on *cost drift*: the underlying data-generating process is assumed to be stationary, while the decision costs governing the loss function may evolve. From this perspective, the proposed setting can be seen as an instance of learning under a fixed concept but evolving evaluation criteria.

### C. Non-Stationary Reinforcement Learning and Bandits

Online decision-making is a long-standing topic of study in artificial intelligence, which has led in particular to Reinforcement Learning (RL) algorithms and Multi-Armed Bandit (MAB) algorithms. In robust RL, the goal is to obtain policies

that perform well under environment variations or distributional shifts, [18], [19], [20].

A closely related line of work studies non-stationary rewards for MAB, with algorithms such as discounted or sliding-window UCB [21] and associated regret bounds in drifting settings [22], which trade off reactivity to change and robustness to noise via forgetting mechanisms.

The present work builds on these ideas by introducing bandit-based trigger models for ECTS, equipped with decay or window mechanisms to handle cost drift. These approaches are evaluated alongside supervised and RL-based methods, within controlled non-stationary cost regimes.

#### IV. BRINGING ECTS TO ONLINE LEARNING

The trigger function may be adaptive, learning from each processed time series and the cost incurred by the corresponding decision. A key aspect of online adaptation is the *update time*, defined as the moment at which feedback becomes available to update the trigger. Depending on the algorithm and the information accessible after a decision, three update regimes can be distinguished: (i) *Delayed updates* correspond to settings where the learning signal is only available at time  $T$ , after the full time series has been observed. This occurs when the outcome associated with a decision can only be assessed once the entire process has unfolded. (ii) *Instant updates* refer to settings where the cost of the decision is immediately observable, allowing the trigger to be updated right after stopping. This regime naturally arises in RL approaches. Finally, (iii) *No updates* denotes methods in which the trigger is not adapted from feedback, but instead relies on explicit cost information specified at deployment time.

##### A. Delayed Updates

In standard ECTS, training typically assumes access to complete time series up to the deadline  $T$ . Several methods explicitly exploit full-horizon information to construct their learning signal. When deployed online, such approaches still require observing the entire time series and realized costs to form new examples to learn from. This update regime is therefore applicable only when the continuation of the series remains observable and is not affected by the trigger decision.

*Proba-Threshold* [3] is a separable baseline that triggers a prediction when the maximum posterior probability,  $p(y|\mathbf{x}_t)$ , exceeds a learned threshold. An online variant maintains a set of candidate thresholds, each characterized by an empirical estimate of the average loss observed when using it. During deployment, decisions are made by selecting the decision threshold that minimizes this estimated average loss. Another variant exploits a decay parameter in order to forget previous observations more quickly: the *decay-Proba-Threshold* variant includes such a mechanism (see Section V). Higher thresholds generally lead to later stopping times, and evaluating their associated loss may require access to the full time series.

*Calimera* [6] is an anticipation-based ECTS method that estimates the difference between cost at *current* time  $t$  and the *best reachable future cost* from the next timestep  $t+1$  until the

end of the time series  $T$ . As soon as this estimate is positive, the decision is triggered, indicating that the minimum cost is being exceeded. It relies on an ensemble of regressors learned by backward induction. We introduce a variant, referred to as *Deep-Calimera*, which replaces the original ensemble of regressors with a single deep regression model, making it more suitable for online updates. This architectural choice has direct implications for target construction. Unlike the original *Calimera* method, where the regression targets are defined iteratively from future model *predictions*, *Deep-Calimera* rather exploits *realized* costs in order to compute its targets (see line 8 of Algorithm 1). This model takes as input five features derived from the classifier output plus the time index  $t$  (i.e., the same as *Alert* [7]). *Deep-Calimera* thus requires access to full time series and associated costs.

