

A Scalable and Efficient Iterative Method for Copying Machine Learning Classifiers

Nahuel Statuto

NAHUEL.STATUTO@ESADE.EDU

*Department of Operations, Innovation and Data Sciences
Universitat Ramon Llull, ESADE
Sant Cugat del Vallès, 08172, Catalonia, Spain*

Irene Unceta

IRENE.UNCETA@ESADE.EDU

*Department of Operations, Innovation and Data Sciences
Universitat Ramon Llull, ESADE
Sant Cugat del Vallès, 08172, Catalonia, Spain*

Jordi Nin

JORDI.NIN@ESADE.EDU

*Department of Operations, Innovation and Data Sciences
Universitat Ramon Llull, ESADE
Sant Cugat del Vallès, 08172, Catalonia, Spain*

Oriol Pujol

ORIOL_PUJOL@UB.EDU

*Departament de Matemàtiques i Informàtica
Universitat de Barcelona
Barcelona, 08007, Catalonia, Spain*

Editor: Maya Gupta

Abstract

Differential replication through copying refers to the process of replicating the decision behavior of a machine learning model using another model that possesses enhanced features and attributes. This process is relevant when external constraints limit the performance of an industrial predictive system. Under such circumstances, copying enables the retention of original prediction capabilities while adapting to new demands. Previous research has focused on the single-pass implementation for copying. This paper introduces a novel sequential approach that significantly reduces the amount of computational resources needed to train or maintain a copy, leading to reduced maintenance costs for companies using machine learning models in production. The effectiveness of the sequential approach is demonstrated through experiments with synthetic and real-world datasets, showing significant reductions in time and resources, while maintaining or improving accuracy.

Keywords: Sustainable AI, transfer learning, environmental adaptation, optimization, and model enhancement.

1. Introduction

Machine learning has become widespread in many industries. Supervised algorithms have been shown to automate tasks with a higher precision and lower costs when compared to previous methods (Gómez et al., 2018; Gharibshah and Zhu, 2021; Shehab et al., 2022; Feizabadi, 2022).

However, maintaining the performance of industrial machine learning models requires constant monitoring. The environment where these models are deployed is subject to continuous change due to factors that are both internal and external to the companies implementing them. These factors may include new production needs, technological updates, novel market trends, or regulatory amendments. Neglecting these changes can result in model degradation and decreased performance. In the worst-case scenario, failure to adapt to these changes can render a model obsolete. To prevent this, regular model monitoring is essential in any industrial machine learning application. Once deployed in production, models are frequently checked for signs of performance deviation, which can occur as early as a few months after deployment. When such signs appear, models are either fully or partially retrained and substituted (Wu et al., 2020). Unfortunately, this process is both time-consuming and costly. The nature of modern model architectures is complex. Hence, the task of managing and updating an existing model demands significant computational resources (Chen, 2019). Handling multiple models simultaneously poses a significant challenge for companies. Thus, ensuring long-term sustainability remains one of the main obstacles faced by industrial machine learning practitioners today (Paley et al., 2022).

To address this, differential replication through copying (Unceta et al., 2020a) has been proposed as an efficient and effective solution for adapting trained machine learning models. This approach builds upon earlier concepts on knowledge distillation (Hinton et al., 2015b; Buciluă et al., 2006b), and enables reuse of a model’s acquired knowledge to train a new one tailored to the specific environmental requirements (Unceta et al., 2020b). It can, therefore, bring numerous benefits in terms of cost-effectiveness and resource optimization.

Differential replication allows for environmental adaptation by projecting an existing decision function onto a new hypothesis space that is compatible with the new requirements. Typically, this process involves using the class probabilities produced by the given decision function as soft targets to train a new model in the new hypothesis space (Liu et al., 2018; Wang et al., 2020). In the case of differential replication through copying, information about the original model’s behavior is acquired via a hard-membership query interface, and training is done using synthetic samples labeled by the original model.

Theoretically, the copying problem can be viewed as a dual optimization problem where both the copy model parameters and the synthetic samples are optimized simultaneously. Previous practical implementations of this problem have simplified it, generating and labeling a large set of synthetic data points in a single step, and then using them to optimize the copy model parameters. This approach has been successful in validating copying on several datasets. However, it is memory-intensive, computationally expensive and requires pre-setting several hyperparameters.

This article introduces a novel iterative approach to differential replication through copying. The goal of copying is to find the simplest model in the copy model hypothesis space that attains the maximum fidelity to the target model being copied. In situations where the original training data are unavailable, this requires finding the best synthetic set that accurately represents the original decision boundary and the simplest model that guarantees perfect fidelity within the capacity constraints of the copy model space. The proposed iterative scheme achieves this by performing two steps at each iteration: (1) generating and selecting synthetic data for copying based on a compression measure of uncertainty, and (2) learning the target copy model using the optimized synthetic dataset while controlling

its complexity to achieve close-to-optimal results. This process allows for control over the amount of data/memory needed and the convergence speed to reach a steady performance. Results show that the proposed model requires 85 – 93% less data samples/memory and has an average 80% improvement in convergence speed when compared to the previous single-pass approach, with no significant degradation in performance.

The main contributions of this article are: (1) to the best of our knowledge, this is the first algorithm to specifically address the problem of copying as described in (Unceta et al., 2020a); (2) the proposed formulation and algorithms allow for explicit control of memory requirements and convergence speed while maintaining accuracy; and (3) the resulting algorithm is accurate, fast, and memory-efficient, as validated by successful results. This article presents an algorithm that includes two hyperparameters. It also proposes a version of this algorithm that automatically sets one of them and dynamically adapts it during the learning process, ensuring fast convergence to an accurate solution. The final algorithm overcomes some of the limitations of current copying methods and opens up opportunities for its use in various real-life applications.

The rest of this paper is organized as follows: Section 2 introduces differential replication through copying, providing an overview of relevant methods, and describes the single-pass approach as its simplest implementation. Section 3 presents the sequential approach. It starts by demonstrating the convergence of this approach to the optimal solution. It then introduces a sample removal policy based on an uncertainty measure. It ends by proposing a regularization term to prevent model forgetting. Section 4 empirically demonstrates the validity of the sequential approach through experiments on a large set of 58 UCI problems. The performance of the sequential copying framework is measured in terms of accuracy, convergence speed, and efficiency. The section ends with a discussion of the results. Finally, Section 5 summarizes our findings and outlines future research directions.

2. Background on copying

The problem of *environmental adaptation* introduced by (Unceta et al., 2020b) describes scenarios where a machine learning model that was designed under a set of constraints must fulfill new requirements imposed by changes in its environment. Specifically, it refers to situations where a model needs to adapt from a *source scenario*, s , where it was trained, to a *target scenario*, t , where it is being deployed.

2.1 The problem of environmental adaptation

Formally speaking, environmental adaptation is defined as follows: consider a task \mathcal{T} and a domain \mathcal{D} . A trained model $h \in \mathcal{H}_s$ is designed to solve \mathcal{T} in \mathcal{D} under a set of constraints \mathcal{C}_s and a compatible hypothesis space \mathcal{H}_s . The problem of environmental adaptation arises when the original set of constraints \mathcal{C}_s evolves to a new set \mathcal{C}_t . Under these circumstances, a potentially different hypothesis space \mathcal{H}_t has to be defined in the same domain \mathcal{D} and for the same task \mathcal{T} , such that it is compatible with the new constraints set \mathcal{C}_t ¹. Unless the considered model h is compatible with the new set of constraints, it will no longer be

1. In some cases the source and target hypothesis spaces \mathcal{H}_s and \mathcal{H}_t may be the same. However, in the most general case, where the new set of constraints defines a new set of feasible solutions, they are not.

| | | | |
|---|-----------------|---|---|
| Source Scenario for \mathcal{T} in \mathcal{D} | | → | Target Scenario for \mathcal{T} in \mathcal{D} |
| maximize for $h \in \mathcal{H}_s$ | $P(y x; h)$ | | maximize for $h \in \mathcal{H}_t$ |
| subject to | \mathcal{C}_s | | subject to |
| | | | $P(y x; h)$ |
| | | | \mathcal{C}_t |

Table 1: The problem of environmental adaptation looks for the solution in the same domain and task when the feasibility constraints change.

a feasible solution. Hence, environmental adaptation refers to the need to adapt h to the constraints introduced by \mathcal{C}_t , as shown in Table 1.

This problem is different from *domain adaptation* and *transfer learning* (Chen and Bühlmann, 2021; Chen et al., 2022; Raffel et al., 2020). Domain adaptation refers to adapting a model from one source domain \mathcal{D}_s to a related target domain \mathcal{D}_t due to a change in the data distributions. Environmental adaptation preserves the domain, changing the constraints instead. Transfer learning requires reusing knowledge from one task \mathcal{T}_s to solve a related task \mathcal{T}_t (Pan and Yang, 2010). Environmental adaptation preserves the task.

Environmental adaptation can be addressed through various methods, including re-training the existing model with a new set of constraints (Barque et al., 2018), using wrappers (Mena et al., 2019, 2020), teacher-student networks and distillation mechanisms (Buciluă et al., 2006a; Hinton et al., 2015a; Szegedy et al., 2016; Yang et al., 2019), label regularization (Müller et al., 2019; Yuan et al., 2020), label refinement (Bagherinezhad et al., 2018), or edited, augmented (Song et al., 2008; Chen et al., 2020; Duan et al., 2018), or synthetic data subsets (Buciluă et al., 2006a; Zeng and Martinez, 2000). A comprehensive overview of all these methods is available in (Unceta et al., 2020b). Here, we focus in differential replication, a general solution to the environmental adaptation problem.

Differential replication is the process whereby the decision boundary of an existing model is projected to a new hypothesis space that is compatible with the target scenario. In cases where there is no access to the original training dataset or the trained model internals, differential replication through copying can be used to solve the environmental adaptation problem Unceta et al. (2020a). In classification settings, copying involves obtaining a new classifier that displays the same performance and decision behavior as the original, without necessarily belonging to the same model family.

In the following sections, we introduce the problem of differential replication through copying and explore potential approaches for its practical implementation.

2.2 Differential replication through copying

Consider a classifier $f_{\mathcal{O}} \in \mathcal{H}_s$ trained on an unknown dataset with input space dimensionality d and output space cardinality n_c . Thus, $f_{\mathcal{O}} : \mathbb{R}^d \rightarrow \{0, 1\}^{n_c}$. In the most restrictive case, $f_{\mathcal{O}}$ is a hard decision classifier that outputs one-hot encoded label predictions, meaning that for any data point, it returns an n_c -output vector with all elements as zeros except for the target label position i , which has a value of 1.

The goal of copying is to obtain a new classifier $f_C \in \mathcal{H}_t$ parameterized by $\theta \in \Theta$ that mimics f_O across the sample space. In the empirical risk minimization framework, we can consider an empirical risk function that measures the discrepancy between two classifiers. We can then formulate the copying problem as a dual optimization of θ and a set of synthetic data points S over which the empirical risk is evaluated, since we do not have access to any training data. This problem can be written as follows:

$$\begin{aligned} & \underset{\theta, S}{\text{minimize}} && \Omega(\theta) \\ & \text{subject to} && \|R_{emp}^{\mathcal{F}}(f_C(\theta), f_O) - R_{emp}^{\mathcal{F}}(f_C(\theta^\dagger), f_O)\| < \varepsilon, \end{aligned} \tag{1}$$

for a defined tolerance ε and a complexity measure $\Omega(\theta)$ (e.g., the ℓ_p -norm of the parameters). The empirical risk function $R_{emp}^{\mathcal{F}}$ measures the difference between the original model f_O and the optimized copy model f_C and is referred to as the *empirical fidelity error*. θ^\dagger is the solution to the following unconstrained optimization problem:

$$\theta^\dagger = \underset{\theta, S}{\operatorname{argmin}} R_{emp}^{\mathcal{F}}(f_C(\theta), f_O). \tag{2}$$

The solution to Equation 1 is a combination of synthetic data and copy model parameters that minimize capacity and empirical risk. If $f_O \in \mathcal{H}_t$, the solution of the unconstrained problem (Equation 2) will always result in $R_{emp}^{\mathcal{F}}(f_C(\theta^\dagger), f_O) = 0$, as the labeling problem for any set of data points using the original hard-label classifier is separable in \mathcal{H}_t .

2.3 The single-pass approach

The *single-pass* approach (Unceta et al., 2020b) aims to find a sub-optimal solution to the copying problem by dividing it into two separate steps. Firstly, a synthetic set S^* is found and then, the copy parameters θ^* are optimized using this set. The process is as follows:

- **Step 1: Synthetic sample generation.** An exhaustive synthetic set S^* is created by randomly sampling from a probability density function P_S that covers the operational space of the copy. The operational space is the region of the input space where the copy is expected to mimic the behavior of the original model. The synthetic set can be expressed as:

$$S^* = \{z_j | z_j \sim P_S, j = 1, \dots, N\}. \tag{3}$$

The simplest choice for P_S is a uniform distribution, or a normal or Gaussian distribution if the original data is normalized. (See (Unceta et al., 2020c) for more information on different options for P_S .)

- **Step 2: Building the copy.** The optimal parameter set for the copy is obtained by minimizing the empirical risk of the copied model over the synthetic set S^* obtained in Step 1:

$$\theta^* = \underset{\theta}{\operatorname{argmin}} R_{emp}^{\mathcal{F}}(f_C(\theta), f_O) \Big|_{S=S^*}. \tag{4}$$

The single-pass approach is a simplified solution for the problem modeled in Equation 1. Step 1 generates data without any optimization, but it requires a large dataset. Step 2 focuses on solving the unconstrained version of the copying problem defined in Equation 2. This approach can be used when the classifier complexity can be directly modeled. However, in the general case, it requires setting many critical parameters and selecting an appropriate model to ensure that the unconstrained problem is a good approximation of the general setting described in Equation 1. To guarantee good performance, a sufficiently large synthetic dataset must be generated.

The implementation of the single-pass approach has limitations. Firstly, learning using a one-shot approach with a single model may be limited by the available memory if unable to handle the full dataset. Secondly, keeping a large number of data samples in memory during training is resource-intensive and doesn't guarantee performance. In addition, blindly learning the entire synthetic dataset using a single model can result in inefficiencies. On the other hand, using an online learning strategy can reduce memory usage, but leads to slow convergence to the optimal solution, making the process time-consuming.

