

A Comprehensive Empirical Evaluation on Online Continual Learning

Albin Soutif–Cormerais
Computer Vision Center
Universitat Autònoma de Barcelona
Barcelona, Spain
albin@cvc.uab.cat

Andrea Cossu
Scuola Normale Superiore
Pisa, Italy
andrea.cossu@sns.it

Vincenzo Lomonaco
Department of Computer Science
University of Pisa
Pisa, Italy
vincenzo.lomonaco@unipi.it

Antonio Carta
Department of Computer Science
University of Pisa
Pisa, Italy
antonio.carta@unipi.it

Julio Hurtado
Department of Computer Science
University of Pisa
Pisa, Italy
julio.hurtado@di.unipi.it

Joost van de Weijer
Computer Vision Center
Universitat Autònoma de Barcelona
Barcelona, Spain
joost@cvc.uab.cat

Hamed Hemati
University of St. Gallen
Saint-Gall, Switzerland
hamed.hemati@unisg.ch

Abstract

Online continual learning aims to get closer to a live learning experience by learning directly on a stream of data with temporally shifting distribution and by storing a minimum amount of data from that stream. In this empirical evaluation, we evaluate various methods from the literature that tackle online continual learning. More specifically, we focus on the class-incremental setting in the context of image classification, where the learner must learn new classes incrementally from a stream of data. We compare these methods on the Split-CIFAR100 and Split-TinyImagenet benchmarks, and measure their average accuracy, forgetting, stability, and quality of the representations, to evaluate various aspects of the algorithm at the end but also during the whole training period. We find that most methods suffer from stability and underfitting issues. However, the learned representations are comparable to i.i.d. training under the same computational budget. No clear winner emerges from the results and basic experience replay, when properly tuned and implemented, is a

very strong baseline. We release our modular and extensible codebase at https://github.com/AlbinSou/ocl_survey based on the avalanche framework to reproduce our results and encourage future research.

1. Introduction

In recent years, we have witnessed a surge of interest and progress in deep continual learning methodologies. These methods are able to learn continually from a stream of non-stationary data, thereby relaxing the principal assumption of having access to *independent and identically distributed* (i.i.d.) samples, often made in statistical learning [21]. In classic (or batch) continual learning, the common assumption is that the data stream is composed of distinct, explicitly defined *tasks* or *domains*, and that the method can detect the task boundaries or easily switch between domains. However, in many real-world scenarios, the data stream may not have clear task boundaries or domain labels, and the method may need to quickly adapt to changes in the data distribution [3, 23]. Moreover, these methods have access to a batch

of data at each task, normally in the form of a dataset, providing them with a locally i.i.d. access to the data.

Online Continual Learning (OCL) [30] is a more challenging and realistic setting of continual learning, where similar to online learning [35], the method learns from *each arriving data point* in the stream. OCL methods *update the model with a high frequency*, often without access to task labels or boundaries, and with a limited computational and memory budget. Due to the non-stationary nature of the stream, OCL methods need to balance stability and plasticity. Furthermore, since the model is used for inference at each time step (*anytime inference*), the learning algorithm must not suffer from stability issues at any point of the learning process.

State-of-the-art OCL methods use a rehearsal buffer to mitigate forgetting [10]. Several improvements of basic experience replay have been proposed to improve the sampling from the buffer [2, 25], the loss function [6, 5, 33], weights update [9] or the classification layer [33]. A common limitation of these works, and even recent empirical surveys [32] is that they focus only on forgetting and final accuracy. However, OCL methods have several other objectives that should be measured with appropriate metrics.

To encourage progress in Online Continual Learning, in this paper, we provide a *comprehensive empirical evaluation* of OCL methods. We exploit recent proposals that provide better metrics to measure forgetting [41], continual stability [26], and quality of the representations [14, 6, 23]. The experimental results on these metrics highlight the strength and limitations of state-of-the-art methods in OCL.

In this work, we contribute to the body of literature in this area as follows:

- We formally define *Online Continual Learning* (Section 2) and a comprehensive set of metrics (Section 4) to measure the performance across different dimensions: accuracy, forgetting, continual stability, and quality of the latent representations.
- We conduct a comprehensive empirical evaluation on *Class-Incremental* scenarios on two main benchmarks (*Split-CIFAR100* and *Split-TinyImagenet*) evaluating 9 different approaches against 5 metrics. The main findings can be found in the “recommendations box” at the top of this page.
- We release all the code to reproduce our results, compare and easily prototype new OCL strategies. The code has been developed within the Deep Continual Learning Avalanche library for maximum flexibility and reproducibility.

2. Online Continual Learning

In *Online Continual Learning (OCL)*, a model learns from a stream of non-stationary experiences of data. In

a classification problem, at each timestep t a mini-batch $(x_t, y_t) \sim p_t(x, y)$ is available, where p_t is the underlying distribution of data and may change over time. The function $f: \mathbb{R}^n \rightarrow \mathbb{R}^c$ is the prediction function and maps inputs to the unnormalized class probabilities (logits). An OCL algorithm is a function $\mathcal{A}: (x_t, y_t), f_{t-1}, \mathcal{M}_{t-1} \rightarrow f_t, \mathcal{M}_t$, where f_t is the model at time t and $\mathcal{M}_t = \{(x_i, y_i)\}$ is a replay buffer, *i.e.*, a small set of samples from the past stream stored for rehearsal.