---

#### Algorithm 1: *Deep-Calimera* update algorithm

---

**Input:**  $(\mathbf{x}_T, y)$ ,  $\{p(y|\mathbf{x}_t)\}_{(1 \leq t \leq T)}$ ,  $C_m^{(u)}$ ,  $C_d^{(u)}$ , DNN  $f$

- 1  $D \leftarrow \{\}$  \*/
- /\* generate labeled examples
- 2  $\hat{y}_T \leftarrow \arg \max_{k \in \mathcal{Y}} p(k|\mathbf{x}_T)$
- 3  $A_T \leftarrow p(y|\mathbf{x}_T) \cdot C_m^{(u)}(\hat{y}_T|\cdot) + C_d^{(u)}(T)$ ,  $y'_T \leftarrow 0$
- 4 **for**  $t = T - 1$  **to** 1 **do**
- 5    $\hat{y}_t \leftarrow \arg \max_{k \in \mathcal{Y}} p(k|\mathbf{x}_t)$
- 6    $A_t \leftarrow p(y|\mathbf{x}_t) \cdot C_m^{(u)}(\hat{y}_t|\cdot) + C_d^{(u)}(t)$
- 7    $y'_t \leftarrow \mathbb{1}(y'_{t+1} < 0)y'_{t+1} + A_{t+1} - A_t$
- 8    $X'_t \leftarrow \text{generate\_features}(p(y|\mathbf{x}_t), t)$
- 9    $D \leftarrow D \cup \{X'_t, y'_t\}$
- /\* update  $f$  on  $D$  \*/
- 10 minimize  $(\sum_{t=1}^{T-1} \|f(X'_t) - y'_t\|^2)$

---

##### B. Instant Updates

Instant-update approaches adjust the triggering model immediately after a decision, without observing the remaining part of the time series. They are therefore applicable when subsequent measurements are unobserved or influenced by the trigger decision itself. RL-based algorithms are naturally suited to this setting [23].

MAB models have been used as triggering mechanisms in early-exit neural networks [24]. We adopt a similar construction by defining each arm as a candidate confidence threshold, yielding a bandit analogue of *Proba-Threshold*. The *UCBI* algorithm [25] balances exploration and exploitation across thresholds values. For each processed time series, the selected threshold determines the stopping time and the incurred deployment-time loss, which serves as bandit feedback. Another variant, *SW-UCBI* [21], forces the trigger to forget past observations by processing only the ones contained in a sliding-window of fixed size. Furthermore, the triggers training examples are used here to initialize the average arms rewards values, before updating online over deployment examples. Indeed, it has been shown that UCB-type algorithms could effectively benefit from offline history to initialize bandit statistics, as done in the *HUCBI* method [26]. Algorithm 2 describes the *HUCBI* update function.

*Alert* [7] is a DQN-based trigger model with a separable design, trained offline. In this work, an online variant is obtained

---

**Algorithm 2: HUCBI update algorithm**

---

**Input:**  $(\mathbf{x}_T, y)$ ,  $\{p(y|\mathbf{x}_t)\}_{(1 \leq t \leq T)}$ ,  $C_m^{(u)}$ ,  $C_d^{(u)}$ ,  $c \in \mathbb{R}^+$ ,  
arms (decision thresholds)  $K = \{\delta_1, \dots, \delta_I\}$ ,  
per-arms average rewards  $\bar{R} = \{\bar{r}_1, \dots, \bar{r}_I\}$ ,  
per-arms examples counts  $n = \{n_1, \dots, n_I\}$ ,  
number of training examples  $N$