To overcome these limitations, a new approach, using an alternating optimization scheme, is introduced. This approach provides a fast and memory-efficient solution to the copying problem, solving Equations 1 and 2.

3. The sequential approach

In this section, we introduce two theorems to show that the sequential approach converges to the single-pass approach when conditions are optimal in terms of both parameters and behavior. Next, we provide a practical implementation of the sequential approach and various optimizations to achieve low memory usage and fast convergence. Specifically, we propose that a perfect copy should be able to compress the synthetic data points in its parameters and suggest epistemic uncertainty as a reliable metric for data compression. Based on this, we develop a data selection strategy that filters samples based on their level of compression by the copy model. This leads to a reduced number of data points needed for each learning step. Finally, we introduce an automatic hyper-parameter tuning policy to ensure optimal implementation of the sequential approach in practice. We refer to the resulting algorithm as the *sequential approach* to copying.

3.1 An alternating optimization algorithm for copying

We start by introducing a preliminary alternating optimization algorithm for solving Equation 1. This algorithm alternates between two optimization steps at each iteration t :

- **Step 1: Sample Optimization.** The optimal synthetic set at iteration t is selected by maximizing the empirical fidelity error for the previous model solution θ_{t-1}^* . That is,

$$S_t^* = \arg \max_S R_{emp}^{\mathcal{F}, S}(f_C(\theta), f_O) \Big|_{\theta = \theta_{t-1}^*} \quad (5)$$

- **Step 2: Copy Parameter Optimization.** The optimal copy parameters θ_t^* at iteration t over samples S_t^* are obtained by:

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \Omega(\theta) && (6) \\ & \text{subject to} && \left\| R_{emp}^{\mathcal{F},S}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}) - R_{emp}^{\mathcal{F},S}(f_{\mathcal{C}}(\theta^\dagger), f_{\mathcal{O}}) \right\|_{S=S_t^*} && < \varepsilon. \end{aligned}$$

The algorithm starts each iteration t by selecting a set of synthetic data points that maximize the empirical risk. These are the points that the copy model from the previous iteration $t - 1$ did not model correctly. By reducing the empirical risk to zero, we can minimize the loss function. In Step 2, we minimize the empirical loss over S_t while keeping the copy model complexity as low as possible. The rest of this section focuses on solving Equations 5 and 6 under various assumptions.

3.2 Step 1: Sample optimization

We start by introducing a formal sequential sample generation scheme and examining its convergence properties. This will serve as the basis for constructing the final algorithm that solves Equation 5. To achieve this, we prove two theorems showing that the sequential copying process converges in both parameters and behavior to the single-pass approach under optimal conditions.

3.2.1 THE SEQUENTIAL FRAMEWORK

Consider a sequence of finite sets S_t such that

$$S_t \subseteq S_{t+1} \subseteq \dots \subseteq S \tag{7}$$

for $t \in \mathbb{N}$, where $|S| = \aleph^0$, which approaches the set S as t increases towards infinity. The sequential framework is based on the notion that, if S_t converges to S , then the optimal copy parameters θ_t^* obtained by optimizing over S_t will converge to θ^* , the optimal parameters over S . This approximation can be iteratively obtained by drawing samples from a given probability density function.

To prove this, we cast the unconstrained copying problem in terms of the loss function and express it as the solution to the following empirical problem:

$$\theta^* = \arg \min_{\theta} R_{emp}^{\mathcal{F},S}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}) = \arg \min_{\theta} \overbrace{\frac{1}{|S|} \sum_{z \in S} \ell(f_{\mathcal{O}}(z), f_{\mathcal{C}}(z; \theta))}^{F(\theta):=}, \tag{8}$$

where $S = \{z | z \sim P_S\}$ is a synthetic dataset of size $|S| = \aleph^0$ and $\ell(\cdot)$ is a non-negative and bounded loss function.

Next, we present Theorem 1, which shows that the solution to Equation 8 for S can be approximated using the sequence of iterative values S_t .

Algorithm 1 Preliminary version of the sequential approach

1: **Input:** `int` T , `int` n , **Classifier** $f_{\mathcal{O}}$
2: **Output:** Optimal copy parameters
3: $S_0 \leftarrow \{(z_j, f_{\mathcal{O}}(z_j)) \mid z_j \sim P_{\mathcal{S}}, j = 1 \dots n\}$
4: $\theta_0^* \leftarrow \arg \min_{\theta} \frac{1}{|S_0|} \sum_{z \in S_0} \ell(f_{\mathcal{O}}(z), f_{\mathcal{C}}(z; \theta))$
5: **for** $t = 1$ to T **do**
6: $S_t \leftarrow S_{t-1} \cup \{(z_j, f_{\mathcal{O}}(z_j)) \mid z_j \sim P_{\mathcal{S}}, j = 1 \dots n\}$
7: $\theta_t^* \leftarrow \arg \min_{\theta} \frac{1}{|S_t|} \sum_{z \in S_t} \ell(f_{\mathcal{O}}(z), f_{\mathcal{C}}(z; \theta)) \Big|_{\theta^0 = \theta_{t-1}^*}$
8: **end for**
9: **return** θ_t^*

Theorem 1 Let $S_i \subseteq S_{i+1} \subseteq \dots \subseteq S$ be a sequence of subsets converging to S with $|S_i| = i \cdot M$, $M \in \mathbb{N}$. Then, the sequence of functions $\{F_i\}_i$ defined as $F_i(\theta) = \frac{1}{|S_i|} \sum_{z \in S_i} \ell(f_{\mathcal{C}}(z, \theta), f_{\mathcal{O}}(z))$, converges uniformly to $F(\theta) = \frac{1}{|S|} \sum_{z \in S} \ell(f_{\mathcal{C}}(z, \theta), f_{\mathcal{O}}(z))$, where $\ell(\cdot)$ is a non-negative and bounded loss function, $|\ell(\cdot)| \leq \mathcal{K}$.

Theorem 2 proves the convergence of parameters in Theorem 1.

Theorem 2 Under the conditions of Theorem 1, the sequence of parameters $\{\theta_t^*\}_t$ defined as $\theta_t^* = \arg \min_{\theta \in \Theta} F_i(\theta)$, converges to $\theta^* = \arg \min_{\theta \in \Theta} F(\theta)$, where Θ is the complete set of parameters.

The full mathematical proof of convergence of the copy parameter sequence can be found in Appendix A.

Definition: The *sequential approach* learning algorithm is an iterative process that optimizes the copy parameters θ_t incrementally over sets of synthetic data points S_t , $t = 1 \dots T$.

The preliminary version of the sequential approach is outlined in Algorithm 1. We begin by generating an initial synthetic set S_0 of size n and optimizing the copy parameters accordingly. As stated in line 6, at each iteration t the new synthetic set S_t adds n samples to the previous set S_{t-1} . The cardinality of S_t is therefore given by $|S_t| = t \cdot n$ and grows linearly with t . This linear growth policy was chosen as the simplest strategy² to build the subsets in Equation 7. Other strategies can also be used. Line 7 describes the optimization of the algorithm considering the solution of the previous step.

For demonstration, we test the sequential approach on a toy binary classification dataset, the *spirals* problem. We first train a Gaussian kernel SVM on this dataset, achieving perfect accuracy. Then, we copy the learned boundary using a fully-connected neural network with three hidden ReLU layers of 64, 32, and 10 neurons, and a SoftMax output. We train it for 1000 epochs with a batch size of 32 samples. Figure 1a) shows the original decision boundary and data points.

2. check Theorem 1's proof and comments to see other growth policies that still ensure convergence.

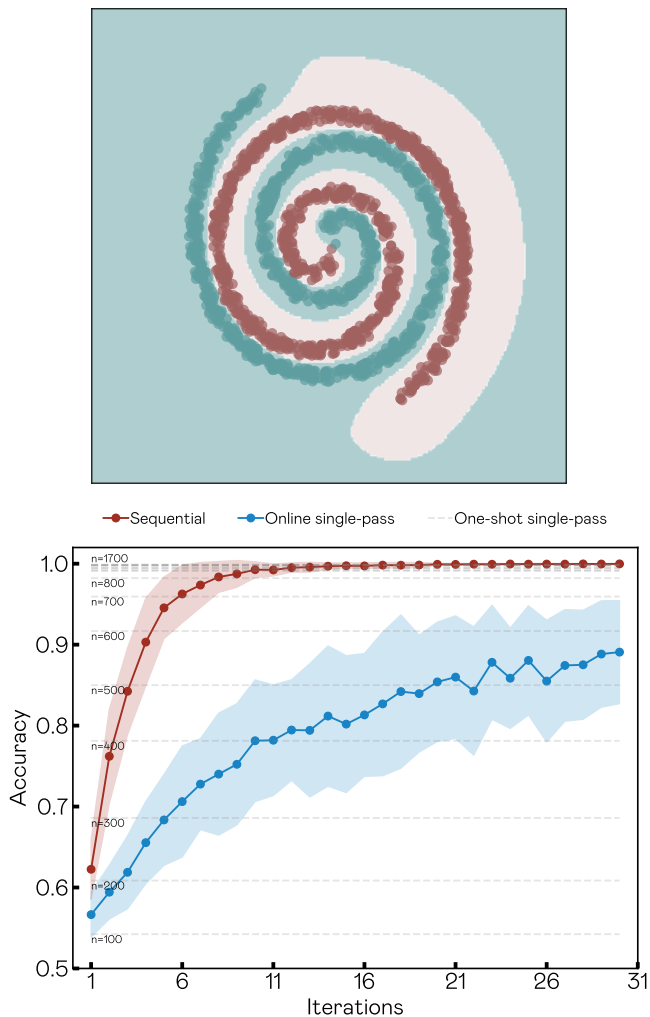


Figure 1: a) Decision boundary learned by a Gaussian kernel SVM trained on the *spirals* dataset. b) Copy accuracy averaged over 30 independent runs for the sequential approach (orange) and the online single-pass approach (blue) with $n = 100$. The grey dashed lines depict the accuracy obtained using a one-shot single-pass approach for various values of n . The shaded areas indicate a standard deviation from the mean.

We compare the results of Algorithm 1 with those of the single-pass approach in both its unique-step and online implementations. Figure 1b) shows the average accuracy computed at different iterations for copies trained using different methods. The online and sequential approaches are trained by generating $n = 100$ new samples at each iteration. The grey dashed lines represent the accuracy achieved by single-pass copies trained in a one-shot for different values of n . The reported values represent accuracy measured over the original test data points.

The plot highlights differences in convergence speeds among the three approaches. The one-shot single-pass approach with $n = 500$ has an accuracy of 0.85, while the purely online single-pass model reaches an accuracy of around 0.85 after 30 iterations and 3000 samples. In comparison, the sequential approach has an accuracy of 0.94 at $t = 5$ with 500 samples, indicating faster convergence. However, both the single-pass and sequential approaches reach an accuracy of ≈ 1 after being exposed to 1000 samples.

| | One-shot single-pass | Online single-pass | Sequential |
|----------|--|--|----------------------------|
| Memory | Fixed large value Needs to be estimated | Fixed small value | Increases monotonically |
| Accuracy | High but N -dependent | Increase with iteration Upper bound is N -dependent | Increase with iteration |
| Time | $T \propto O(t \cdot N)$ | $T \propto O(t \cdot N)$ | $T \propto O(t^2 \cdot N)$ |

Table 2: Comparison table that summarizes the principal differences between both three learning approaches: one-shot *single-pass*, online *single-pass* and *sequential*.

Table 2 compares the three approaches in terms of memory, accuracy, and computational time. The one-shot single-pass approach requires a large amount of data to achieve high accuracy and estimating the number of training data points can lead to under or over-training. The online single-pass approach alleviates memory allocation issues but has limited accuracy and is dependent on the value of n . While it is guaranteed to converge to the optimal solution, this is an asymptotic guarantee. The sequential approach combines the benefits of both: it increases the number of points over time, so there’s no need to estimate this parameter upfront, and stops copying once a desired accuracy is reached, as proven by Theorem 1’s monotonic accuracy increase with iterations. The main drawback of this method is its computational cost. The sequential approach grows quadratically with the number of data points, unlike the linear growth of both the single-pass and pure online approaches. This makes the sequential approach highly time-consuming in its current implementation. This comes as no surprise given the construction of the sequence of the set $\{S_t\}$. We discuss potential improvements to address this issue in what follows.

3.2.2 A MODEL COMPRESSION MEASURE

The purpose of a parametric machine learning model is to compress the data patterns relevant to a specific task. In the copying framework, the original model is considered the ground truth, and the goal is to imitate its decision behavior. Because this model produces hard classification labels, its decision boundary creates a partition of the feature space that results in a separable problem. Hence, any uncertainty measured during the copying process should only come from the model, not the data: a perfect copy should have no aleatoric uncertainty while compressing all the data patterns.

Assumption: Any uncertainty measured during the copying process is epistemic³.

It follows from this assumption that when the epistemic uncertainty over a given data point is minimum, then the data point is perfectly compressed by the model. This is, we can assume that a data point is perfectly learned by the copy when its uncertainty is zero.

Estimating uncertainty in practice requires that we measure the difference between the predictive distribution and the target label distribution. This can be done using various ratios or predictive entropy. We refer the reader to (DeVries and Taylor, 2018; Senge et al., 2014; Nguyen et al., 2018) for full coverage of this topic. For the sake of simplicity, here we consider the simplest approach to measuring uncertainty. We relax the hard-output constraint for the copy. Instead, we impose that, for any given data point, the copy returns a n_c -output vector of class probability predictions, such that

$$f_C : \mathbb{R}^d \times \Theta \longrightarrow [0, 1]^{n_c}.$$

For the sake of simplicity, we model the *uncertainty*, ρ , of the copy for a given data point z as the normalized euclidean norm of the distance between the n_c -vectors output by the original model and the copy, such that

$$\rho(z, \theta) = \frac{\|f_C(z, \theta) - f_{\mathcal{O}}(z)\|_2}{\sqrt{n_c}} \in [0, 1]. \quad (9)$$

A small ρ value indicates that the copy has low uncertainty (or strong confidence) in the class prediction output for z and that z is properly classified. Conversely, a large ρ value indicates that, despite the copy model having strong confidence in the class predicted for z , this prediction is incorrect. Or alternatively, despite the prediction being correct, there is a large dispersion among the output class probabilities.