Unlike Offline Continual Learning, OCL is a challenging setting where only a few new samples are available at each step, which can be stored in a very limited amount. The following are some properties that characterize OCL setups:

Online. At each timestep only a small mini-batch (x_t, y_t) is available (10 in our experiments).

Task Labels. In a task-aware setting, the models know that samples belong to a set of known tasks and a task label is available to associate each sample to its own task. We assume a task-agnostic setting where task labels are not available.

Task Boundaries. Even in the absence of task labels, many continual learning methods assume knowledge about task boundaries, expecting to know when the data distribution switches to a new task. In a boundary-agnostic setting, this information is not available¹. In this paper, we test both boundary-agnostic and boundary-aware methods.

Anytime Inference. In most OCL applications, models should be able to train but also to perform inference online, after every training step [23]. As a result, methods that require expensive finetuning steps before inference such as GDumb [38] are not considered OCL methods in this paper.

3. Methods

Similar to Offline Continual Learning, in OCL, most state-of-the-art results are obtained by rehearsal-based methods. Because of this, most approaches use rehearsal, which is the main reason we focus on them for our empirical study. Rehearsal-based methods keep a separate memory \mathcal{M} of fixed size to store past samples, updated after each mini-batch. Most rehearsal-based approaches, and all the chosen methods, follow the pseudocode shown in Figure 1. In the following paragraphs, we will describe in detail each line, explaining some methods-specific additions and the main reason for selecting each approach.

¹This is commonly referred to as task-free, and it is a common assumption in OCL. We propose the term boundary-agnostic to highlight that task labels and boundaries are two different kinds of knowledge about the properties of the stream

MAJOR FINDINGS OF OUR PERFORMANCE EVALUATION ON ONLINE CONTINUAL LEARNING

- Good stability does not necessarily transfer to higher accuracy (See Table 2, Figure 3 and Section 6 **Stability**).
- There is no best-performing OCL method across all metrics or memory sizes (See Table 2).
- OCL methods suffer from under-fitting in the common experimental setup (See Figure 3 and Section 6 **Forgetting**).
- Well properly tuned ER is a very competitive baseline obtaining better results than most existing methods (See **memory batch size** discussion 6 and Section 5 **Implementation**).
- The quality of the representation is very close to the one learned on the i.i.d stream, indicating that learning a good classifier is one of main problems. (See Section 6 **Representation quality**).



Figure 1: Pseudocode of replay-based OCL methods. Each method can be obtained from basic experience replay (ER) by modifying one of its fundamental components: sampling, loss, classifier, weight update.

Name	Elements	Year	Boundary
AGEM [9]	Modified Update	2018	
ER [10]	-	2019	
ER + LwF [28]	Distillation Loss	2019	✓
ER-ACE [6]	Modified Cross-Entropy Loss	2021	
MIR [2]	Modified sampling	2019	
SCR [33]	Contrastive Loss, NMC	2021	
RAR [25]	Adversarial Augmentations	2022	
DER++ [5]	Distillation Loss	2020	
GDumb [38]	Offline finetuning on the buffer	2020	✓

Table 1: Summary of methods tried in the survey along with their particularities (release year, access to task boundaries).

Sampling. Usually, each new mini-batch is reused for multiple training passes, each time sampling a different mini-batch from the memory and applying stochastic augmentations to both old and new data. This is justified by the theoretical analysis in [44]. MIR [2] finds the maximally

interfered samples, *i.e.*, those that maximally decrease their loss after an SGD step on the new data, and select those for rehearsal. Instead, RAR [25] generates new samples using targeted adversarial attacks that are designed to be in close proximity to the decision boundary of the classifier.

Loss. Most methods use a supervised loss on both new and memory data, such as the cross-entropy. ER-ACE [6] use different loss on new and old data due to the different nature of the two samples. DER++ [5] uses logits distillation, storing the logits together with the raw data in the memory. While the objective is a knowledge distillation loss, the teacher used to compute the logits depend on the time when the sample was stored. SCR [33] uses a contrastive loss. In CL settings with large batches, it was shown that contrastive losses suffer less from forgetting than supervised ones [13, 31, 19]. However, many contrastive losses require large batch sizes, big and diverse datasets, and more time to converge, which may not be possible in an online setting. Notice that distillation often requires task bound-

aries to know when to update the teacher.

Classifier. Most methods use a linear classifier trained by backpropagation. Another popular choice is the NCM (Nearest-Class-Mean) classifier, which computes a prototype for each class and uses the distance between prototypes to classify at inference time. For example, SCR uses an NCM classifier. Usually, the NCM classifier is only used during inference, while a separate linear classifier is trained via backpropagation during training.

Model update. Most methods use end-to-end backpropagation using both the new and the memory samples. Instead, A-GEM [9] uses the gradient from the memory to impose constraints on the update of the model, exploiting the fact that interference can be measured using the cosine similarity between gradients of different tasks.

Baseline methods. GDumb updates the memory via class-balanced reservoir sampling [11] and, before each inference step, retrains the entire model using only the memory data. GDumb is not an effective online method because it cannot do anytime inference due to the high cost of re-training at each step. However, it is a useful baseline that is surprisingly effective.