- 1  $u \leftarrow \sum_{i=1}^I n_i$ ,  $\hat{t} \leftarrow T$ ,  $t \leftarrow 1$   
/\* select arm and get trigger time \*/
- 2  $\hat{k} \leftarrow \arg \max_{1 \leq i \leq I} (\bar{r}_i + c \sqrt{\frac{2 \ln(N+u)}{N+n_i}})$
- 3 **while**  $t \leq \hat{t}$  **do**
- 4      $a_t \leftarrow \mathbb{1}(\arg \max_{k \in \mathcal{Y}} p(k|\mathbf{x}_t) > \delta_{\hat{k}})$
- 5     **if**  $a_t = 1$  **then**
- 6          $\hat{t} \leftarrow t$
- 7     **else**
- 8          $t \leftarrow t + 1$
- 9      $\hat{y}_t \leftarrow \arg \max_{k \in \mathcal{Y}} p(k|\mathbf{x}_t)$
- 10     $r_{\hat{k}} \leftarrow \text{get\_reward}(\hat{y}_t, y, \hat{t}, a_t)$
- 11     $n_{\hat{k}} \leftarrow n_{\hat{k}} + 1$ ; // update arm statistics
- 12     $\bar{r}_{\hat{k}} \leftarrow \bar{r}_{\hat{k}} + \frac{r_{\hat{k}} - \bar{r}_{\hat{k}}}{N+u}$

---

by enabling the exploration parameter  $\epsilon$  during deployment and continuing to update the Q-network based on temporal-difference errors (see Algorithm 3). Each episode corresponds to time series processed until a decision at time  $\hat{t}$  or when the deadline  $T$  is reached, and the cumulative time-dependent loss defines the reward signal. This enables the *Alert* to track evolving deployment-time costs using only partial time series.

---

**Algorithm 3: Alert update algorithm**

---

**Input:**  $(\mathbf{x}_T, y)$ ,  $\{p(y|\mathbf{x}_t)\}_{(1 \leq t \leq T)}$ ,  $C_m^{(u)}$ ,  $C_d^{(u)}$ ,  $\epsilon \in [0, 1]$ ,  
 $\gamma \in [0, 1]$ , DQN  $Q$

- 1  $D \leftarrow \{\}$ ,  $\hat{t} \leftarrow T$ ,  $t \leftarrow 1$   
/\* generate transitions \*/
- 2 **while**  $t \leq \hat{t}$  **do**
- 3      $X'_t \leftarrow \text{generate\_features}(p(y|\mathbf{x}_t), t)$
- 4      $a_t \leftarrow \begin{cases} \arg \max_a Q(X'_t, a) \text{ with proba. } 1 - \epsilon \\ \sim \text{Bernoulli}(0.5) \text{ with proba. } \epsilon \end{cases}$
- 5      $\hat{y}_t \leftarrow \arg \max_{k \in \mathcal{Y}} p(k|\mathbf{x}_t)$
- 6      $r_t \leftarrow \text{get\_reward}(\hat{y}_t, y, t, a_t)$
- 7     **if**  $a_t = 1$  **then**
- 8          $\hat{t} \leftarrow t$ ,  $X'_{t+1} \leftarrow \emptyset$
- 9     **else**
- 10          $X'_{t+1} \leftarrow \text{generate\_features}(p(y|\mathbf{x}_{t+1}), t+1)$
- 11          $t \leftarrow t + 1$
- 12      $D \leftarrow D \cup \{X'_t, a_t, r_t, X'_{t+1}\}$   
/\* update  $Q$  on  $D$  \*/
- 13 minimize  $(\sum_{t=1}^{\hat{t}} \|Y_t - Q(X'_t, a_t)\|^2)$
- 14 with  $Y_t = r_t + \gamma \max_a Q(X'_{t+1}, a)$

---

### C. No Updates (Cost-Agnostic Training)

A final category consists of methods that do not adapt their triggering model based on deployment-time feedback, but instead rely on explicit specification of costs at deployment.

*Economy- $\gamma$*  [8] is a non-myopic ECTS method that estimates expected future costs and triggers a decision when the current expected cost is minimal over the remaining time

horizon. The triggering model is learned in a cost-agnostic manner and exploits the misclassification cost matrix and the delay cost function as an input at deployment-time. This separation allows the method to naturally handle cost drift, provided that updated cost information is available at inference time (with a one-step lag in this work).