This uncertainty measure can be used to evaluate how well individual data points are compressed by the copy model in the sequential framework. In the limit, this measure can also be used to assess how well the copy model replicates the original decision behavior. Therefore, we can introduce a new (non-negative and bounded) loss function such that for a given set of data points S_t , we define the empirical risk of the copy as

$$R_{emp}^{\mathcal{F}, \mathcal{S}}(f_C(\theta_j), f_{\mathcal{O}}) = \frac{1}{|S_t|} \sum_{z \in S_t} \rho^2(z, \theta). \quad (10)$$

The results of using the loss function are displayed in Figure 2. The plot illustrates the average value of ρ_t per iteration over all data points and runs for the different copying approaches introduced in Figure 1b). The value of ρ_t is significantly low for the sequential learning approach at iteration $t = 15$. Conversely, the single-pass and online learning methods exhibit higher uncertainty levels throughout the process. This difference can be attributed to the fact that the sequential approach continuously exposes the model to the same data points, reducing uncertainty iteration by iteration. The other two methods, however, only expose the model once to the same data, causing higher uncertainty levels.

3. Epistemic uncertainty corresponds to the uncertainty coming from the distribution of the model parameters. It can be reduced as the number of training samples increases. This allows us to single out the optimal model parameters. Thus, any measurement of uncertainty comes from the mismatch between the model parameters and the desired parameters, i.e. a perfect copy will have zero epistemic uncertainty.

The one-shot single-pass approach exposes the model to the whole dataset at once, while the sequential approach never exposes the model to the same data points more than once. As a result, the sequential approach has higher redundancy and reduces uncertainty faster when compared also to the online method, where the learner is never exposed to the same data points more than once.




Figure 2: Average uncertainty across the entire dataset for both the sequential (orange) and online single-pass (blue) copying approaches, averaged over 30 independent runs. The uncertainty achieved by the one-shot single-pass approach, with different values of n , is represented by dashed lines. The shaded area indicates a one standard deviation deviation from the mean.

3.2.3 A SAMPLE SELECTION POLICY FOR COPYING

Including an uncertainty measure in the training algorithm enables us to assess the degree of compression for each data point. The higher the compression, the better the copy model is at capturing the pattern encoded by each data point for the given task. This estimation can be used for selecting those points that contribute the most to the learning process. By filtering out the rest of the samples, we can reduce the number of resources consumed when copying. Hence, we enforce a policy that uses uncertainty as a criterion for sample selection.

At each iteration, data points with an uncertainty lower than a threshold δ are removed from the learning process (refer to Algorithm 2). The procedure starts with building a new sequential set by randomly sampling n points and adding them to S_{t-1} in line 3. Then, in

line 4, the uncertainty measure is used to select points with $\rho_j \geq \delta$, forming the filtered set S_t that is used to optimize the copy parameters.

Figure 3 presents the results for the sequential training with $n = 100$ and different values of the sample selection parameter δ in the *spirals* dataset. For comparison, results for the online single-pass approach and the pure sequential approach are also shown. Figure 3a) displays the average accuracy at each iteration. The results demonstrate that sequential training with sample selection performs better than online training, but falls short of the pure sequential setting.

Figures 3b) and c) show the change in ρ and number of synthetic data points used for training over increasing iterations, respectively. The online single-pass approach shows a constant uncertainty, while the pure sequential approach approaches 0. Contrary to what one might expect, the sample selection policy leads to an overall increase in uncertainty, while also reducing the number of data points used for training. Eventually, the number of samples Nn converges to a fixed value after a few iterations.

Algorithm 2 Sample selection policy

- 1: **Input:** Sample Set S_{t-1} , **int** n , **float** δ , **Distribution** P_Z , **Measurement** $\rho(\cdot)$, **Classifier** $f_{\mathcal{O}}$, **Classifier** $f_{\mathcal{C}}(\theta_{t-1}^*)$
 - 2: **Output:** Filtered set
 - 3: $\Xi \leftarrow S_{t-1} \cup \{(z_j, f_{\mathcal{O}}(z_j)) \mid z_j \sim P_Z, j = 1 \dots n\}$
 - 4: $S_t \leftarrow \{z \mid z \in \Xi, \rho(z, \theta_{t-1}^*) \geq \delta\}$ ▷ Filter according to uncertainty
 - 5: **return** S_t
-

We consider all three plots in Figure 3 at once. During the first iterations, when copies are still not sufficiently tuned to the data, there is no filtering effect. The number of points Nn grows almost equally for all the settings for a few iterations, as shown in Figure 3c). During this time, accuracy increases gradually, as shown in Figure 3a), and ρ decreases, as displayed in Figure 3b), as expected for a model that is compressing information. The differences arise after a few iterations when the filtering effect begins and the number of data points decreases in settings where the sample selection policy is enforced. At this stage, the copy models have a confidence level close to δ , which allows the removal of samples. As a result, accuracy stops growing and becomes flat, while average uncertainty starts to rise. With regard to the number of points, the curves quickly stabilize by $t = 18$. The model with the lowest threshold, $\delta = 10^{-10}$, accumulates the larger number of points (≈ 250). This is less than 10% of the number of points required by the pure sequential method, but still reaches a reasonable accuracy. On the other hand, models trained with $\delta = 10^{-8}$ and $\delta = 10^{-6}$ have lower accuracy compared to the online single-pass setting.

This algorithm addresses the problem in Equation 5 but with limitations. The uncertainty metric ρ is linked to the empirical fidelity error \mathcal{R}_{emp}^F . The points with high uncertainty also have the highest empirical fidelity error. However, removing points from set S_t violates the assumptions in Theorems 1 and 2. When the copy model has high confidence in the synthetic data, sample removal changes its behavior from sequential (adding points incrementally) to online (fixed number of points per iteration). For parametric models, this shift results in *catastrophic forgetting*, as shown in Figure 3 where sequential models start to act like online models. The removal of samples increases uncertainty in the copy

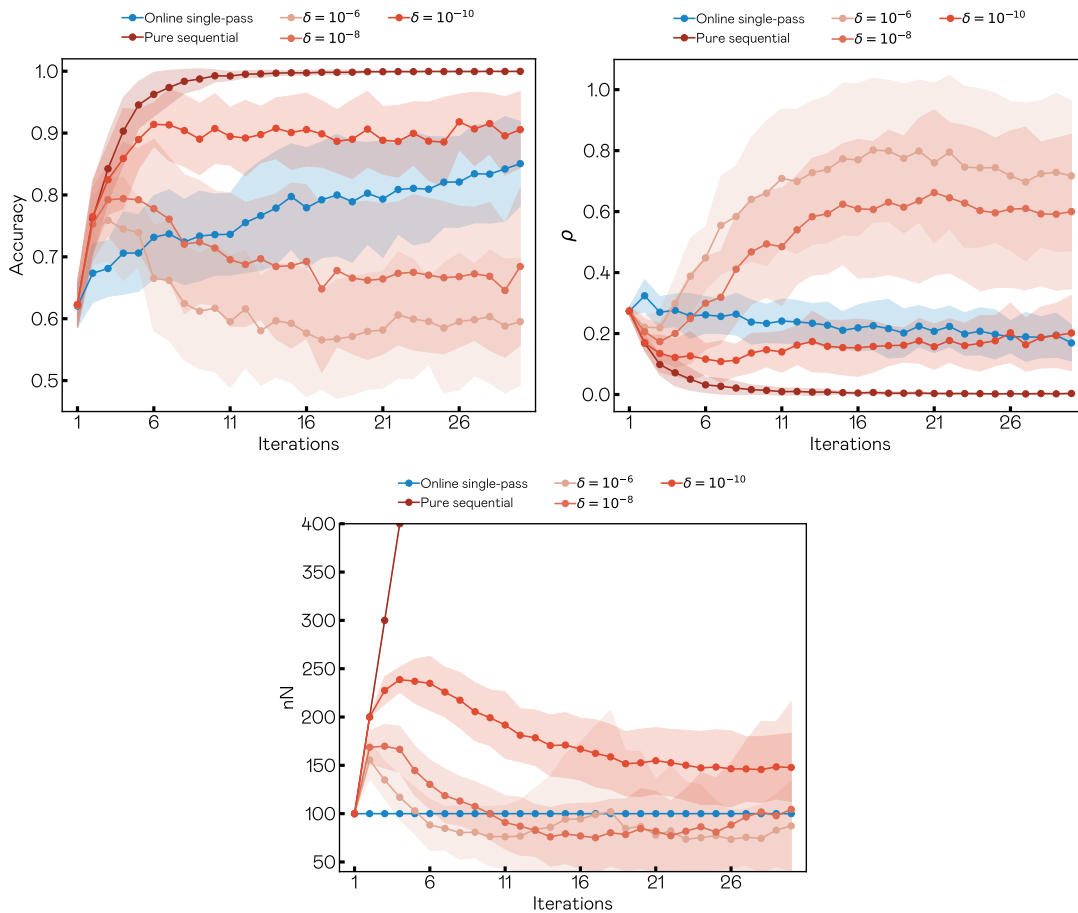


Figure 3: a) Accuracy and b) uncertainty per iteration for the sequential approach with different uncertainty thresholds, averaged over 30 independent runs. Results for the online setting are also shown in blue for comparative purposes at every iteration. c) Number of data points used at each iteration. The y -axis is restricted to 400 in order to make curves observable. The number of points of the sequential learning (orange line) grows linearly until $30 \cdot 100 = 3000$.

models, causing them to forget prior information. To address this, we introduce several improvements in the implementation of Step 2 in the alternating optimization algorithm.

3.3 Step 2: Optimizing copy model parameters

Let us briefly turn our attention to Equation 6. This equation models the challenge of controlling the capacity of the copy model while minimizing the loss function. To tackle this issue, we introduce a capacity-enhancement strategy. For this purpose, it's worth noting that if we assume that the copy model has enough capacity, then Equation 6 can be simplified for a given iteration t as follows

Algorithm 3 Empirical risk minimization implementation

- 1: **Input:** Sample Set S , float ε , Classifier $f_{\mathcal{O}}$
 - 2: **Output:** Copy parameters at step k
 - 3: $k \leftarrow 0, \epsilon_k \leftarrow C, C \gg \varepsilon$ $\triangleright C$ takes an arbitrarily large value.
 - 4: **while** $R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta_k), f_{\mathcal{O}}; \Omega) \geq \varepsilon$ **do**
 - 5: $\theta_k^* \leftarrow \arg \min_{\theta} R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}; \Omega)$ subject to $R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}; \Omega) \geq \epsilon_k$ $\Big|_{\theta_k^0 = \theta_{k-1}^*}$
 - 6: $k \leftarrow k + 1$
 - 7: $\epsilon_k \leftarrow \min(\epsilon_{k-1}/2, \varepsilon)$ \triangleright Reduce the value of epsilon.
 - 8: Increase Ω
 - 9: **end while**
 - 10: **return** θ_k^*
-

$$\begin{aligned} & \underset{\theta}{\text{minimize}} && \Omega(\theta) && (11) \\ & \text{subject to} && \left\| R_{emp}^{\mathcal{F}, S}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}) \right\|_{S=S_t^*} &< \varepsilon. \end{aligned}$$

Having the above simplification in mind, our proposed scheme, outlined in Algorithm 3, aims to control the copy model’s capacity while minimizing the loss function. It iteratively solves the parameter optimization problem in stages, ensuring that the empirical risk decreases at each step. At iteration t , the copy model starts with a small capacity Ω and solves the optimization problem for increasing upper bounds ϵ_k of the target value ε . For increasing values of k , we increase the capacity and reduce the upper bound. The larger the capacity, the more the empirical risk decreases and the smaller the complexity of the approximation to the target value.

In practice, we train copies using the loss function in Equation 10. Using stochastic subgradient optimization models (neural networks with “Stochastic Gradient Descent”), we control model complexity with an early stopping criterion. Thus, delaying early stopping increases model complexity. To improve convergence, hyperparameters are updated at each iteration.

3.3.1 FORCING THE MODEL TO REMEMBER

With the previous scheme in mind, we revisit the argument in Section 3.2.3. As noted, the sample removal policy conflicts with some of the assumptions made in the theorems introduced in Section 3.2, leading to a forgetting effect in Step 2. To restore the convergence properties, we first point out that, as shown by Theorem 2, parameters converge to their optimal value in the sequential approach. This implies that the difference between two consecutive terms in the sequence also converges to zero, such that

$$\left\| \theta_{t+1}^* - \theta_t^* \right\| \longrightarrow 0. \tag{12}$$

Algorithm 4 Memory aware empirical risk minimization implementation

- 1: **Input:** Sample Set S , float ε , Classifier $f_{\mathcal{O}}$, Parameters θ_{t-1}^*
 - 2: **Output:** Copy parameters at step k
 - 3: $k \leftarrow 0$, $\epsilon_k \leftarrow C$, $C \gg \varepsilon$ $\triangleright C$ takes an arbitrarily large value.
 - 4: **while** $R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta_k), f_{\mathcal{O}}; \Omega) \geq \varepsilon$ **do**
 - 5:

$$\theta_k^* \leftarrow \arg \min_{\theta} R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}; \Omega) + \lambda \cdot \|\theta - \theta_{t-1}^*\|$$

$$\text{subject to } R_{emp}^{\mathcal{F}, S^k}(f_{\mathcal{C}}(\theta), f_{\mathcal{O}}; \Omega) \geq \epsilon_k \left| \begin{array}{l} \\ \theta_k^0 = \theta_{k-1}^* \end{array} \right.$$
 - 6: $k \leftarrow k + 1$
 - 7: $\epsilon_k \leftarrow \min(\epsilon_{k-1}/2, \varepsilon)$ \triangleright Reduce the value of epsilon
 - 8: Increase Ω
 - 9: **end while**
 - 10: **return** θ_k^*
-

The asymptotic invariant in Equation 12 can be forced to preserve the compressed model obtained from previous iterations even after filtering the data points. We add a regularization term to the loss function at iteration t , which minimizes the left term

$$\mathcal{L}_t = R_{emp}^{\mathcal{F}, S^t}(f_{\mathcal{C}}(\theta_t), f_{\mathcal{O}}; \Omega) + \lambda \cdot \|\theta_t - \theta_{t-1}^*\| \tag{13}$$

This regularization term originates from the derived theorems and can also be found in the literature under the name of Elastic Weight Consolidation (EWC), though derived heuristically from a different set of assumptions (Kirkpatrick et al., 2017). Algorithm 4 outlines the implementation of this strategy.