4. Metrics

To evaluate the performance of each strategy, we use a comprehensive set of metrics recently introduced in the OCL literature. We have taken special care to include metrics to measure stability and knowledge accumulation. Following [26], we perform continual evaluation after every timestep t from the stream (after training on each mini-batch). In addition, we also evaluate the performance at task boundaries. We indicate the accuracy of a model f at training iteration t on evaluation task E_i with $A(E_i, f_t)$, we denote k the current training task index.

Stability. The *Worst-Case Accuracy* (WC-ACC) [26] provides a trade-off between the accuracy on iteration t and the average minimum accuracy over all tasks learned before the current task T_k . Formally:

$$\text{WC-ACC}_t = \frac{1}{k}A(E_k, f_t) + \left(1 - \frac{1}{k}\right)\text{min-ACC}_{T_k}. \quad (1)$$

The metric *min-ACC* is defined in [26] with respect to the last task T_k and can be computed as:

$$\text{min-ACC}_{T_k} = \frac{1}{k-1} \sum_{i=1}^{k-1} \min_{|T_k| < n \leq t} A(E_i, f_n), \quad (2)$$

where $|T_k|$ represents the last training iteration for task T_k .

Average Anytime Accuracy. The *Average Anytime Accuracy* [6] (AAA), sometimes called Area Under the Curve accuracy [23] (A_{AUC}), is a generalization of the average incremental accuracy [39] to the continual evaluation setting, and is defined as:

$$AAA_t = \frac{1}{t} \sum_{j=1}^t \frac{1}{k} \sum_{i=1}^k A(E_i, f_j), \quad (3)$$

Average Accuracy/Forgetting. Similarly to most of the works in CL, we report the *Average Accuracy* and the *Average Forgetting* as defined by [30]. The Average Accuracy is computed on k tasks from a model at step t by $AA = \frac{1}{k} \sum_{i=1}^k A(E_i, f^t)$. Similarly, the Average Forgetting is computed by $AF = \frac{1}{k} \sum_{i=1}^k A(E_i, f^{|T_i|}) - A(E_i, f^t)$.

Cumulative Accuracy/Forgetting. When applied to class-incremental learning, as done in this paper, the above-described Average Forgetting metric measures both the forgetting and the drop in performance because of the increasingly difficult classification task. Soutif et al. [41] propose a forgetting measure tailored to class-incremental learning called cumulative forgetting. The *Cumulative Accuracy* for a model f^t is computed on the concatenation of all evaluation tasks seen up to task k (E_{Σ}^k), and only considering logits up to the classes seen in task k (C_{Σ}^k). It can be computed as

$$b_k^t = \frac{1}{|E_{\Sigma}^k|} \sum_{x,y \in E_{\Sigma}^k} 1_{y(\arg \max_{c \in C_{\Sigma}^k} f^t(x)_c)}, \quad (4)$$

where $1_y(\hat{y})$ is the indicator function that is 1 if $y = \hat{y}$ and 0 otherwise. We then compute the *Average Cumulative Forgetting* [41] across all tasks, which is simply computed from the Cumulative Accuracy as $F_{\Sigma}^t = \frac{1}{t-1} \sum_{k=1}^{t-1} \max_{i=1, \dots, t} b_k^i - b_k^t$.

Representation quality. In this setup, the model backbone is frozen and a linear classifier is trained on top of it using all the training data seen so far (linear probing) [14], we then report the accuracy obtained with this classifier (Probed Accuracy). This metric allows us to evaluate the quality of the representations computed with the incrementally learned backbone. This metric is computed only at task boundaries.

5. Experimental setup

Benchmarks. We present results on two continual learning classification benchmarks. **Split-Cifar100** is created from the Cifar-100 dataset [24], that contains 50 000 training images and 10 000 test images of size 32x32, equally divided into 100 classes that are refined from 20 super-classes.

Split-TinyImagenet [1] is a reduced version of the Imagenet dataset [16], containing 100 000 training images resized to 64x64 and splitted in 200 classes. We split these datasets into 20 tasks using a random class order (task composition change for each random seed).

Model and training details. In all the experiments, we use a slim version of Resnet18 as done in previous works [30]. We use the SGD optimizer without momentum nor weight decay. We tune the learning rate for each method using hyperparameter selection protocol defined later in this section. Since all of the compared methods make use of a replay buffer, we choose to follow a common protocol for learning that consists in performing several training passes on the same batch of data, using a different batch of memory. This protocol has been used in several works [2, 44, 22]. We choose to always apply input transformations since they have been proven to drastically reduce the overfitting on the buffer samples and to be efficient in combination with several iterations per incoming batch [44]. We use a batch size of 10 for the current data for both benchmarks. Each training batch is then constituted of 10 images from current data and 10 images from the replay buffer, except for SCR which needs to sample a bigger batch from memory (of size 118). The number of training passes on each mini-batch is kept fixed for all methods (not tuned), and set to 3 on Split-Cifar100 and 9 on Split-TinyImagenet. On both of the benchmarks we use random cropping and random horizontal flip as input transformation, except for SCR which uses more advanced transformations. In the main results, we use a fixed memory size of 2000 on Split-Cifar100 and 4000 on Split-TinyImagenet. Additionally, we present results for more extreme amounts of memory (low and high) on Split-Cifar100 in Figure 3.