A key assumption of *Economy- $\gamma$*  is access to detailed cost information at deployment-time, rather than a single aggregated signal such as a realized time-dependent loss. While this enables cost anticipation, it also restricts the range of scenarios in which the method can be directly applied.

**Reproducibility:** all the approaches presented in this paper are available as open-source code at Github.

## V. EXPERIMENTAL PROTOCOL

This section describes the experimental protocol used to evaluate the proposed online ECTS methods under cost non-stationarity. We detail the evaluation metrics, the cost drift scenarios, the data generation process, and the evaluated methods. All these elements are fixed prior to the analysis of results presented in Section VI.

### A. Evaluation Metrics

All online methods are evaluated using an *eval-then-update* (prequential) protocol [27]. At each deployment step  $u$ , a time series and its associated deployment-time costs are first used to evaluate the current model. The observed loss is then used to update the trigger according to its update strategy.

To study the impact of different trade-offs between misclassification and delay costs, we use the standard weighted formulation of the average cost:

$$\text{AvgCost}_\alpha = \frac{1}{N} \sum_{i=1}^N \alpha C_m(\hat{y}_i | y_i) + (1 - \alpha) C_d(\hat{t}_i), \quad (5)$$

where  $\alpha \in [0, 1]$  controls the relative importance of accuracy and earliness. Throughout this work, the same cost functions as [3], [6], [8] are used:  $C_m(\hat{y}|y) = \mathbb{1}(\hat{y} \neq y)$ ,  $C_d(t) = \frac{t}{T}$ .

In the presence of cost non-stationarity, two complementary evaluation metrics are considered. First, we measure the cumulative regret over  $U$  deployment steps,

$$R_U = \sum_{u=1}^U \mathcal{L}^{(u)}(y, \hat{y}, \hat{t}) - \mathcal{L}^{(u)}(y, \hat{y}, t^*),$$

which quantifies how each method adapts to changing costs. Second, we report  $\text{AvgCost}_\alpha$  computed on a separate hold-out set with drifted costs, while keeping the trigger model learned at  $u$  frozen during this evaluation. This latter metric provides an interpretable and robust performance estimate under a fixed data distribution and is applicable to supervised and RL-based approaches. Bandit-based methods, though, are primarily compared using regret-based metrics, as they inherently rely on continued exploration.

## B. Cost Drift Scenarios

To evaluate the robustness of online ECTS methods to non-stationary costs, we consider controlled cost variation scenarios affecting both misclassification and delay costs. Cost non-stationarity occurs across deployment steps  $u$ , while the time index  $t \in [0, T]$  remains associated with the progression within an individual time series.

We first distinguish two temporal patterns of cost variation:

- **Abrupt changes (AC)**, where costs switch suddenly from the training regime to the deployment one.
- **Periodic variations (PV)**, where costs evolve smoothly over deployment and return to the initial training regime.

Independently of their temporal pattern, we consider two types of cost non-stationarity:

- **Drift (D)**, where deployment costs differ from training ones and evolve across deployment steps  $u$ .
- **Stochasticity (S)**, where the cost incurred for each processed time series is a realization of an underlying generative process, potentially varying across deployment steps  $u$ . In this case, the costs used during training represent partial information about this process (e.g., its mode or expectation), and the variability of realized costs is only observed at deployment.

Combining these two dimensions yields four experimental scenarios:

- **AC\_D**: abrupt deterministic cost drift: the  $\alpha$  parameter drops from 0.8 in training to 0.4 during deployment.
- **PV\_D**: periodic deterministic cost drift: the  $\alpha$  parameter goes from 1 to 0.1 and goes back to 1.
- **AC\_S**: abrupt stochastic cost variations: costs associated to a subgroup of classes are sampled from a log-normal distribution with mode equal to 1 and scale equal to 5, clipped between 0 and 500.
- **PV\_S**: same as above with mode equal to 1 and scale parameter evolving from 0.25 to 10 to 0.25 again.