Continuing with the spirals example, the experimental results for the sequential copying process with four different λ values and a threshold $\delta = 10^{-8}$ are shown in Figure 4 (see also Figure 3 for comparison). As before, Figure 4a) reports the change in accuracy for increasing iterations, while Figures 4b) and c) present the results for the uncertainty ρ and the number of data points used nN . Again, we can combine the information from all three panels to better understand the behavior displayed by the copy model when *forced to remember*. Initially, the model has not been trained, so the number of points increases while the uncertainty decreases. Once the removal threshold δ is reached, the regularization term becomes relevant. For low λ values, the ρ -term of Equation 12 dominates the cost function, and the optimization process learns the new data points. As a result, the copy parameters adapt to new data points and the model exhibits forgetting. Conversely, for *large* λ values, the regularization term will dominate the optimization. The model is therefore forced to retain more data points to ensure the ρ term can compete with the regularization term. In this setting, the number of data points still grows, but in a sub-linear way. The most desirable behavior is a constant number of points during the training process. For this particular example, we observe this behavior for $\lambda = 0.05$ and $\lambda = 0.1$.

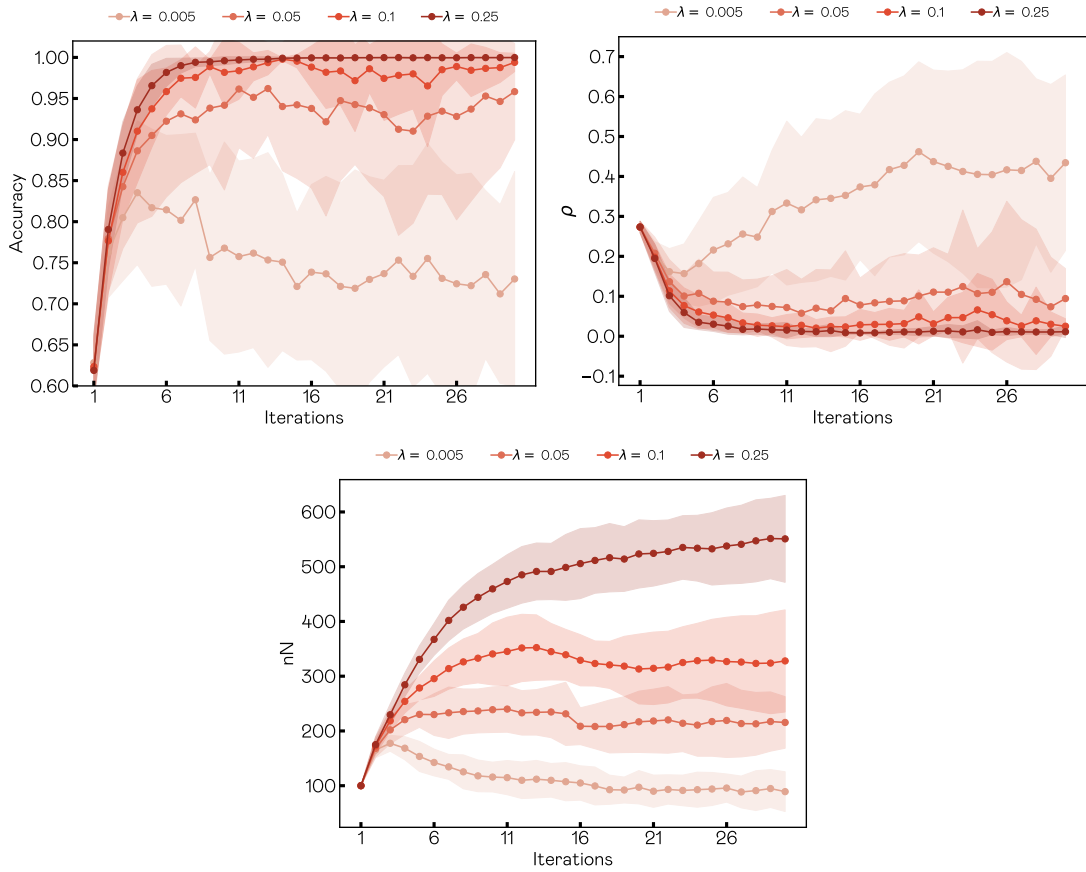


Figure 4: a) Copy accuracy averaged over 30 independent runs for four different lambda values: $\lambda = 0.005, 0.05, 0.1, 0.25$ with the same drooping threshold $\delta = 10^{-8}$. b) Average uncertainty at each iteration. c) Number of data points used at each iteration. Shaded region shows $\pm\sigma$

3.3.2 AUTOMATIC LAMBDA

In the previous section, we observed an increment in the number of data points when using large lambda values, to compensate for the memorization effect introduced by the regularization term. This is to consolidate the previous knowledge acquired by the model. In contrast, small lambda values promote short-lived data compression. This is desirable at the beginning of the learning process or when the data distribution suffers a shift. To retain the best of both regimes, we propose a heuristic to dynamically adapt the λ value to the needs of the learning process. Our underlying intuition is that, whenever the amount of data required increases, the memory term prevents the copy from adapting to new data. This observation signals that the value of λ must decrease. Equivalently, when we observe a decrement in the number of data points after data filtering, this suggests that the model can classify most of them correctly. This indicates that we must stabilize it to avoid unnecessary model drift due to future disturbances in the data. Hence, we must increase the λ parameter.

Algorithm 5 Sequential approach with alternating optimization

```

1: Input: int  $T$ , int  $n$ , float  $\varepsilon$ , Classifier  $f_{\mathcal{O}}$ 
2: Output: Optimal copy parameters
3:  $S_0 \leftarrow \{(z_j, f_{\mathcal{O}}(z_j)) \mid z_j \sim P_{\mathcal{Z}}, j = 1 \dots n\}$ 
4:  $\theta_0^* \leftarrow \arg \min_{\theta} \frac{1}{|S_0|} \sum_{z \in S_0} \rho^2(z, \theta)$ 
5: for  $t = 1$  to  $T$  do
6:    $S_t^* \leftarrow \text{Step1}(S_{t-1}^*, n, \delta, P_{\mathcal{Z}}, \rho(), f_{\mathcal{O}}, f_{\mathcal{C}}(\theta_{t-1}^*))$  ▷ Algorithm 2
7:    $\lambda \leftarrow \begin{cases} \lambda/2 & \text{if } |S_t| \geq |S_{t-1}|, \\ 1.5 \cdot \lambda & \text{otherwise.} \end{cases}$  ▷ Adaptive lambda
8:    $\theta_t^* \leftarrow \text{Step2}(S_t, \varepsilon, f_{\mathcal{O}}, \theta_{t-1}^*)$  ▷ Algorithm 4
9: end for
10: return  $\theta_t^*$ 

```

Thus, considering the set of data at iteration t , S_t , we force the described behavior by automatically updating the value of λ parameter as follows

$$\lambda = \begin{cases} \lambda/2 & \text{if } |S_t| \geq |S_{t-1}|, \\ 1.5 \cdot \lambda & \text{otherwise.} \end{cases}$$

The updated optimization process, including the modifications discussed, is presented in Algorithm 5, the final algorithm proposed in this article. Our implementation models data trends by computing the difference between the number of points in the previous iteration and the number of points after data filtering. The filtering process occurs at the beginning of each iteration after generating a new set of n samples.

As before, we demonstrate this improvement for the *spirals* example. We repeat the experiments using Algorithm 5 with $n = 100$ and three different dropping thresholds: $\delta = 10^{-6}$, 10^{-8} and 10^{-10} as we did in Subsection 3.2.3 where the dropping procedure was first introduced. Recall that, as discussed in Figure 3 only the setting corresponding to the smallest δ value managed to perform better than the online approach, yet still performed worse than the pure sequential case, where the model keeps all data points. The remaining thresholds lost too many data points, and their performance was worse than the online approach for a similar number of data points. The results obtained for the full Algorithm 5 are depicted in Figure 5. The accuracy obtained using the automatic regularization term is comparable with the desired optimal *pure sequential* case, where the model keeps all data points. Even for the most conservative approach, where we use a very small dropping threshold, the number of points used after 30 iterations is 1/3 smaller than those required for the pure sequential setting. Moreover, even when deliberately forcing a very volatile setup by using a large value of delta, $\delta = 10^{-6}$, we obtain an accuracy larger than 0.95 for an almost constant number of data points equal to 200.

4. Experiments

We present the results of our experimental study on the sequential copy approach applied to a set of heterogeneous problems. Our results are analyzed using various performance metrics and compared to the single-pass approach in both one-shot and online settings.

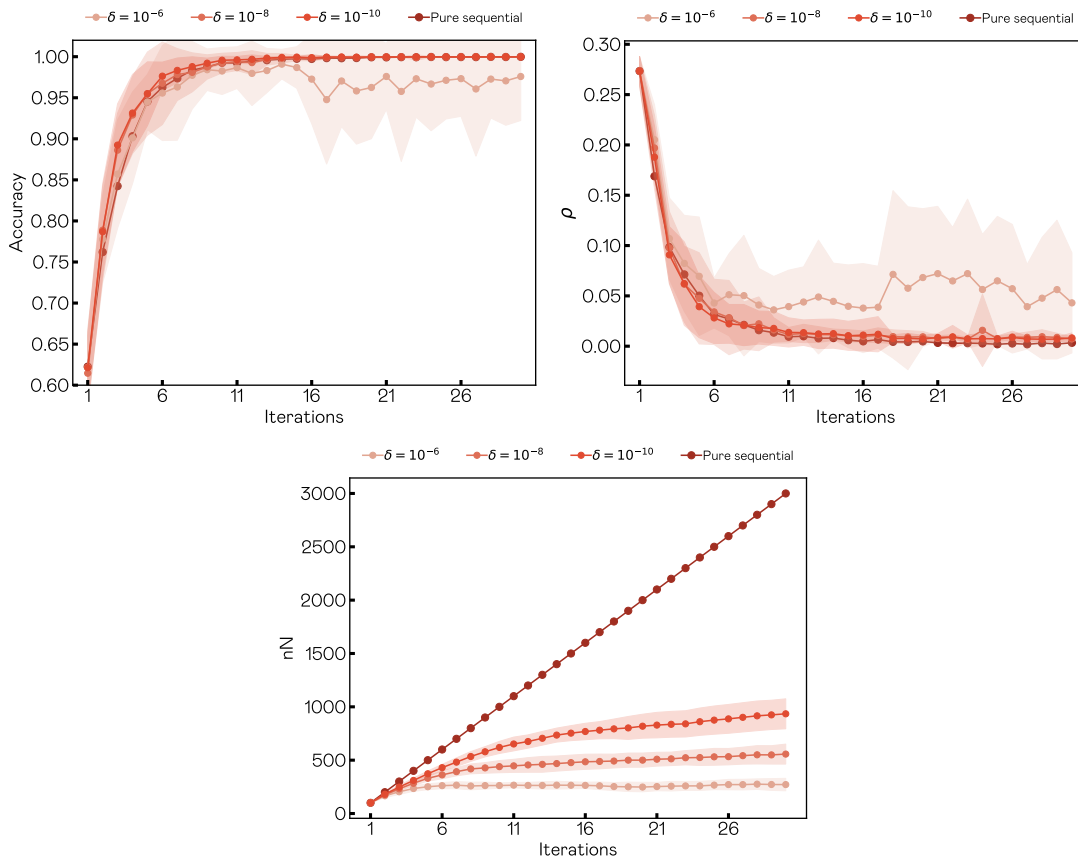


Figure 5: a) Copy accuracy averaged over 30 independent runs for three different drooping thresholds: 10^{-6} , 10^{-8} , and 10^{-10} . b) Average uncertainty at each iteration. c) Number of data points used at each iteration. The shaded region shows $\pm\sigma$.

Before presenting the results, we provide a clear and reproducible description of the data and experimental setup.

4.1 Experimental set-up

Data. We use 58 datasets from the UCI Machine Learning Repository database (Dheeru and Karra Taniskidou, 2017) and follow the experimental methodology outlined in (Unceta et al., 2020a). For a more detailed explanation of the problem selection and data preparation process, we refer the reader to the mentioned article.

Pre-processing. We convert all nominal attributes to numerical and rescale variables to have zero mean and unit variance. We split the pre-processed data into stratified 80/20 training and test sets. We sort the datasets in alphabetical order and group them in sets of 10, and assign each group one of the following classifiers: AdaBoost (*adaboost*), artificial neural network (*ann*), random forest (*rfc*), linear SVM (*linear_svm*), radial basis kernel

SVM (*rbf_svm*) and gradient-boosted tree (*xgboost*). We train all models using a 3-fold cross-validation over a fixed parameter grid. A full description of the 58 datasets, including general data attributes, and their assigned classifier, can be found in Table 3.

Models. We copy the resulting original classifiers using a fully-connected neural network with three hidden layers, each consisting of 64, 32, and 10 *ReLU* neurons and a SoftMax output layer. We use no pre-training or drop-out and initialize the weights randomly. We implement the sequential copy process as described in Algorithm 5 above.

Parameter setting. We train these models sequentially for 30 iterations. At each iteration t , we generate $n = 100$ new data points by randomly sampling a standard normal distribution $\mathcal{N}(0,1)$. We use Algorithm 5, discarding any data points for which the instantaneous copy model has an uncertainty ρ below a defined threshold δ . We adjust the weights at each iteration using the *Adam* optimizer with a learning rate of $5 \cdot 10^{-4}$. For each value of t , we use 1000 epochs with balanced batches of 32 data points. We use the previously defined normalized uncertainty average as the loss function and evaluate the impact of the δ parameter by running independent trials for $\delta \in \{5 \cdot 10^{-4}, 10^{-4}, 5 \cdot 10^{-5}, 10^{-5}, 5 \cdot 10^{-6}, 10^{-6}, 5 \cdot 10^{-7}, 10^{-7}, 5 \cdot 10^{-8}, 10^{-8}, 10^{-9}, 10^{-10}\}$. Additionally, we allow the λ parameter to be updated automatically, starting from a value of 0.5.

Hardware. We perform all experiments on a server with 28 dual-core AMD EPYC 7453 at 2.75 GHz, and equipped with 768 Gb RDIMM/3200 of RAM. The server runs on Linux 5.4.0. We implement all experiment in Python 3.10.4 and train copies using TensorFlow 2.8. For validation purposes, we repeat each experiment 30 times and present the average results of all repetitions.