Hyperparameter selection. In order to make sure every method tried performs at its best, we use a hyperparameter selection mechanism. The constraints of the online setting usually do not allow for efficient hyperparameter selection. In [32], the first few tasks of the training stream are used as validation tasks, which is borrowed from the protocol defined in [9]. We also follow this protocol, using the first 4 tasks (out of 20), to optimize the hyperparameters by looking at the accuracy on the validation set. Contrary to what is done in [9], we do not allow the learner to learn offline on these 4 tasks, but rather put the learner in the setting of online learning. We use the tree-structured Parzen estimator algorithm [4] to guide the search of hyperparameters, by running 200 trials for each method. In the Appendix (See Section. ??) we provide the list of tuned parameters for each method as well as the ranges used and the best values found.

Evaluation. As done in [6, 23], and studied more in depth in [15]. We perform *continual evaluation*, meaning we evaluate the model on held-out validation data (5% of the whole stream) after learning on each new mini-batch. On top of this, we also perform a more classical evaluation on the test stream after training on each task.

Methods. We report results for all methods presented in Section 3. On top of that, we add an i.i.d. reference method, which uses the same methodology as the one of ER but learns on an i.i.d. stream instead of a class-incremental one.

Implementation. All methods were implemented using the Avalanche framework [8], an open source continual learning framework that provides the tools to implement new strategies and benchmarks in continual learning. While some methods were already present, we adapted some of the existing methods and added some new methods to the framework to conduct this study. The implementation of each method is modular, which allows reuse of common components (e.g., reservoir buffers) and highlights the similarities and differences between methods. Despite our efforts to make the comparison as fair as possible, there are some implementation choices that made it difficult to reconcile all methods. We list them in the Appendix and discuss their potential impact on the results.

6. Results

Final Average Accuracy. On **Split-Cifar100**, we found that the final performance of all compared methods is very similar to the one of vanilla replay (ER) (See Table 2). Except for the performance of AGEM, most methods performed quite competitively in the OCL setting (only around 5% away from the i.i.d. reference method), the best one being the introduced baseline combining ER and LwF, but only by a tiny margin. On **Split-TinyImagenet**, the results are similarly close (See Table 2), with the exception of ER-ACE performing better on that benchmark, and RAR and SCR underperforming.

Stability (WC-Acc and AAA). On **Split-Cifar100**, in terms of stability, the results vary much more between each method, and the final performance is not necessarily correlated with the stability of the method. For instance, the final accuracy for MIR is 1% above the final accuracy for SCR, however, SCR has about 9% better WC-Acc than MIR. In Figure 3, we illustrate the difference in stability between SCR and ER on Split-Cifar100 with 2000 memory. In this figure, the accuracy on the previous task drops significantly when shifting tasks in the case of ER, indicating low stability, this is not the case for SCR. In general, we observe that the AAA is moderately correlated to the WC-Acc, even

Method	Split-Cifar100 (20 Tasks)				Split-TinyImagenet (20 Tasks)			
	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow	Probed Acc \uparrow	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow	Probed Acc \uparrow
<i>i.i.d</i>	35.3 \pm 1.5	-	-	45.8 \pm 0.6	26.5 \pm 0.6	-	-	34.3 \pm 0.5
GDumb	18.5 \pm 0.5	-	-	-	13.1 \pm 0.4	-	-	-
AGEM	3.1 \pm 0.2	10.4 \pm 0.6	2.9 \pm 0.3	18.7 \pm 0.8	2.6 \pm 0.2	7.3 \pm 0.5	2.6 \pm 0.2	23.3 \pm 0.6
ER	28.2 \pm 1.2	36.6 \pm 2.0	12.5 \pm 0.6	44.9 \pm 0.9	21.2 \pm 0.6	33.9 \pm 1.7	15.2 \pm 0.5	35.6 \pm 0.6
ER + LwF	30.4 \pm 0.8	39.2 \pm 2.0	15.3 \pm 0.9	44.4 \pm 0.8	22.7 \pm 1.1	34.4 \pm 2.4	17.0 \pm 0.7	33.8 \pm 0.9
MIR	29.4 \pm 1.9	33.1 \pm 3.2	11.6 \pm 1.6	43.4 \pm 0.7	21.3 \pm 0.8	31.0 \pm 1.8	15.2 \pm 0.5	33.0 \pm 0.4
ER-ACE	29.9 \pm 0.6	38.5 \pm 1.8	14.9 \pm 0.9	42.4 \pm 0.6	23.6 \pm 0.7	35.0 \pm 1.5	16.8 \pm 0.7	34.2 \pm 0.3
DER++	29.3 \pm 0.9	37.5 \pm 2.5	13.4 \pm 0.7	44.0 \pm 0.8	22.9 \pm 0.5	34.2 \pm 4.0	16.3 \pm 0.3	31.5 \pm 0.9
RAR	28.2 \pm 1.4	38.2 \pm 1.6	14.9 \pm 0.7	42.3 \pm 0.9	15.7 \pm 0.9	27.8 \pm 2.8	10.1 \pm 0.9	29.8 \pm 0.9
SCR	28.3 \pm 0.8	42.1 \pm 2.1	20.3 \pm 0.4	37.0 \pm 0.3	16.9 \pm 0.4	30.7 \pm 1.5	12.3 \pm 0.5	22.5 \pm 0.4

Table 2: Last step results on Split-Cifar100 (20 Tasks) with 2000 memory (Left) and for Split-TinyImagenet (20 Tasks) with 4000 memory (Right). For each metric, we report the average and standard deviation over 5 seeds

though they are not strongly linked in theory (it is possible to have low WC-Acc and high AAA).