## C. Data Generation

Experiments are conducted on the MNIST-1D synthetic dataset [28], a ten-class classification problem based on univariate time series of fixed length  $T = 40$ . Each generated time series contains a class-specific pattern representing the shape of a digit, embedded at a random position. As the pattern is randomly located within the time series, classification performance improves, on average, as more time steps are observed, which is a necessary condition for ECTS [3]. A total of 20,000 time series are generated using the default parameters. Among them, 25% are used for the training phase (classifier, calibrator, trigger model), 50% are reserved for deployment-time evaluation, and the remaining 25% form the hold-out set.

## D. Evaluated Methods

Our experimental study compares the online ECTS triggering strategies introduced in Section IV, namely *Proba-Threshold*, *decay-Proba-Threshold*, *HUCB1*, *SW-HUCB1*,

*Deep-Calimera*, *Alert*, and *Economy- $\gamma$* . In addition, two static baselines are included to provide reference points for evaluating the benefit of online adaptation: (i) *no\_adapt*: a static *Proba-Threshold* model that does not adapt at deployment, consistently applying the threshold learned during training. (ii) *silver*: a *Proba-Threshold* model retrained from scratch at each deployment step  $u$  using the true deployment-time cost configuration.

## E. Implementation details

The classification module consists of an ensemble of 20 MiniROCKET estimators [29], i.e. one classifier per 5% of the series, followed by a calibration step, based on isotonic regression. During deployment, incoming time series are processed in batches of size 16. The *Deep-Calimera* trigger is implemented as a two-layer fully connected neural network with Leaky ReLU activations and a hidden dimension of size 64. For both *Alert* and *Deep-Calimera*, the deployment-time learning rate is set to  $10^{-3}$ . Unless stated otherwise, all hyperparameters are kept identical to those reported in the original papers. For *Proba-Threshold*, an exponential decay with parameter  $\gamma = 0.01$  is applied to weight recent observations more strongly. For *HUCB1*, we include the *SW-HUCB1* variant [21], which computes bandit statistics over a sliding window of fixed size (set to 1000 in our experiments).

## VI. RESULTS

In this section, we analyze the ECTS online performances under the two types of cost non-stationarity presented in Section V.

### A. Cost balance drift (D)

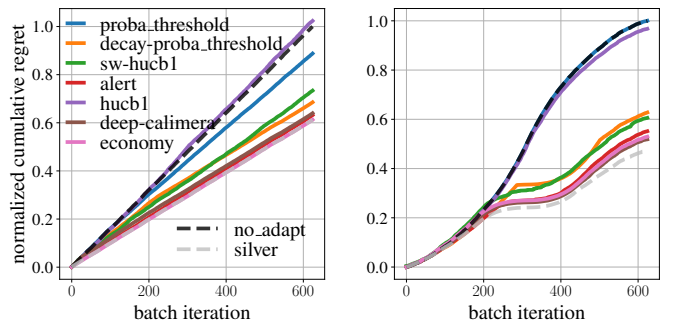


Fig. 1: Normalized cumulated regret.

In this subsection, we evaluate two types of  $\alpha$  drift: (i) abrupt drift (**AC\_D**) and (ii) periodic variations drift (**PV\_D**). The regret analysis reported in Figure 1 shows that most ECTS methods are able to adapt to both drift scenarios, with the notable exceptions of base *Proba-Threshold* and *HUCB1*. These results highlight the importance of forgetting mechanisms for handling non-stationarity: methods incorporating




Fig. 2: **PV\_D**: Hold-out *AvgCost*, decomposed over both earliness and error rate axis. The thick *AvgCost* line corresponds to the  $\alpha$  weighted sum of earliness and error rate. Purple curve describe the evolution of the  $\alpha$  parameter.

such mechanisms, namely *decay-Proba-Threshold* and *SW-HUCB1*, substantially outperform their full-memory counterparts. Among the approaches closest to the *silver* baseline in terms of cumulative regret, which is not applicable in practice, *Economy*, *Alert*, and *Deep-Calimera* achieve the best performance. In particular, *Economy* performs especially well in this setting, where cost matrices exhibit limited variation between deployment steps. Finally, the strong performance of *Alert* is particularly encouraging, as being an *instant-updates* method, it relies on less information than *Deep-Calimera*, which delays its updates.