4.2 Metrics and performance plots

We evaluate the copy performance using three metrics, including copy accuracy (\mathcal{A}_C). The copy accuracy is calculated as the accuracy of the copy on the original test set⁴. Given a model f , we define the copy accuracy as the fraction of correct predictions made by this model on a data set of labelled pairs, $\mathcal{D} = (x_j, y_j)$, as:

$$\mathcal{A}_C^f = \frac{1}{N} \sum_j 1^N \mathbb{I}[f(x_j) == y_j], \quad (14)$$

where N is the number of samples and $\mathbb{I}[\text{cond}]$ is the indicator function that returns 1 if the condition is true, i.e., when the model predicts the right outcome. We use the definition above to compare the results obtained when using the sequential approach, $\mathcal{A}_C^{\text{seq}}$, with those obtained when training copies based on the single-pass approach, $\mathcal{A}_C^{\text{single}}$.

Another performance metric is the area under the normalized convergence accuracy curve (*conv*). This metric measures the convergence speed of the sequential approach and is defined as the area under the curve of the copy accuracy increment per iteration, normalized by the maximum copy accuracy. The *conv* value is in the range of 0 to 1 and represents

4. Computing this value requires the original test data to be known and accessible. We consider this assumption to be reasonable, as it provides a lower bound for all the other metrics.

| Dataset | Classes | Samples | Features | Original | \mathcal{A}_C |
|--------------------------------|---------|---------|----------|----------------------|-----------------|
| abalone | 3 | 4177 | 8 | <i>adaboost</i> | 0.545 |
| acute-inflammation | 2 | 120 | 6 | <i>adaboost</i> | 1.0 |
| acute-nephritis | 2 | 120 | 6 | <i>adaboost</i> | 1.0 |
| bank | 2 | 4521 | 16 | <i>adaboost</i> | 0.872 |
| breast-cancer-wisc-diag | 2 | 569 | 30 | <i>adaboost</i> | 0.921 |
| breast-cancer-wisc-prog | 2 | 198 | 33 | <i>adaboost</i> | 0.7 |
| breast-cancer-wisc | 2 | 699 | 9 | <i>adaboost</i> | 0.914 |
| breast-cancer | 2 | 286 | 9 | <i>adaboost</i> | 0.69 |
| breast-tissue | 6 | 106 | 9 | <i>adaboost</i> | 0.545 |
| chess-krvkp | 2 | 3196 | 36 | <i>ann</i> | 0.995 |
| congressional-voting | 2 | 435 | 16 | <i>ann</i> | 0.609 |
| conn-bench-sonar-mines-rocks | 2 | 208 | 60 | <i>ann</i> | 0.833 |
| connect-4 | 2 | 67557 | 42 | <i>ann</i> | 0.875 |
| contrac | 3 | 1473 | 9 | <i>ann</i> | 0.573 |
| credit-approval | 2 | 690 | 15 | <i>ann</i> | 0.79 |
| cylinder-bands | 2 | 512 | 35 | <i>ann</i> | 0.777 |
| echocardiogram | 2 | 131 | 10 | <i>ann</i> | 0.815 |
| energy-y1 | 3 | 768 | 8 | <i>ann</i> | 0.974 |
| energy-y2 | 3 | 768 | 8 | <i>ann</i> | 0.922 |
| fertility | 2 | 100 | 9 | <i>random_forest</i> | 0.9 |
| haberman-survival | 2 | 306 | 3 | <i>random_forest</i> | 0.613 |
| heart-hungarian | 2 | 294 | 12 | <i>random_forest</i> | 0.763 |
| hepatitis | 2 | 155 | 19 | <i>random_forest</i> | 0.742 |
| ilpd-indian-liver | 2 | 583 | 9 | <i>random_forest</i> | 0.615 |
| ionosphere | 2 | 351 | 33 | <i>random_forest</i> | 0.944 |
| iris | 3 | 150 | 4 | <i>random_forest</i> | 0.933 |
| magic | 2 | 19020 | 10 | <i>linear_svm</i> | 0.801 |
| mammographic | 2 | 961 | 5 | <i>random_forest</i> | 0.803 |
| miniboone | 2 | 130064 | 50 | <i>random_forest</i> | 0.936 |
| molec-biol-splice | 3 | 3190 | 60 | <i>random_forest</i> | 0.944 |
| mushroom | 2 | 8124 | 21 | <i>linear_svm</i> | 0.979 |
| musk-1 | 2 | 476 | 166 | <i>linear_svm</i> | 0.812 |
| musk-2 | 2 | 6598 | 166 | <i>linear_svm</i> | 0.958 |
| oocytes_merluccius_nucleus_4d | 2 | 1022 | 41 | <i>linear_svm</i> | 0.771 |
| oocytes_trisopterus_nucleus_2f | 2 | 912 | 25 | <i>linear_svm</i> | 0.803 |
| parkinsons | 2 | 195 | 22 | <i>linear_svm</i> | 0.923 |
| pima | 2 | 768 | 8 | <i>linear_svm</i> | 0.721 |
| pittsburg-bridges-MATERIAL | 3 | 106 | 7 | <i>linear_svm</i> | 0.909 |
| pittsburg-bridges-REL-L | 3 | 103 | 7 | <i>linear_svm</i> | 0.667 |
| pittsburg-bridges-T-OR-D | 2 | 102 | 7 | <i>rbf_svm</i> | 0.857 |
| planning | 2 | 182 | 12 | <i>rbf_svm</i> | 0.703 |
| ringnorm | 2 | 7400 | 18 | <i>rbf_svm</i> | 0.983 |
| seeds | 3 | 210 | 7 | <i>rbf_svm</i> | 0.881 |
| spambase | 2 | 4601 | 57 | <i>rbf_svm</i> | 0.926 |
| statlog-australian-credit | 2 | 690 | 14 | <i>rbf_svm</i> | 0.681 |
| statlog-german-credit | 2 | 1000 | 24 | <i>rbf_svm</i> | 0.765 |
| statlog-heart | 2 | 270 | 13 | <i>rbf_svm</i> | 0.852 |
| statlog-image | 7 | 2310 | 18 | <i>rbf_svm</i> | 0.952 |
| statlog-vehicle | 4 | 846 | 18 | <i>xgboost</i> | 0.765 |
| synthetic-control | 6 | 600 | 60 | <i>rbf_svm</i> | 1.0 |
| teaching | 3 | 151 | 5 | <i>xgboost</i> | 0.548 |
| tic-tac-toe | 2 | 958 | 9 | <i>xgboost</i> | 0.974 |
| titanic | 2 | 2201 | 3 | <i>xgboost</i> | 0.778 |
| twonorm | 2 | 7400 | 20 | <i>xgboost</i> | 0.976 |
| vertebral-column-2classes | 2 | 310 | 6 | <i>xgboost</i> | 0.839 |
| vertebral-column-3classes | 3 | 310 | 6 | <i>xgboost</i> | 0.806 |
| waveform-noise | 3 | 5000 | 40 | <i>xgboost</i> | 0.843 |
| waveform | 3 | 5000 | 21 | <i>xgboost</i> | 0.843 |
| wine | 3 | 178 | 11 | <i>xgboost</i> | 0.944 |

Table 3: Description of datasets.

the fraction of time required for the system to reach a steady convergence state. For T iterations, the value of $conv$ is defined as follows

$$conv = \frac{1}{T} \frac{\int_0^T \mathcal{A}_C^{\text{seq}}(t) dt}{\max_{t \in [0, T]} \mathcal{A}_C^{\text{seq}}(t)}, \quad (15)$$

where $\mathcal{A}_C^{\text{seq}}(t)$ corresponds to the copy accuracy increment iteration by iteration.

Intuitively, $conv$ metric measures the time required for the system to reach a steady convergence state. Given T iterations of the sequential approach, a convergence speed of $conv$ means that the algorithm reaches the steady state at step $2(1 - conv)T$. A $conv$ value

of 90% tells us that we can reach convergence as fast as $0.2T$. This is, the system requires only 20% of the allocated time to converge.

Finally, we also introduce the efficiency metric, eff . This metric evaluates the computational cost of the copying process in terms of the number of synthetic data points used for training. We compute it by comparing the actual number of points used in the sequential approach with the theoretical number of points that would be used if no sample removal policy was applied. We define eff as:

$$eff = 1 - \frac{\int_0^T \eta(t) dt}{\int_0^T n \cdot t dt}, \quad (16)$$

where $\eta(t)$ is the number of points used in the sequential approach at each iteration t , as shown in Figure 5(c).

The eff metric models the expected number of samples required for copying. This value can be roughly approximated by $(1 - eff)/2$. A eff value of 90% indicates that on average, only 5% of the available data points are used in the process. Note that both the pure sequential approach, where no sample removal policy is used, and the single-pass approach have 0 efficiency because they both use all the available data points for training.

We evaluate the performance of the sequential approach by combining various metrics into a single representation. We consider models with copy accuracy within 5% of the single-pass result ($\mathcal{A}_C^{seq}/\mathcal{A}_C^{single} | nT > 0.95$). Out of these configurations, we select the ones with the highest copy accuracy (\mathcal{A}_C^{seq}), best efficiency (eff), and fastest convergence ($conv$), referred to as the *Best accuracy*, *Best efficiency*, and *Best convergence*, respectively. To visualize the results, we present four plots: (1) a comparison of the copy accuracy between the single-pass approach and the *Best accuracy* model, (2) a comparison of the copy accuracy and efficiency between the *Best accuracy* and *Best efficiency* results, (3) a comparison of the copy accuracy and convergence rate between the *Best accuracy* and *Best convergence* configurations, and (4) a demonstration of the relationship between convergence and efficiency.

As an example, Figure 6 shows these plots for three toy datasets: *spirals*, *moons*, and *yin-yang*. In Figure 6a), we can compare the performance degradation or improvement of the sequential and single-pass approaches. In Figure 6b), a comparison between the gain in efficiency (square marker) and the the best-performing configuration can be seen for each dataset. The most efficient configuration typically displays a significant efficiency gain. Figure 6c) compares the *Best accuracy* and *best convergence* operational points. We observe that both configurations are nearly indistinguishable in terms of both accuracy and convergence, i.e. the most accurate model is also the one which converges faster. Finally, in Figure 6d) we observe that the *Best efficiency* configurations yield a significant improvement in efficiency, while the loss in convergence is not too big.

The biggest gains are therefore obtained in terms of efficiency. This is a relevant result, because it shows that the sequential approach to copying can provide significant advantages for memory usage and computational resource allocation.

4.3 Results

We report the metrics for each UCI dataset as introduced above. Table 4 lists the values of efficiency (eff), convergence rate ($conv$), and accuracy \mathcal{A}_C^{seq} for the three operational points:

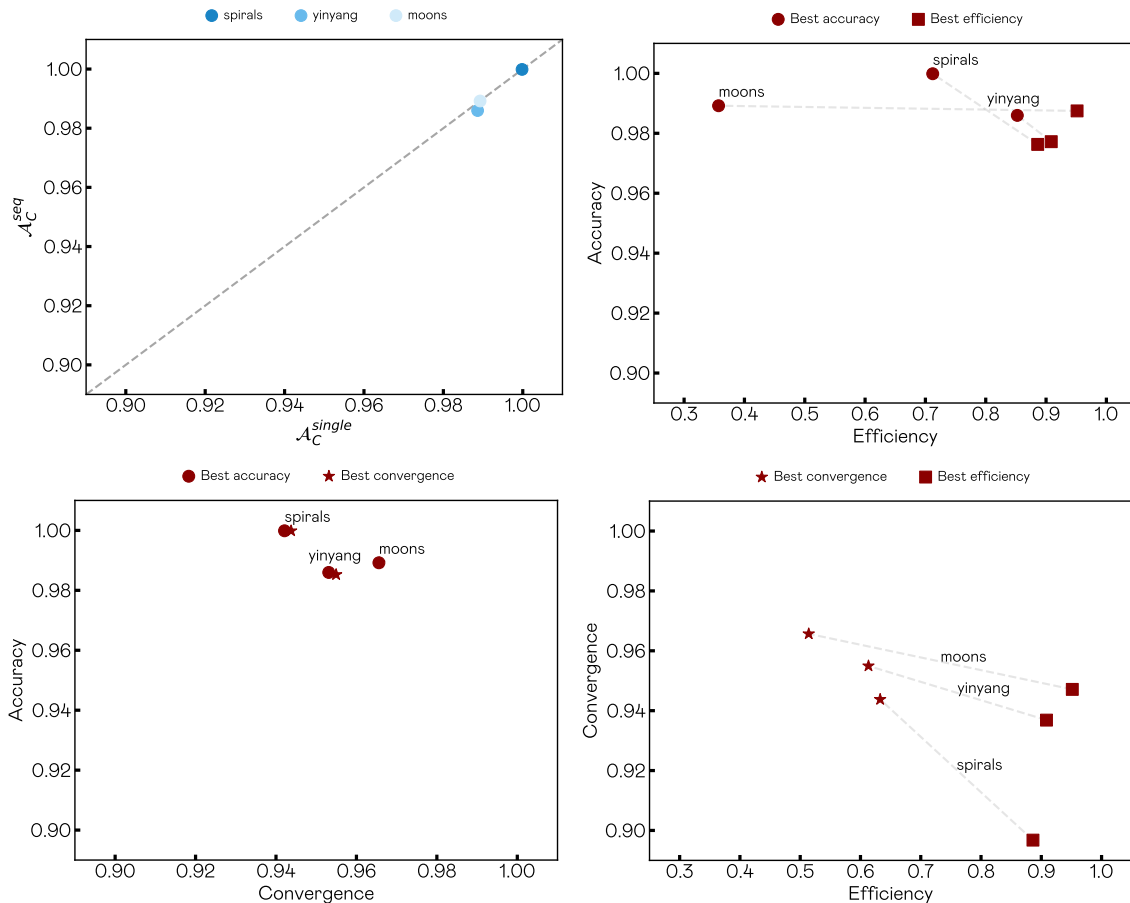


Figure 6: a) Comparison of accuracy between single-pass and sequential approaches on the *spirals*, *yinyang*, and *moons* datasets. The dashed line marks equal accuracy between the two methods. Points above (below) the line indicate better accuracy for the sequential (single-pass) approach. b), c), and d) Comparison of *Best accuracy* (circles), *Best efficiency* (squares) and *Best convergence* (stars) operational points for each dataset. The gray dashed lines connect the points of the same dataset and display the linear fit of intermediate solutions.