Representation quality. Surprisingly, we find that the probed accuracy for most methods is close to the one of the *i.i.d.* reference method (45.8% on Split-Cifar100 and 34.3% on Split-Tinyimagenet). This suggests that the representation learned by methods on the continual stream is not much worse than the one learned on the *i.i.d.* stream. Therefore, the significant performance difference (in *Acc*) between incremental learning methods and the *i.i.d.* reference is most likely caused by a deterioration of the incrementally learned classifier. In the Appendix, we provide results with 500 and 8000 memory on Split-Cifar100. Even when using 500 memory, we observe a similar conclusion, with the gap to probing augmenting only a bit (from 1% to 2%), showing that only a low amount of memory (5 per class in that case) is efficient to get a decent representation strength. In general, we find that the probed accuracy is not strongly linked with the stability metrics. In Table 2, we see that ER has in both cases the best Probed Accuracy, while it has lower than average stability metrics compared to the other methods.

Forgetting. In Figure 3, we display both the classical forgetting and the cumulative forgetting defined in Section 4 [41]. We see that while the classical forgetting indicates a high amount of forgetting that is increasing across the stream, the cumulative forgetting gives a different picture, indicating that for all methods, some backward transfer is achieved, and this remains quite constant across all tasks. Two curves exhibit slightly distinct behavior, namely the ones of SCR and MIR. These distinctive behaviors can also be observed in Figure 2. In the case of MIR, the average accuracy is initially low, and later increases to match

the one of ER-ACE, resulting in high backward transfer. Whereas for SCR, the accuracy is initially high, but later meets the one of other methods, resulting in neutral forgetting (no forgetting, but no backward transfer). Classical forgetting however is not very relevant in the class-incremental learning setting because it increases consequently to an increase in the difficulty of the task (more and more classes to consider in the classification problem), which makes it hard to interpret, as such, we advise against its use in class-incremental learning (online or not online). In definitive, these forgetting numbers indicate that there is some backward transfer happening, we believe this is mainly due to the fact that the network is underfitted in online learning, due to the low number of training iterations, making it easy to gain additional performance on a task when training on subsequent tasks.

Effect of the memory batch size. As explained in Section 5, the memory batch size is set to the same as the current batch size in our experiments, except for SCR, which requires a higher one. We believe that this difference explains the good performance of SCR in the early training regime (see Figure 2) and in the high memory size setting. We perform additional experiments (See Table ??) where we use the same memory batch size for ER as the one of SCR. On Split-Cifar100 with 8000 memory, changing the memory batch size from 10 to 118 is sufficient to make ER match the performance of SCR (jumping from 34.9% to 43.0% final accuracy). This confirms our belief that this parameter is important to take into account when interpreting results.

Effect of the memory size. In Figure 3, we report the final performance of ER, the *i.i.d.* reference method, and the GDumb baseline when using more extreme (low and high)

memory amounts on Split-Cifar100. When high amounts of memory are used, the GDumb baseline can surpass the performance of the continual learning methods if no special care is given to the memory batch size. This one needs to be adapted in order to obtain the best performances with continual learning methods.

7. Related Work

Existing surveys. Surveys on continual learning have focused on different aspects. Parisi et al. [37] provide a survey on lifelong learning and draw parallels with how biological systems prevent catastrophic forgetting. The survey of Lesort et al. [27] studies continual learning focusing on robotics applications. More recent surveys have focused on popular settings for continual learning. De Lange et al. [15] studies task-incremental learning, and Masana et al. [34] investigates class-incremental learning. More similar to the proposed work in this paper, is the work of Mai et al. [32] who propose an empirical evaluation of several online continual learning methods. However, other than them, our paper aims to compare a variety of competitive replay methods that each use different approaches to tackle the setting of online continual learning (as described in Figure 1). On top of that, we evaluate and compare the performance of each method from different points of view, analyzing both the final performance but also the stability (by evaluating on the validation set after each mini-batch). We also evaluate the learned representation by probing the features with the full dataset at the end of the training, as in [14].

Existing Libraries. This survey contributes to the implementation of multiple methods in the Avalanche Library [8]. Avalanche is an end-to-end library based on Pytorch with the goal of providing a codebase for fast prototyping, training, and reproducible evaluation of continual learning algorithms. Besides this, other libraries have been implemented with different objectives and qualities. SequeL [17] is a new library that focuses on developing methods not only in PyTorch, but also in JAX. It provides a simple interface for running experiments in both. However, the newness of the library and the need to implement the methods in both languages make it difficult to use. On the other hand, Sequoia [36] is a library that attempts to unite as many continual learning settings as possible under a common tree. The root is the most difficult problem to learn, and the leaves and branches are different settings. Lastly, Continuum [18] is a library that focuses mostly on the benchmark aspects of continual learning and provides tools to easily split the datasets and iterate on the resulting tasks.