Figure 2 shows the decomposition of the hold-out *AvgCost* into its temporal (earliness) and misclassification (error rate) components for the **PV\_D** experiment. As  $\alpha$  decreases, corresponding to increasing temporal pressure, most methods adapt their triggering behavior, as evidenced by the downward trend of the earliness (light-blue) curve, indicating that decisions are made earlier, on average. This shift is accompanied by an increase in the error rate, reflecting the higher uncertainty associated with earlier predictions. When  $\alpha$  is restored to its initial value, most methods revert to their original operating regime; however, *decay-Proba-Threshold* exhibits a slower recovery and requires additional time to return to its initial behavior, as its memory of past events gives it inertia.

### B. Stochastic costs (S)

The second form of cost non-stationarity considered in this work concerns stochastic cost realizations. This experiment evaluates the robustness of ECTS methods to noise and their ability to adapt under stochastic perturbations. The regret curves reported in Figure 3 show that, for both **AC\_S** and **PV\_S** scenarios, threshold-based methods achieve comparable performance levels, similar to that of *Economy*. However, these approaches are consistently outperformed by non-myopic models, namely *Alert* and *Calimera*. While *Deep-Calimera* attains lower cumulative regret than *Alert*, this improvement comes at the cost of a more restrictive setting, as it requires

access to the complete time series for updates. In contrast, *Alert* can update its model online and still demonstrates strong performance. Although *Economy* is also non-myopic, its performance degrades in this scenario due to the lack of temporal continuity in the cost values, which are independently drawn at each iteration. Consequently, two consecutive cost functions may differ substantially, a setting for which *Economy* is not designed. Finally, bandit-based methods are adversely affected by the iterative reward normalization based on the current sampled maximum cost, which is required to satisfy boundedness assumptions but results in non-homogeneous updates.

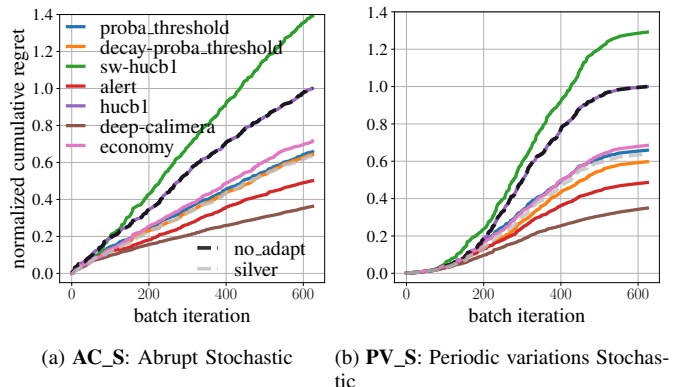


Fig. 3: Normalized cumulated regret. Costs of classes  $\{1, 4, 7\}$  (straight digits) are stochastic. Similar results were obtained with others noised classes. Averaged over five random seeds.

### C. Computation time

As ECTS methods are intended for real-time deployment, the runtimes of both inference and trigger update procedures are reported in Table I, measured under the same experimental conditions as in **PV\_S**. As expected, *Alert* and *Deep-Calimera*

incur higher computational overhead, whereas bandit-based and threshold-based methods exhibit lower runtimes.