Best accuracy, *Best efficiency*, and *Best convergence*. These results are compared with the original accuracy (\mathcal{A}_O) and the one-shot single-pass copy accuracy (\mathcal{A}_C^{single}).

We observe that not all copies can perfectly reproduce the original accuracy. This effect is observed in 19 datasets for both single-pass and sequential models. Possible reasons include a mismatch between the copy’s capacity and the boundary complexity, or a misalignment between the sampling region and the original data distribution. This observation is in line with results in the literature but outside of the scope of this article. Conversely, in 8 datasets, copies outperform the original classifier performance, specifically in 4 cases. This unexpected result may be due to statistical noise and requires further investigation.

The relevant results for our proposal show that the sequential copy process at the *Best accuracy* operational point matches the single-pass approach in 53 of 58 problems, performs worse in 4 datasets and better in 1. On average, the copy process converges in 11.5% of the allotted time with a convergence metric of $conv = 0.942$ and efficiency of $eff = 0.716$. This requires an average of 14% of the samples used by the single-pass. A graphic display of the results comparing both approaches is shown in Figure 7. The most notable results have been highlighted in darker colors to ease interpretation. Values plotted in the diagonal correspond to cases where both approaches yield comparable results. Most of the sequential copies recover most of the single-pass accuracy, even when training on smaller synthetic datasets. In some datasets, however, copies based on the sequential approach fall below the diagonal. This effect is observed when the amount of memory (n) is not enough to describe the decision boundary entirely. Finally, we observe some cases where points lie above the diagonal, signaling that sequential copies improve over the single-pass results.




Figure 7: Comparison of accuracy between the one-shot single-pass and sequential approaches. Points over the dashed line exhibit the same accuracy using both methods whereas points above (below) the line are those where the sequential (single-pass) method gives better performance. Color gets darker with the distance to the dashed line.

The results of the *Best efficiency* and *Best convergence* operational points have an average accuracy degradation of 5%, which is not statistically significant. For the *Best efficiency* operational point, 51 datasets have a degradation in performance, but there’s a statistically significant improvement in efficiency in 47 datasets ($eff=0.882$) using only 6% of the data points. The convergence speed increases by 17.4% in the allotted time, meaning 3% more time compared to the *Best accuracy* configuration.

For the *Best convergence* operational point, 20 out of 58 datasets show a performance degradation, which is not statistically significant. There’s a statistical efficiency loss in 10

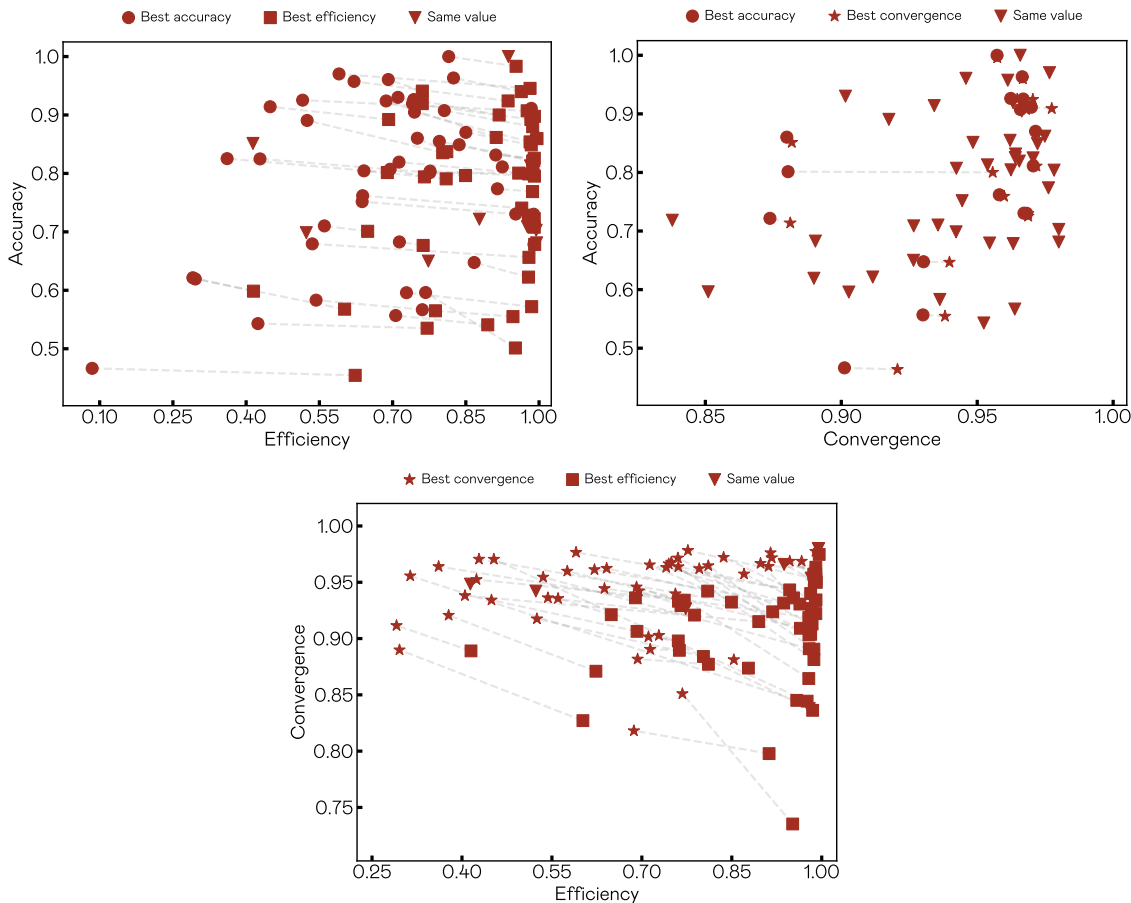


Figure 8: Comparison of *Best accuracy* (circles), *Best efficiency* (squares) and *Best convergence* (stars) operational points for each dataset in terms of a) accuracy and efficiency, b) accuracy and convergence and c) convergence and efficiency. The inverted triangles identify those cases where the considered operational points are the same. The gray dashed lines connect the points of the same dataset and display the linear fit of intermediate solutions.

datasets, and it requires an average of 15% of the allotted time. This operational point barely improves the convergence speed compared to *Best accuracy*. It differs from it in average in 20 datasets, with a required time to reach the steady state of 11.2% of the allotted time. This small time difference is due to the automatic lambda setting. Large λ values and small δ values ensure fast convergence speeds. The automatic lambda algorithm starts with a large lambda value for fast convergence at the initial steps, then reduces it in subsequent iterations to improve accuracy. This results in similar metrics for both *Best accuracy* and *Best convergence* configurations.