Additional methods. In addition to the methods implemented, there are other methods proposed in recent years.

In Online Bias Correction [12], the authors explain how experience replay biases the model output towards recent observations. With this, they propose a way to modify the classifier output and mitigate the bias. Following the same idea of reducing the bias, Guo et al. [20] propose OCM based on mutual information maximization. Here, the authors deal with the bias reduction caused by cross entropy and they encourage the preservation of previous knowledge. Another approach that relies on a buffer is Proxy-based Contrastive Replay [29]. Here, the authors propose a way to complement a buffer loss and a contrastive loss. Using a Visual Transformer in conjunction with a focal contrastive learning strategy, Wang et al. [42] suggest mitigating the stability-plasticity dilemma.

8. Conclusions

In this study, we examined the performance of various Online Continual Learning (OCL) methods, focusing on performance, stability, representation quality, and forgetting. Our analysis revealed intriguing insights. Firstly, we found that stability does not always translate to higher accuracy, challenging the notion that a stable model guarantees superior performance in the OCL setting. Additionally, we observed that the quality of the representation learned by continual learning methods does not differ strongly from the one obtained by learning on the i.i.d. stream, indicating that the main challenge faced by continual learning methods is to learn a good classifier. We also found that methods were prone to underfitting in the OCL setting, challenging the common assumption that continual learning methods suffer from forgetting; we here claim that they keep improving their performance on previous tasks as they learn on subsequent tasks. In general, we found all compared methods to perform very similarly to the common Experience Replay (ER) method. We also investigate some small implementation differences and conclude that sometimes small details in implementations can make a method shine using the existing metrics, but that it is often possible to obtain these same results by slightly modifying the baseline ER hyperparameters or implementation details, highlighting the necessity to implement these methods in a unified framework like avalanche so that they can be more fairly compared. Finally, we found that no single OCL method proved to be universally superior across all metrics or memory sizes, highlighting the absence of a one-fits-all solution. Considering these findings, we identify several promising research directions for online continual learning.

We advocate for stronger connections between normal i.i.d. online learning and online continual learning, given their similar representation strengths. As the ER strategy is already common to both settings and proven competitive,

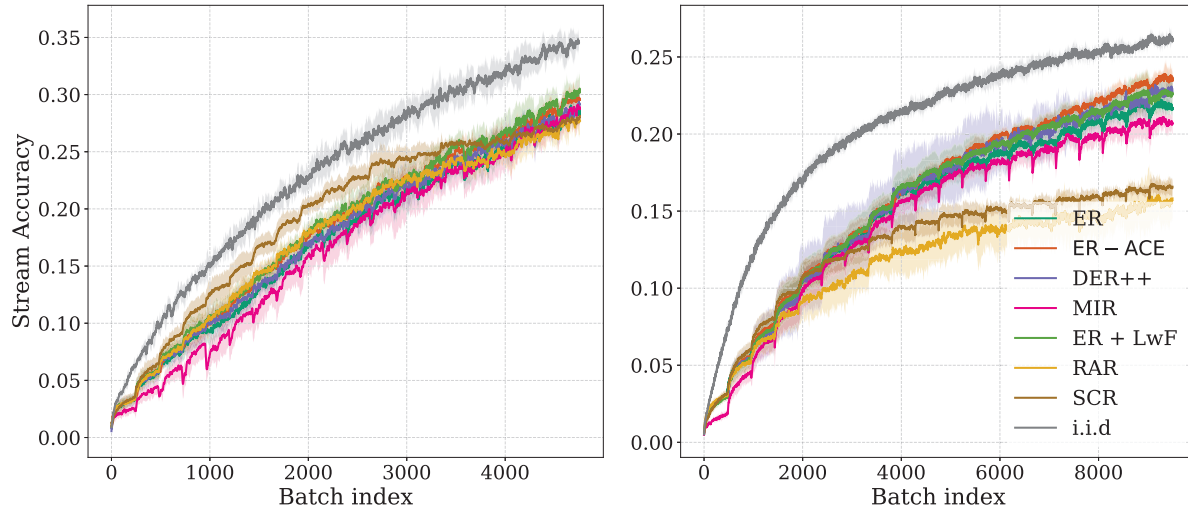


Figure 2: Validation stream accuracy for each of the methods, compared to the one of the i.i.d. reference method, on Split-Cifar100 (20 tasks), using a fixed memory of 2000 exemplars. The accuracy is reported after training on each mini-batch, we display mean and standard deviation across 5 seeds.