Trigger	Inference	Update	Total ↓
<i>HUCBI</i>	0.079	0.374	0.453
<i>Economy</i>	0.490	0	0.490
<i>SW-HUCBI</i>	0.119	0.970	1.089
<i>(decay-)Proba-Threshold</i>	0.076	2.546	2.622
<i>Alert</i>	0.780	5.657	6.437
<i>Deep-Calimera</i>	0.707	8.026	8.733

TABLE I: Mean inference and update time per batch (ms).

## VII. CONCLUSION

This paper addresses ECTS under deployment-time cost non-stationarity, a practically relevant setting that is largely overlooked in existing work. Beyond the specific ECTS setting, this work contributes to the broader problem of early and adaptive decision-making under evolving cost constraints.

We investigated two complementary sources of non-stationarity: deterministic drift in misclassification and delay costs, and stochastic deployment-time cost realizations. Focusing on separable ECTS approaches, we proposed online adaptations of representative triggering strategies that update only the trigger during deployment.

Our contributions include online extensions of several anticipation, bandit, and RL-based triggers, together with controlled experimental protocols explicitly modeling cost drift and stochasticity. In particular, RL-based triggers show strong and stable performance across all scenarios, while non-myopic methods such as *Deep-Calimera* achieve lower costs when delayed updates are feasible, thus at the expense of more restrictive applicability.

This work opens several directions for future research, including evaluation on real-world datasets, improving end-to-end ECTS methods to make it competitive, and jointly addressing cost non-stationarity, data distribution shifts and concept drifts. To facilitate further research, all experiments are fully reproducible and are released as open source. We believe that explicitly accounting for cost non-stationarity is a necessary step toward deploying reliable and adaptive neural decision systems in real-world, long-running applications.