| Dataset | \mathcal{A}_C | $\mathcal{A}_C^{\text{single}}$ | | | Best accuracy | | | Best efficiency | | | Best convergence | | |
|--------------------------------|-----------------|---------------------------------|---------------|--------------|-----------------|---------------|--------------|------------------------------|---------------|--------------|------------------------------|---------------|--------------|
| | | \mathcal{A}_C | conv | eff | \mathcal{A}_C | conv | eff | $\mathcal{A}_C^{\text{seq}}$ | conv | eff | $\mathcal{A}_C^{\text{seq}}$ | conv | eff |
| abalone | 0.545 | 0.465±0.095 | 0.085±0.01 | 0.466±0.087 | 0.901±0.135 | 0.623±0.027 | 0.871±0.143 | 0.454±0.087 | 0.378±0.019 | 0.921±0.136 | 0.464±0.079 | | |
| acute-inflammation | 1.0 | 1.0±0.092 | 0.937±0.032 | 1.0±0.063 | 0.966±0.012 | 0.937±0.032 | 0.966±0.012 | 1.0±0.063 | 0.937±0.032 | 0.966±0.012 | 1.0±0.063 | | |
| acute-nephritis | 1.0 | 1.0±0.073 | 0.815±0.014 | 1.0±0.049 | 0.957±0.017 | 0.953±0.007 | 0.936±0.037 | 0.983±0.074 | 0.87±0.024 | 0.957±0.019 | 0.996±0.049 | | |
| bank | 0.872 | 0.808±0.05 | 0.776±0.041 | 0.801±0.064 | 0.88±0.047 | 0.958±0.005 | 0.845±0.11 | 0.801±0.125 | 0.314±0.023 | 0.956±0.03 | 0.8±0.043 | | |
| breast-cancer-wisc-diag | 0.921 | 0.904±0.125 | 0.414±0.018 | 0.851±0.086 | 0.948±0.07 | 0.414±0.018 | 0.948±0.07 | 0.851±0.086 | 0.414±0.018 | 0.948±0.07 | 0.851±0.086 | | |
| breast-cancer-wisc-prog | 0.7 | 0.73±0.069 | 0.976±0.038 | 0.709±0.113 | 0.927±0.113 | 0.976±0.038 | 0.844±0.113 | 0.709±0.117 | 0.976±0.038 | 0.927±0.113 | 0.709±0.117 | | |
| breast-cancer-wisc | 0.914 | 0.934±0.125 | 0.516±0.019 | 0.925±0.093 | 0.967±0.017 | 0.976±0.003 | 0.827±0.113 | 0.907±0.241 | 0.453±0.02 | 0.97±0.025 | 0.925±0.114 | | |
| breast-cancer | 0.69 | 0.736±0.055 | 0.638±0.02 | 0.944±0.048 | 0.944±0.048 | 0.978±0.014 | 0.919±0.087 | 0.728±0.09 | 0.638±0.02 | 0.944±0.048 | 0.752±0.054 | | |
| breast-tissue | 0.545 | 0.541±0.102 | 0.706±0.015 | 0.557±0.091 | 0.93±0.118 | 0.895±0.019 | 0.915±0.14 | 0.541±0.116 | 0.405±0.024 | 0.938±0.104 | 0.555±0.082 | | |
| chess-krvkp | 0.995 | 0.983±0.032 | 0.691±0.032 | 0.961±0.014 | 0.946±0.014 | 0.761±0.026 | 0.933±0.017 | 0.941±0.031 | 0.691±0.032 | 0.946±0.014 | 0.961±0.031 | | |
| conn-bench-sonar-mines-rocks | 0.833 | 0.826±0.097 | 0.695±0.032 | 0.942±0.046 | 0.807±0.101 | 0.765±0.037 | 0.929±0.053 | 0.794±0.101 | 0.695±0.032 | 0.942±0.046 | 0.807±0.101 | | |
| connect-4 | 0.875 | 0.56±0.061 | 0.424±0.025 | 0.543±0.043 | 0.952±0.041 | 0.771±0.039 | 0.934±0.052 | 0.535±0.043 | 0.424±0.025 | 0.952±0.041 | 0.543±0.043 | | |
| contract | 0.573 | 0.568±0.032 | 0.761±0.012 | 0.567±0.026 | 0.964±0.023 | 0.947±0.009 | 0.943±0.032 | 0.555±0.026 | 0.761±0.012 | 0.964±0.023 | 0.567±0.026 | | |
| credit-approval | 0.79 | 0.803±0.032 | 0.925±0.013 | 0.811±0.017 | 0.971±0.013 | 0.988±0.003 | 0.951±0.038 | 0.84±0.072 | 0.917±0.015 | 0.972±0.013 | 0.811±0.02 | | |
| cylinder-bands | 0.777 | 0.736±0.079 | 0.56±0.022 | 0.936±0.051 | 0.936±0.051 | 0.649±0.032 | 0.921±0.052 | 0.701±0.06 | 0.56±0.022 | 0.936±0.051 | 0.71±0.058 | | |
| echocardiogram | 0.815 | 0.831±0.049 | 0.911±0.009 | 0.964±0.03 | 0.964±0.03 | 0.937±0.013 | 0.963±0.056 | 0.826±0.111 | 0.911±0.009 | 0.964±0.03 | 0.831±0.034 | | |
| energy-y1 | 0.974 | 0.962±0.046 | 0.621±0.015 | 0.961±0.016 | 0.961±0.016 | 0.937±0.009 | 0.932±0.028 | 0.924±0.048 | 0.621±0.015 | 0.961±0.016 | 0.957±0.038 | | |
| energy-y2 | 0.922 | 0.899±0.042 | 0.745±0.014 | 0.905±0.041 | 0.966±0.019 | 0.982±0.004 | 0.904±0.054 | 0.854±0.15 | 0.745±0.014 | 0.966±0.019 | 0.905±0.041 | | |
| fertility | 0.9 | 0.91±0.034 | 0.806±0.018 | 0.908±0.042 | 0.966±0.032 | 0.987±0.009 | 0.881±0.195 | 0.88±0.286 | 0.749±0.026 | 0.968±0.028 | 0.908±0.041 | | |
| haberman-survival | 0.613 | 0.651±0.062 | 0.867±0.017 | 0.648±0.067 | 0.978±0.004 | 0.964±0.005 | 0.865±0.159 | 0.623±0.136 | 0.756±0.029 | 0.94±0.047 | 0.647±0.051 | | |
| heart-hungarian | 0.763 | 0.769±0.057 | 0.639±0.023 | 0.958±0.036 | 0.958±0.036 | 0.978±0.006 | 0.909±0.087 | 0.741±0.094 | 0.576±0.02 | 0.96±0.037 | 0.769±0.051 | | |
| hepatitis | 0.742 | 0.811±0.058 | 0.982±0.025 | 0.954±0.101 | 0.954±0.101 | 0.982±0.025 | 0.954±0.101 | 0.813±0.164 | 0.982±0.025 | 0.954±0.101 | 0.813±0.164 | | |
| lipd-indian-liver | 0.615 | 0.682±0.049 | 0.535±0.02 | 0.955±0.047 | 0.955±0.047 | 0.979±0.004 | 0.891±0.096 | 0.656±0.087 | 0.535±0.02 | 0.955±0.047 | 0.679±0.049 | | |
| ionosphere | 0.944 | 0.932±0.127 | 0.449±0.037 | 0.934±0.051 | 0.934±0.051 | 0.692±0.036 | 0.906±0.062 | 0.892±0.126 | 0.449±0.037 | 0.934±0.051 | 0.914±0.12 | | |
| iris | 0.933 | 0.962±0.064 | 0.825±0.013 | 0.967±0.024 | 0.967±0.024 | 0.964±0.005 | 0.931±0.058 | 0.94±0.113 | 0.898±0.014 | 0.967±0.029 | 0.906±0.059 | | |
| magic | 0.801 | 0.804±0.022 | 0.777±0.014 | 0.803±0.004 | 0.978±0.004 | 0.991±0.003 | 0.953±0.014 | 0.795±0.096 | 0.777±0.014 | 0.978±0.004 | 0.804±0.01 | | |
| manograph | 0.803 | 0.823±0.036 | 0.641±0.027 | 0.804±0.06 | 0.962±0.04 | 0.81±0.02 | 0.942±0.046 | 0.791±0.065 | 0.641±0.027 | 0.962±0.04 | 0.804±0.06 | | |
| miniboone | 0.936 | 0.736±0.098 | 0.523±0.023 | 0.942±0.103 | 0.942±0.103 | 0.523±0.027 | 0.942±0.103 | 0.699±0.096 | 0.523±0.027 | 0.942±0.103 | 0.699±0.096 | | |
| molec-biol-splice | 0.944 | 0.59±0.016 | 0.543±0.031 | 0.936±0.021 | 0.936±0.021 | 0.936±0.031 | 0.921±0.024 | 0.565±0.02 | 0.543±0.031 | 0.936±0.021 | 0.583±0.019 | | |
| mushroom | 0.979 | 0.962±0.09 | 0.711±0.022 | 0.902±0.05 | 0.902±0.05 | 0.761±0.022 | 0.898±0.052 | 0.919±0.07 | 0.711±0.022 | 0.902±0.05 | 0.93±0.072 | | |
| muskm-1 | 0.812 | 0.745±0.065 | 0.773±0.041 | 0.927±0.092 | 0.927±0.092 | 0.65±0.09 | 0.927±0.092 | 0.65±0.09 | 0.773±0.041 | 0.927±0.092 | 0.65±0.09 | | |
| muskm-2 | 0.958 | 0.805±0.1 | 0.979±0.003 | 0.838±0.241 | 0.838±0.241 | 0.979±0.003 | 0.838±0.241 | 0.718±0.228 | 0.979±0.003 | 0.838±0.241 | 0.718±0.228 | | |
| oocytes.merluccius.nucleus.4d | 0.771 | 0.579±0.096 | 0.728±0.03 | 0.903±0.129 | 0.903±0.129 | 0.985±0.003 | 0.836±0.166 | 0.572±0.138 | 0.728±0.03 | 0.903±0.129 | 0.596±0.105 | | |
| oocytes.trisopterus.nucleus.2f | 0.803 | 0.741±0.061 | 0.714±0.022 | 0.89±0.054 | 0.89±0.054 | 0.763±0.02 | 0.89±0.057 | 0.677±0.051 | 0.714±0.022 | 0.89±0.054 | 0.683±0.054 | | |
| parkinsons | 0.923 | 0.878±0.114 | 0.751±0.023 | 0.88±0.089 | 0.88±0.089 | 0.86±0.096 | 0.877±0.081 | 0.837±0.102 | 0.693±0.024 | 0.882±0.087 | 0.851±0.096 | | |
| pima | 0.721 | 0.722±0.023 | 0.989±0.003 | 0.968±0.024 | 0.968±0.024 | 0.99±0.004 | 0.958±0.032 | 0.723±0.037 | 0.947±0.008 | 0.969±0.017 | 0.725±0.02 | | |
| pittsburg-bridges-MATERIAL | 0.909 | 0.909±0.027 | 0.984±0.005 | 0.97±0.021 | 0.97±0.021 | 0.911±0.054 | 0.934±0.076 | 0.898±0.172 | 0.99±0.043 | 0.977±0.039 | 0.909±0.081 | | |
| pittsburg-bridges-REL-L | 0.667 | 0.671±0.033 | 0.99±0.003 | 0.963±0.031 | 0.963±0.031 | 0.679±0.061 | 0.95±0.058 | 0.679±0.076 | 0.99±0.003 | 0.963±0.031 | 0.679±0.061 | | |
| pittsburg-bridges-T-OR-D | 0.857 | 0.857±0.0 | 0.994±0.006 | 0.975±0.002 | 0.975±0.002 | 0.862±0.021 | 0.975±0.001 | 0.86±0.011 | 0.994±0.006 | 0.975±0.002 | 0.862±0.021 | | |
| planning | 0.703 | 0.703±0.0 | 0.994±0.025 | 0.98±0.003 | 0.98±0.003 | 0.994±0.025 | 0.98±0.003 | 0.703±0.018 | 0.994±0.025 | 0.98±0.003 | 0.703±0.018 | | |
| ringnorm | 0.983 | 0.905±0.037 | 0.687±0.029 | 0.818±0.026 | 0.818±0.026 | 0.924±0.035 | 0.798±0.033 | 0.862±0.04 | 0.687±0.029 | 0.924±0.035 | 0.924±0.035 | | |
| seeds | 0.881 | 0.87±0.113 | 0.85±0.012 | 0.971±0.021 | 0.971±0.021 | 0.87±0.069 | 0.913±0.094 | 0.849±0.13 | 0.76±0.011 | 0.972±0.023 | 0.869±0.069 | | |
| spambase | 0.926 | 0.923±0.032 | 0.741±0.035 | 0.963±0.009 | 0.963±0.009 | 0.919±0.024 | 0.925±0.048 | 0.893±0.119 | 0.741±0.035 | 0.963±0.009 | 0.919±0.024 | | |
| statlog-australian-credit | 0.681 | 0.681±0.0 | 0.995±0.023 | 0.98±0.0 | 0.98±0.0 | 0.995±0.023 | 0.98±0.0 | 0.681±0.0 | 0.995±0.023 | 0.98±0.0 | 0.681±0.0 | | |
| statlog-german-credit | 0.765 | 0.723±0.016 | 0.952±0.008 | 0.967±0.019 | 0.967±0.019 | 0.73±0.02 | 0.922±0.075 | 0.708±0.131 | 0.966±0.007 | 0.969±0.021 | 0.73±0.02 | | |
| statlog-heart | 0.852 | 0.848±0.02 | 0.836±0.011 | 0.972±0.013 | 0.972±0.013 | 0.849±0.022 | 0.933±0.036 | 0.822±0.05 | 0.836±0.011 | 0.972±0.013 | 0.849±0.022 | | |
| statlog-image | 0.952 | 0.505±0.067 | 0.768±0.022 | 0.851±0.073 | 0.851±0.073 | 0.596±0.062 | 0.735±0.103 | 0.501±0.085 | 0.768±0.022 | 0.851±0.073 | 0.596±0.062 | | |
| statlog-vehicle | 0.765 | 0.629±0.068 | 0.291±0.012 | 0.912±0.078 | 0.912±0.078 | 0.415±0.014 | 0.889±0.069 | 0.599±0.069 | 0.291±0.012 | 0.912±0.078 | 0.621±0.07 | | |
| synthetic-control | 1.0 | 0.693±0.083 | 0.878±0.024 | 0.874±0.06 | 0.874±0.06 | 0.722±0.08 | 0.827±0.076 | 0.722±0.08 | 0.853±0.029 | 0.881±0.06 | 0.714±0.072 | | |
| teaching | 0.548 | 0.597±0.095 | 0.296±0.018 | 0.89±0.112 | 0.89±0.112 | 0.619±0.095 | 0.602±0.126 | 0.568±0.106 | 0.296±0.018 | 0.89±0.112 | 0.619±0.095 | | |
| tic-tac-toe | 0.974 | 0.876±0.035 | 0.525±0.021 | 0.918±0.023 | 0.918±0.023 | 0.891±0.035 | 0.803±0.02 | 0.884±0.03 | 0.525±0.021 | 0.918±0.023 | 0.891±0.035 | | |
| titanic | 0.778 | 0.778±0.021 | 0.915±0.021 | 0.976±0.009 | 0.976±0.009 | 0.774±0.009 | 0.976±0.009 | 0.769±0.151 | 0.915±0.021 | 0.976±0.009 | 0.774±0.009 | | |
| twonorm | 0.976 | 0.974±0.029 | 0.59±0.02 | 0.977±0.004 | 0.977±0.004 | 0.982±0.003 | 0.94±0.035 | 0.946±0.062 | 0.59±0.02 | 0.977±0.004 | 0.97±0.013 | | |
| vertebral-column-2classes | 0.839 | 0.847±0.066 | 0.796±0.014 | 0.962±0.027 | 0.962±0.027 | 0.855±0.042 | 0.889±0.102 | 0.821±0.128 | 0.796±0.014 | 0.962±0.027 | 0.855±0.042 | | |
| vertebral-column-3classes | 0.806 | 0.825±0.063 | 0.713±0.018 | 0.965±0.039 | 0.965±0.039 | 0.979±0.004 | 0.903±0.066 | 0.799±0.074 | 0.713±0.018 | 0.965±0.039 | 0.819±0.053 | | |
| waveform-noise | 0.843 | 0.841±0.044 | 0.361±0.016 | 0.964±0.013 | 0.964±0.013 | 0.361±0.016 | 0.936±0.035 | 0.802±0.035 | 0.361±0.016 | 0.964±0.013 | 0.825±0.035 | | |
| waveform | 0.843 | 0.833±0.033 | 0.428±0.015 | 0.971±0.013 | 0.971±0.013 | 0.825±0.028 | 0.932±0.023 | 0.796±0.028 | 0.428±0.015 | 0.971±0.013 | 0.825±0.028 | | |
| wine | 0.944 | 0.943±0.098 | 0.743±0.016 | 0.962±0.036 | 0.962±0.036 | 0.926±0.084 | 0.924±0.055 | 0.94±0.084 | 0.811±0.02 | 0.965±0.039 | 0.925±0.084 | | |
| average | 0.839 | 0.800 | 0.716 | 0.942 | 0.942 | 0.794 | 0.913 | 0.776 | 0.701 | 0.944 | 0.793 | | |

Table 4: Results of the copying process for the single-pass and sequential approaches.

The operational points discussed are graphically represented in Figure 8. Figures 8a), b) and c) mark the *Best accuracy* operational point with a circle, the *Best efficiency* operational point with a square and the *Best convergence* operational point with a star. Each dataset is linked by a line connecting the two points, which is a linear fit of the intermediate solutions. Triangles mark the datasets where there is no statistically significant difference between the considered operational points. Figure 8a) displays that, in most cases, the method’s efficiency can be largely increased with only a small degradation in accuracy, as indicated by the relatively flat slopes. Figure 8b) confirms the results, with the *Best accuracy* and *Best convergence* operational points tending to be very similar, as evidenced by the high number of triangles. The method thus showcases fast convergence speed and high accuracy simultaneously. Finally, Figure 8c) displays convergence against efficiency, with the *Best convergence* operational point marked with a circle. A dashed line links this point to the *Best efficiency* operational point in the same dataset. We observe that there is a larger slope in the lines, which indicates a trade-off between the number of data points used and the speed of convergence. This is to be expected, because the largest the number of points used in the training, the smaller the amount of iterations the system will probably require. However, the method still shows fast convergence, ranging between 85% and 98%. This indicates that the method is not only fast but it also requires very small amount of points.

5. Conclusions

In this paper, we proposed a sequential approach for replicating the decision behavior of a machine learning model through copying. Our approach offers a unique solution to the problem of balancing memory requirements and convergence speed and is the first to tackle the problem modeled in Equation 1 in the context of copying. To this end, we moved the copying problem to a probabilistic setting and introduced two theorems to demonstrate that the sequential approach converges to the single-pass approach as the number of samples used for copying increases.

We also studied the duality of compression and memorization in the copy model and showed that a perfect copy can compress all the data in the model parameters. To assess this effect, we employed epistemic uncertainty as a reliable measure of data compression for copying. This measure is only valid for copies and can not be extrapolated to standard learning procedures. This is because, contrary to the standard learning case, there is no aleatoric uncertainty when copying. Therefore, all measured uncertainty corresponds to that coming from the model itself. As such, we can devise copy models that effectively compress all the data with guarantees. With this in mind, we identified the phenomenon of catastrophic forgetting in copies, a well-known effect that appears in online learning processes. To mitigate this effect, we introduced a regularization term derived from an invariant in one of the theorems that enable the process to become more stable.

To reduce the computational time and memory used, we introduced a sample selection policy. This policy controls the compression level that each data sample undertakes. If new data points are already well compressed and represented by the copy at the considered iteration, it is unnecessary to feed them back to this model. As a result of this process, very few data are required during the learning process. Moreover, we observed that the number of samples required to converge to an optimal solution stabilizes at a certain amount.

Additionally, we introduced a regularization term for the copy-loss function to prevent noise during the learning process. This term prevents copies from diverging from one iteration to another. To control the hyperparameter governing this regularization term, we devised an automatic adjustment policy. This policy resorts to a simple but stable meta-learning algorithm that allows us to weigh the dynamic adjustment of the regularization term. As a result, there is no need for hyperparameter tuning.

Our empirical validation on 58 UCI datasets and six different machine learning architectures showed that the sequential approach can create a copy with the same accuracy as the single-pass approach, while offering faster convergence and more efficient use of computational resources. The sequential approach provides a flexible solution for companies to reduce the maintenance costs of machine learning models in production by choosing the most suitable copying setting based on available computational resources for memory and execution time.

Acknowledgments

This work was funded by Huawei Technologies Duesseldorf GmbH under project TC20210924032, and partially supported by MCIN/ AEI/10.13039/501100011033/ under project PID2022-136436NB-I00.

References

- Hessam Bagherinezhad, Maxwell Horton, Mohammad Rastegari, and Ali Farhadi. Label refinery: Improving imagenet classification through label progression. arXiv:1805.02641, 2018. [Online] Available from: <https://arxiv.org/abs/1805.02641>.
- Mariam Barque, Simon Martin, Jérémie Etienne Norbert Vianin, Dominique Genoud, and David Wannier. Improving wind power prediction with retraining machine learning algorithms. In *Proceedings of the International Workshop on Big Data and Information Security, Jakarta, Indonesia*, pages 43–48, Jakarta, Indonesia, 2018. IEEE Computer Society, Washington, DC, USA.
- Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 535–541, Philadelphia, PA, USA, 2006a. Association for Computing Machinery, New York, NY, USA. ISBN 1595933395. URL <https://doi.org/10.1145/1150402.1150464>. doi: 10.1145/1150402.1150464.
- Cristian Buciluă, Rich Caruana, and Alexandru Niculescu-Mizil. Model compression. In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 535–541, Philadelphia, PA, USA, 2006b. Association for Computing Machinery, New York, NY, USA. ISBN 1595933395. URL <https://doi.org/10.1145/1150402.1150464>. doi: 10.1145/1150402.1150464.