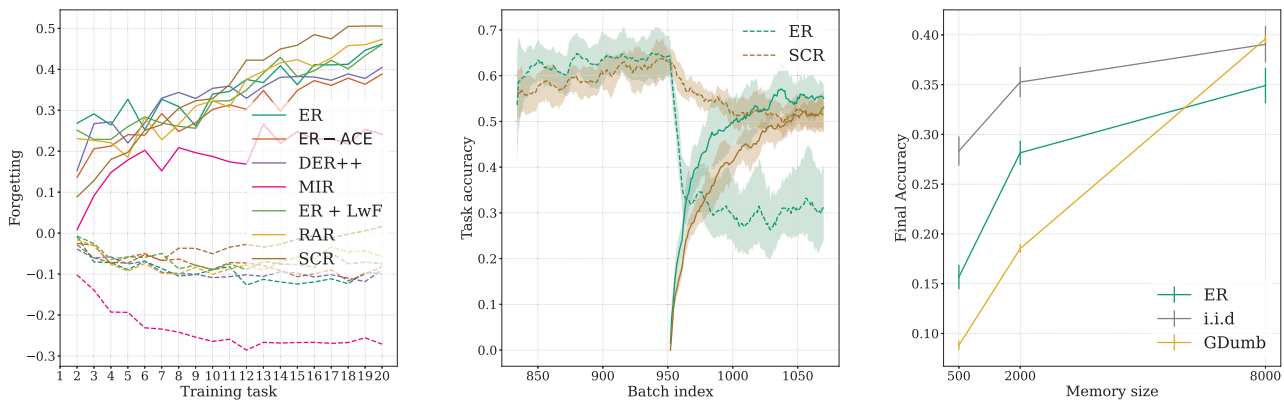


Figure 3: **Left:** Forgetting (full lines), and Cumulative Forgetting (dotted lines) on Split-Cifar100 with 2000 memory; **Middle:** Illustration of the difference in stability between ER and SCR on Split-Cifar100 (20 tasks), using 2000 memory. We place ourselves at the task shift between task 4 and 5 and display the accuracy on previous task data (dotted lines) as well as the accuracy on current task data (full lines).; **Right:** Final performance of ER, i.i.d. reference method, and GDumb baseline for 3 different memory sizes on Split-Cifar100

emphasis should be put on tuning its hyperparameters during training, as already attempted in [43] and [7]. Proper hyperparameter tuning in online continual learning remains an open challenge. Additionally, we encourage further exploration of linking stability metrics to training efficiency, as we found that poor stability does not necessarily impact final representation strength. If no direct link exists, enforc-

ing good stability during training may not be essential, and ad-hoc methods [40] could be sufficient to achieve desired stability.

Acknowledgements: We acknowledge the support from the Spanish Government funding for projects PID2022-143257NB-I00, TED2021-132513B-I00. We also acknowl-

edge the support from the Italian Ministry of University and Research (MUR) as part of the FSE REACT-EU - PON 2014-2020 “Research and Innovation” resources – Innovation Action - DM MUR 1062/2021

References

- [1] Stanford, “tiny imagenet challenge, cs231n course.”. <https://tiny-imagenet.herokuapp.com/>, 2015. 5
- [2] Rahaf Aljundi, Eugene Belilovsky, Tinne Tuytelaars, Laurent Charlin, Massimo Caccia, Min Lin, and Lucas Page-Caccia. Online continual learning with maximal interfered retrieval. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2, 3, 5
- [3] Rahaf Aljundi, Min Lin, Baptiste Goujaud, and Yoshua Bengio. Gradient based sample selection for online continual learning. *Advances in neural information processing systems*, 32, 2019. 1
- [4] James Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. *Advances in neural information processing systems*, 24, 2011. 5
- [5] Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020. 2, 3
- [6] Lucas Caccia, Rahaf Aljundi, Nader Asadi, Tinne Tuytelaars, Joelle Pineau, and Eugene Belilovsky. New insights on reducing abrupt representation change in online continual learning. In *International Conference on Learning Representations*, 2021. 2, 3, 4, 5
- [7] Zhipeng Cai, Ozan Sener, and Vladlen Koltun. Online continual learning with natural distribution shifts: An empirical study with visual data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8281–8290, 2021. 8
- [8] Antonio Carta, Lorenzo Pellegrini, Andrea Cossu, Hamed Hemati, and Vincenzo Lomonaco. Avalanche: A pytorch library for deep continual learning. *arXiv preprint arXiv:2302.01766*, 2023. 5, 7
- [9] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. In *International Conference on Learning Representations*, 2018. 2, 3, 4, 5
- [10] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc’Aurelio Ranzato. On tiny episodic memories in continual learning. *arXiv preprint arXiv:1902.10486*, 2019. 2, 3
- [11] Aristotelis Chrysakis and Marie-Francine Moens. Online continual learning from imbalanced data. In *International Conference on Machine Learning*, pages 1952–1961. PMLR, 2020. 4
- [12] Aristotelis Chrysakis and Marie-Francine Moens. Online bias correction for task-free continual learning. In *The Eleventh International Conference on Learning Representations*, 2022. 7
- [13] Andrea Cossu, Tinne Tuytelaars, Antonio Carta, Lucia Passaro, Vincenzo Lomonaco, and Davide Bacciu. Continual Pre-Training Mitigates Forgetting in Language and Vision, May 2022. 3
- [14] MohammadReza Davari, Nader Asadi, Sudhir Mudur, Rahaf Aljundi, and Eugene Belilovsky. Probing representation forgetting in supervised and unsupervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16712–16721, 2022. 2, 4, 7
- [15] Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Aleš Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021. 5, 7
- [16] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [17] Nikolaos Dimitriadis, Francois Fleuret, and Pascal Frossard. Sequel: A continual learning library in pytorch and jax. *arXiv preprint arXiv:2304.10857*, 2023. 7
- [18] Arthur Douillard and Timothée Lesort. Continuum: Simple management of complex continual learning scenarios, 2021. 7
- [19] Enrico Fini, Victor G. Turrissi Da Costa, Xavier Alameda-Pineda, Elisa Ricci, Karteek Alahari, and Julien Mairal. Self-Supervised Models are Continual Learners. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9611–9620, June 2022. 3
- [20] Yiduo Guo, Bing Liu, and Dongyan Zhao. Online continual learning through mutual information maximization. In *International Conference on Machine Learning*, pages 8109–8126. PMLR, 2022. 7
- [21] Trevor Hastie, Robert Tibshirani, Jerome H Friedman, and Jerome H Friedman. *The elements of statistical*