## REFERENCES

- [1] A. Gupta, H. P. Gupta, B. Biswas, and T. Dutta, "Approaches and applications of early classification of time series: A review," *IEEE Transactions on Artificial Intelligence*, vol. 1, no. 1, pp. 47–61, 2020.
- [2] C. Akasiadis, E. Kladis, P.-F. Kamberi, E. Michelioudakis, E. Alevizos, and A. Artikis, "A framework to evaluate early time-series classification algorithms." in *EDBT*, 2024, pp. 623–635.
- [3] A. Renault, A. Bondu, A. Cornuéjols, and V. Lemaire, "Early classification of time series: A survey and benchmark," *Transactions on Machine Learning Research*.
- [4] A. Bondu, Y. Achenchabe, A. Bifet, F. Clérot, A. Cornuéjols, J. Gama, G. Hébrail, V. Lemaire, and P.-F. Marteau, "Open challenges for machine learning based early decision-making research," *ACM SIGKDD Explorations Newsletter*, vol. 24, no. 2, pp. 12–31, 2022.
- [5] L. M. Fleuren, T. L. T. Klausch, C. L. Zwager, L. J. Schoonmade, T. Guo, L. F. Roggeveen, E. L. Swart, A. R. J. Girbes, P. Thorat, A. Ercole *et al.*, "Machine learning for the prediction of sepsis: A systematic review and meta-analysis of diagnostic test accuracy," *Intensive Care Medicine*, vol. 46, no. 3, pp. 383–400, 2020.
- [6] J. M. Bilski and A. Jastrzebska, "Calimera: A new early time series classification method," *Information Processing & Management*, vol. 60, no. 5, p. 103465, 2023.
- [7] A. Renault, A. Bondu, A. Cornuéjols, and V. Lemaire, "Deep reinforcement learning based triggering function for early classifiers of time series," *arXiv preprint arXiv:2502.06584*, 2025.
- [8] Y. Achenchabe, A. Bondu, A. Cornuéjols, and A. Dachraoui, "Early classification of time series: Cost-based optimization criterion and algorithms," *Machine Learning*, vol. 110, no. 6, pp. 1481–1504, 2021.
- [9] A. Dachraoui, A. Bondu, and A. Cornuéjols, "Early classification of time series as a non myopic sequential decision making problem," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2015, Porto, Portugal, September 7-11, 2015, Proceedings, Part I 15*. Springer, 2015, pp. 433–447.
- [10] C. Elkan, "The foundations of cost-sensitive learning," in *International joint conference on artificial intelligence*, vol. 17, no. 1. Lawrence Erlbaum Associates Ltd, 2001, pp. 973–978.
- [11] V. S. Sheng and C. X. Ling, "Thresholding for making classifiers cost-sensitive," in *Aaai*, vol. 6, 2006, pp. 476–481.
- [12] G. Petrides and W. Verbeke, "Cost-sensitive ensemble learning: a unifying framework," *Data Mining and Knowledge Discovery*, vol. 36, no. 1, pp. 1–28, 2022.
- [13] J. Wang, P. Zhao, and S. C. Hoi, "Cost-sensitive online classification," *IEEE Transactions on Knowledge and Data Engineering*, vol. 26, no. 10, pp. 2425–2438, 2013.
- [14] M. Harel, S. Mannor, R. El-Yaniv, and K. Crammer, "Concept drift detection through resampling," in *International conference on machine learning*. PMLR, 2014, pp. 1009–1017.
- [15] J. Lu, A. Liu, F. Dong, F. Gu, J. Gama, and G. Zhang, "Learning under concept drift: A review," *IEEE transactions on knowledge and data engineering*, vol. 31, no. 12, pp. 2346–2363, 2018.
- [16] A. S. Iwashita and J. P. Papa, "An overview on concept drift learning," *IEEE access*, vol. 7, pp. 1532–1547, 2018.
- [17] G. Jaber, A. Cornuéjols, and P. Tarroux, "A new on-line learning method for coping with recurring concepts: The adacc system," in *International conference on neural information processing*. Springer, 2013, pp. 595–604.
- [18] S. Padakandla, P. KJ, and S. Bhatnagar, "Reinforcement learning algorithm for non-stationary environments," *Applied Intelligence*, vol. 50, no. 11, pp. 3590–3606, 2020.
- [19] J. Morimoto and K. Doya, "Robust reinforcement learning," *Neural computation*, vol. 17, no. 2, pp. 335–359, 2005.
- [20] J. Moos, K. Hansel, H. Abdulsamad, S. Stark, D. Clever, and J. Peters, "Robust reinforcement learning: A review of foundations and recent advances," *Machine Learning and Knowledge Extraction*, vol. 4, no. 1, pp. 276–315, 2022.
- [21] A. Garivier and E. Moulines, "On upper-confidence bound policies for non-stationary bandit problems," *arXiv preprint arXiv:0805.3415*, 2008.
- [22] O. Besbes, Y. Gur, and A. Zeevi, "Stochastic multi-armed-bandit problem with non-stationary rewards," *Advances in neural information processing systems*, vol. 27, 2014.
- [23] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [24] D. J. Bajpai and M. K. Hanawal, "Beyond greedy exits: Improved early exit decisions for risk control and reliability," in *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.
- [25] P. Auer, N. Cesa-Bianchi, and P. Fischer, "Finite-time analysis of the multiarmed bandit problem," *Machine learning*, vol. 47, no. 2, pp. 235–256, 2002.
- [26] P. Shivaswamy and T. Joachims, "Multi-armed bandit problems with history," in *Artificial intelligence and statistics*. PMLR, 2012, pp. 1046–1054.
- [27] J. a. Gama, R. Sebastião, and P. P. Rodrigues, "Issues in evaluation of stream learning algorithms," in *SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '09. Association for Computing Machinery, 2009, p. 329–338.
- [28] S. Greydanus and D. Kobak, "Scaling down deep learning with mnist-1d," *arXiv preprint arXiv:2011.14439*, 2020.
- [29] A. Dempster, D. F. Schmidt, and G. I. Webb, "Minirocket: A very fast (almost) deterministic transform for time series classification," in *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 2021, pp. 248–257.