- Shaohan Chen, Nikolaos V. Sahinidis, and Chuanhou Gao. Transfer learning in information criteria-based feature selection. *Journal of Machine Learning Research*, pages 1–105, 2022.
- Shuxiao Chen, Edgar Dobriban, and Jane H. Lee. A group-theoretic framework for data augmentation. *Journal of Machine Learning Research*, 21:1–71, 2020.
- Tao Chen. All versus one: An empirical comparison on retrained and incremental machine learning for modeling performance of adaptable software. In *2019 IEEE/ACM 14th International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*, pages 157–168, 2019. doi: 10.1109/SEAMS.2019.00029.
- Yuansi Chen and Peter Bühlmann. Domain adaptation under structural causal models. *Journal of Machine Learning Research*, pages 1–80, 2021.
- Terrance DeVries and Graham W Taylor. Learning confidence for out-of-distribution detection in neural networks. *arXiv preprint arXiv:1802.04865*, 2018.
- Dua Dheeru and Efi Karra Taniskidou. UCI Machine Learning Repository, 2017.
- Leo L. Duan, James E. Johndrow, and David B. Dunson. Scaling up data augmentation mcmc via calibration. *Journal of Machine Learning Research*, 19:1–34, 2018.
- Javad Feizabadi. Machine learning demand forecasting and supply chain performance. *International Journal of Logistics Research and Applications*, 25(2):119–142, 2022. doi: 10.1080/13675567.2020.1803246. URL <https://doi.org/10.1080/13675567.2020.1803246>.
- Zhabiz Gharibshah and Xingquan Zhu. User response prediction in online advertising. *ACM Comput. Surv.*, 54(3), may 2021. doi: 10.1145/3446662. URL <https://doi.org/10.1145/3446662>.
- Jon Ander Gómez, Juan Arévalo, Roberto Paredes, and Jordi Nin. End-to-end neural network architecture for fraud scoring in card payments. *Pattern Recognition Letters*, 105: 175–181, 2018. ISSN 0167-8655. doi: <https://doi.org/10.1016/j.patrec.2017.08.024>. URL <https://www.sciencedirect.com/science/article/pii/S016786551730291X>. Machine Learning and Applications in Artificial Intelligence.
- Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015a.
- Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In *NIPS Deep Learning and Representation Learning Workshop*, 2015b.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, March 2017. doi: 10.1073/pnas.1611835114. URL <https://doi.org/10.1073/pnas.1611835114>.

- Ruishan Liu, Nicolo Fusi, and Lester Mackey. Teacher-student compression with generative adversarial networks. arXiv:1812.02271, 2018. [Online] Available from: <https://arxiv.org/pdf/1812.02271.pdf>.
- Jose Mena, Axel Brando, Oriol Pujol, and Jordi Vitrià. Uncertainty estimation for black-box classification models: a use case for sentiment analysis. In *Pattern Recognition and Image Analysis. IbPRIA 2019. Lecture Notes in Computer Science*, volume 11867, pages 29–40, Madrid, Spain, 2019. ISBN 978-3-030-31331-9. doi: 10.1007/978-3-030-31332-6_3.
- Jose Mena, Oriol Pujol, and Jordi Vitrià. Uncertainty-based rejection wrappers for black-box classifiers. *IEEE Access*, 8:101721–101746, 2020. doi: 10.1109/ACCESS.2020.2996495.
- Rafael Müller, Simon Kornblith, and Geoffrey Hinton. When does label smoothing help? In *Proceedings of the 33rd Conference on Neural Information Processing Systems*, Vancouver, Canada, 2019. MIT Press, Cambridge, MA, USA.
- Vu-Linh Nguyen, Sébastien Destercke, Marie-Hélène Masson, and Eyke Hüllermeier. Reliable multi-class classification based on pairwise epistemic and aleatoric uncertainty. In *27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pages 5089–5095, 2018.
- Andrei Paleyes, Raoul-Gabriel Urma, and Neil D. Lawrence. Challenges in deploying machine learning: A survey of case studies. *ACM Computing Surveys*, apr 2022. ISSN 0360-0300. doi: 10.1145/3533378. URL <https://doi.org/10.1145/3533378>.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions in Knowledge Data Engineering*, 22:1345–1359, 2010.
- Colin Raffel, Noam Shazeer, Adam Roberts, Sharan Lee, Katherine and Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learnign Research*, pages 1–68, 2020.
- Robin Senge, Stefan Bösner, Krzysztof Dembczyński, Jörg Haasenritter, Oliver Hirsch, Norbert Donner-Banzhoff, and Eyke Hüllermeier. Reliable classification: Learning classifiers that distinguish aleatoric and epistemic uncertainty. *Information Sciences*, 255: 16–29, 2014. ISSN 0020-0255. doi: <https://doi.org/10.1016/j.ins.2013.07.030>. URL <https://www.sciencedirect.com/science/article/pii/S0020025513005410>.
- Mohammad Shehab, Laith Abualigah, Qusai Shambour, Muhannad A. Abu-Hashem, Mohd Khaled Yousef Shambour, Ahmed Izzat Alsalibi, and Amir H. Gandomi. Machine learning in medical applications: A review of state-of-the-art methods. *Computers in Biology and Medicine*, 145:105458, 2022. ISSN 0010-4825. doi: <https://doi.org/10.1016/j.combiomed.2022.105458>. URL <https://www.sciencedirect.com/science/article/pii/S0010482522002505>.
- Xiaomu Song, Guoliang Fan, and Mahesh Rao. SVM-Based data editing for enhanced one-class classification of remotely sensed imagery. *IEEE Geoscience Remote Sensors Letters*, 5:189–193, 2008.

- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, Las Vegas, NV, USA, 2016. IEEE Computer Society, Washington, DC, USA. doi: 10.1109/CVPR.2016.308.
- Irene Unceta, Jordi Nin, and Oriol Pujol. Copying machine learning classifiers. *IEEE Access*, 8(11):160268–160284, 2020a. doi = 10.1109/ACCESS.2020.3020638.
- Irene Unceta, Jordi Nin, and Oriol Pujol. Environmental adaptation and differential replication in machine learning. *Entropy*, 22(10), 2020b. doi = 10.3390/e22101122.
- Irene Unceta, Diego Palacios, Jordi Nin, and Oriol Pujol. Sampling unknown decision functions to build classifier copies. In *Proceedings of 17th International Conference on Modeling Decisions for Artificial Intelligence*, pages 192–204, Sant Cugat, Spain, 2020c. doi: 10.1007/978-3-030-57524-3_16.
- Dongdong Wang, Yandong Li, Liqiang Wang, and Boqing Gong. Neural networks are more productive teachers than human raters: active mixup for data-efficient knowledge distillation from a blackbox model. In *Proceedings of the 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1498–1507, Seattle, WV, USA, 2020. IEEE Computer Society, Washington, DC, USA.
- Yinjun Wu, Edgar Dobriban, and Susan Davidson. DeltaGrad: Rapid retraining of machine learning models. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 10355–10366. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/wu20b.html>.
- Chenglin Yang, Lingxi Xie, Siyuan Qiao, and Alan L. Yuille. Training deep neural networks in generations: A more tolerant teacher educates better students. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):5628–5635, Jul. 2019. doi: 10.1609/aaai.v33i01.33015628. URL <https://ojs.aaai.org/index.php/AAAI/article/view/4506>.
- Li Yuan, Francis E.H. Tay, Guilin Li, Tao Wang, and Jiashi Feng. Revisiting knowledge distillation via label smoothing regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*. IEEE Computer Society, Washington, DC, USA, 2020.
- Xinchuan Zeng and Tony R. Martinez. Using a Neural Network to Approximate an Ensemble of Classifiers. *Neural Processing Letters*, 12:225–237, 2000.

Appendix A. Proof of convergence

We establish the equivalence between solving the optimization problem in Equation 8 and its counterpart as i approaches infinity. To do this, we demonstrate the convergence of the function sequence F_i to $F(\theta)$ for increasing values of i , *i.e.*, $\lim_{i \rightarrow \infty} F_i(\theta) = F(\theta)$. Subsequently, we prove that this convergence implies that the sequence θ_i^* converges to θ^* as i approaches infinity. We do so under several assumptions.

For the sake of completeness, let us introduce the following proposition on the uniform convergence of functions.

Proposition 1: Cauchy convergence criterion A sequence of functions $f_n : X \subseteq \mathbb{R}^n \rightarrow \mathbb{R}$ is uniformly convergent if and only if, for every $\epsilon > 0$, there exists an N such that for all $m, n \geq N$, $|f_n(x) - f_m(x)| < \epsilon$ for all $x \in X$.

Now, let us prove Theorem 1 from the main text:

Theorem 1 Let $S_i \subseteq S_{i+1} \subseteq \dots \subseteq S$ be a sequence of subsets converging to S with $|S_i| = i \cdot M$, $M \in \mathbb{N}$. Then, the sequence of functions $\{F_i\}_i$ defined as $F_i(\theta) = \frac{1}{|S_i|} \sum_{z \in S_i} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z))$, converges uniformly to $F(\theta) = \frac{1}{|S|} \sum_{z \in S} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z))$, where $\ell(\cdot)$ is a non-negative and bounded loss function, $|\ell(\cdot)| \leq \mathcal{K}$.

Proof Let us consider the difference between two consecutive steps

$$\begin{aligned} |F_i(\theta) - F_{i-1}(\theta)| &= \left| \frac{1}{|S_i|} \sum_{z \in S_i} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) - \frac{1}{|S_{i-1}|} \sum_{z \in S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right| \\ &= \left| \frac{1}{|S_i|} \sum_{z \in S_i} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) - \frac{1}{|S_{i-1}|} \sum_{z \in S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right. \\ &\quad \left. + \frac{1}{|S_i|} \sum_{z \in S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) - \frac{1}{|S_i|} \sum_{z \in S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right| \\ &= \left| \frac{1}{|S_i|} \sum_{z \in S_i \setminus S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) + \left(\frac{1}{|S_i|} - \frac{1}{|S_{i-1}|} \right) \sum_{z \in S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right| \end{aligned}$$

then, as $|S_i| \geq |S_{i-1}|$,

$$\begin{aligned} |F_i(\theta) - F_{i-1}(\theta)| &\leq \left| \frac{1}{|S_i|} \sum_{z \in S_i \setminus S_{i-1}} \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right| \\ &\leq \frac{1}{|S_i|} \sum_{z \in S_i \setminus S_{i-1}} \left| \ell(\mathcal{f}_C(z, \theta), \mathcal{f}_O(z)) \right| \\ &\leq \frac{|S_i \setminus S_{i-1}| \cdot \mathcal{K}}{|S_i|}. \end{aligned}$$

Now, given an $\epsilon > 0$, we need to show that

$$\exists N \text{ such that if } i > N \text{ then } \frac{|S_i \setminus S_{i-1}| \cdot \mathcal{K}}{|S_i|} < \epsilon,$$

however, this last step will strongly depend on the growth policy when generating new samples. In order to fulfil the above requirement, the order of $|S_i \setminus S_{i-1}|$ should be smaller than $|S_i|$. In this work, we limit ourselves to study the growth policy stated in the theorem

$$|S_i| = i \cdot M, \quad M \in \mathbb{N}.$$

Then, we can rewrite

$$\begin{aligned} \frac{|S_i \setminus S_{i-1}|}{|S_i|} &= \frac{iM - (i-1)M}{iM} \\ &= \frac{1}{i} < \frac{\epsilon}{\mathcal{K}} \end{aligned}$$

with $N \in \mathbb{N}$ such that $N > \frac{\mathcal{K}}{\epsilon} + 1$, and choosing $i > N$, then the uniform convergence of $F_i(\theta)$ is proved by the Cauchy convergence criterion. \blacksquare

This growth policy is a sufficient condition for proving the theorem but it is not the only one that may be used. It can be easily proved that any policy of the form $|S_i| = i^k \cdot M$ with $M \in \mathbb{N}$ also works, whereas recursive policies of the form $|S_i| = i \cdot |S_{i-1}|$ will not. Additionally, any loss function may easily fulfill the requirements of the theorem. For example, using a sigmoid function of the loss.

As a consequence of this uniform convergence, two properties of the function F naturally arise. Firstly, $F_i(\theta)$ converges point-wise to $F(\theta)$. Secondly, $F(\theta)$ is a continuous function on the complete set of parameters, Θ .

Parameter convergence

Theorem 2 Under the conditions of Theorem 1, the sequence of parameters $\{\theta_i^*\}_i$ defined as $\theta_i^* = \arg \min_{\theta \in \Theta} F_i(\theta)$, converges to $\theta^* = \arg \min_{\theta \in \Theta} F(\theta)$, where Θ is the complete set of parameters.

Proof As previously introduced, we can define the optimal copy parameters for a given value of i , and its corresponding function F_i , according to the following equation

$$\theta_i^* = \arg \min_{\theta \in \Theta} F_i(\theta) \tag{17}$$

for Θ the complete parameter set. We assume that this set is only *well defined* if F_i and F have a unique global minimum. This is a commonly made assumption in the literature. In addition, we also assume that Θ is compact.

From the definition of θ_i^* above it follows that

$$F_i(\theta_i^*) \leq F_i(\theta), \quad \forall \theta \in \Theta, \quad \forall i \in \mathbb{N}$$

and using the point-wise convergence of F_i we obtain

$$F(\hat{\theta}^*) \leq F(\theta), \quad \forall \theta \in \Theta \quad (18)$$

where $\hat{\theta}^*$ is the limit of the sequence $(\theta_i^*)_i$. As a consequence of the compactness of Θ and the continuity of F , we can conclude that $\hat{\theta}^*$ must exist. Moreover, given our assumption that Θ is *well defined*, we can conclude that

$$\hat{\theta}^* = \theta^* = \arg \min_{\theta} F(\theta). \quad (19)$$

■

The proof above shows that we can approximate the true optimal parameters by sequentially estimating their value using a sequence of subsets S_i that uniformly converge to S . In other words, it demonstrates the theoretical feasibility of the sequential approach.