- learning: data mining, inference, and prediction*, volume 2. Springer, 2009. 1
- [22] Huiyi Hu, Ang Li, Daniele Calandriello, and Dilan Görür. One pass imagenet. *CoRR*, abs/2111.01956, 2021. 5
- [23] Hyunseo Koh, Dahyun Kim, Jung-Woo Ha, and Jonghyun Choi. Online continual learning on class incremental blurry task configuration with anytime inference. In *International Conference on Learning Representations*, 2022. 1, 2, 4, 5
- [24] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. 4
- [25] Lilly Kumari, Shengjie Wang, Tianyi Zhou, and Jeff A Bilmes. Retrospective adversarial replay for continual learning. *Advances in Neural Information Processing Systems*, 35:28530–28544, 2022. 2, 3
- [26] Matthias De Lange, Gido M van de Ven, and Tinne Tuytelaars. Continual evaluation for lifelong learning: Identifying the stability gap. In *The Eleventh International Conference on Learning Representations*, 2023. 2, 4
- [27] Timothée Lesort, Vincenzo Lomonaco, Andrei Stoian, Davide Maltoni, David Filliat, and Natalia Díaz-Rodríguez. Continual learning for robotics: Definition, framework, learning strategies, opportunities and challenges. *Information fusion*, 58:52–68, 2020. 7
- [28] Zhizhong Li and Derek Hoiem. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017. 3
- [29] Huiwei Lin, Baoquan Zhang, Shanshan Feng, Xutao Li, and Yunming Ye. Pcr: Proxy-based contrastive replay for online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24246–24255, 2023. 7
- [30] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017. 2, 4, 5
- [31] Divyam Madaan, Jaehong Yoon, Yuanchun Li, Yunxin Liu, and Sung Ju Hwang. Representational Continuity for Unsupervised Continual Learning. In *ICLR*, Apr. 2022. 3
- [32] Zheda Mai, Ruiwen Li, Jihwan Jeong, David Quispe, Hyunwoo Kim, and Scott Sanner. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51, 2022. 2, 5, 7
- [33] Zheda Mai, Ruiwen Li, Hyunwoo Kim, and Scott Sanner. Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3589–3599, 2021. 2, 3
- [34] Marc Masana, Xialei Liu, Bartłomiej Twardowski, Mikel Menta, Andrew D Bagdanov, and Joost Van De Weijer. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5):5513–5533, 2022. 7
- [35] Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. MIT press, 2018. 2
- [36] Fabrice Normandin, Florian Golemo, Oleksiy Ostapenko, Pau Rodriguez, Matthew D Riemer, Julio Hurtado, Khimya Khetarpal, Ryan Lindeborg, Lucas Cecchi, Timothée Lesort, et al. Sequoia: A software framework to unify continual learning research. *arXiv preprint arXiv:2108.01005*, 2021. 7
- [37] German I Parisi, Ronald Kemker, Jose L Part, Christopher Kanan, and Stefan Wermter. Continual lifelong learning with neural networks: A review. *Neural networks*, 113:54–71, 2019. 7
- [38] Ameeya Prabhu, Philip HS Torr, and Puneet K Dokania. Gdumb: A simple approach that questions our progress in continual learning. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 524–540. Springer, 2020. 2, 3
- [39] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 2001–2010, 2017. 4
- [40] Albin Soutif-Cormerais, Antonio Carta, and Joost Van de Weijer. Improving online continual learning performance and stability with temporal ensembles, 2023. 8
- [41] Albin Soutif-Cormerais, Marc Masana, Joost Van de Weijer, and Bartłomiej Twardowski. On the importance of cross-task features for class-incremental learning. *arXiv: 2106.11930*, 2021. 2, 4, 6
- [42] Zhen Wang, Liu Liu, Yajing Kong, Jiaxian Guo, and Dacheng Tao. Online continual learning with contrastive vision transformer. In *European Conference on Computer Vision*, pages 631–650. Springer, 2022. 7
- [43] Yaqian Zhang, Eibe Frank, Bernhard Pfahringer, Albert Bifet, Nick Jin Sean Lim, and Alvin Jia. Closed-loop control for online continual learning, 2022. 8
- [44] Yaqian Zhang, Bernhard Pfahringer, Eibe Frank, Albert Bifet, Nick Jin Sean Lim, and Yunzhe Jia. A simple but strong baseline for online continual learning: Repeated augmented rehearsal. *Advances in Neu-*

ral Information Processing Systems, 35:14771–14783,
2022. [3](#), [5](#)