

IMPROVING ONLINE CONTINUAL LEARNING PERFORMANCE AND STABILITY WITH TEMPORAL ENSEMBLES

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

Antonio Carta
Department of Computer Science
University of Pisa
antonio.cart@unipi.it

Joost Van de Weijer
Computer Vision Center
Universitat Autònoma de Barcelona
joost@cvc.uab.es

ABSTRACT

Neural networks are very effective when trained on large datasets for a large number of iterations. However, when they are trained on non-stationary streams of data and in an online fashion, their performance is reduced (1) by the online setup, which limits the availability of data, (2) due to catastrophic forgetting because of the non-stationary nature of the data. Furthermore, several recent works (Caccia et al., 2022; Lange et al., 2023) showed that replay methods used in continual learning suffer from the *stability gap*, encountered when evaluating the model continually (rather than only on task boundaries). In this article, we study the effect of model ensembling as a way to improve performance and stability in online continual learning. We notice that naively ensembling models coming from a variety of training tasks increases the performance in online continual learning considerably. Starting from this observation, and drawing inspirations from semi-supervised learning ensembling methods, we use a lightweight temporal ensemble that computes the exponential moving average of the weights (EMA) at test time, and show that it can drastically increase the performance and stability when used in combination with several methods from the literature.

1 INTRODUCTION

Learning neural networks with backpropagation has been proven capable of good generalization properties even when using overparametrized networks (Krizhevsky et al., 2017). However, these good learning properties mainly occur when the data is provided in an independent and identically distributed manner. When learning on a stream which distribution varies over time, neural networks are known to suffer from *catastrophic forgetting* (McCloskey & Cohen, 1989; Goodfellow et al., 2014; Kirkpatrick et al., 2017), and tend to forget knowledge acquired in previous learning tasks. The field of *continual learning* aims to address this problem. Generally, incremental learning separates the learning into distinct tasks (identified by a task-ID) that are encountered sequentially by the agent. A variety of settings have been introduced in continual learning in order to evaluate several aspects of the continual learning agent; *task-incremental learning* (De Lange et al., 2021; van de Ven & Tolias, 2018), and *class-incremental learning* (Masana et al., 2022; Belouadah et al., 2021) are among the most popular. In this paper, we focus on the more challenging class-incremental setting, where the learner does not have access to the task-ID at inference time.

Model Ensembling, or the aggregation of predictions coming from different models, is a well studied and popular topic, both in the academic literature and in practical applications (Hansen & Salamon, 1990; Perrone & Cooper, 1995; Dietterich, 2000). It is known to improve performance compared to using a single model. The success of ensembling methods has been found to depend on the functional diversity of the members and the efficiency of the resulting ensemble between them (Goodfellow et al., 2016; Wortsman et al., 2021). In continual learning, the model learns the data task by task, being exposed to one task at a time. Therefore, models at different timesteps represent a functionally diverse set of models, each one locally adapted to the current task. Such functional diversity can be exploited by ensembling techniques. In Figure 1, we show ensembling results on two continual learning benchmarks. Here we apply a temporal ensemble of twenty models (chosen from a large number of models saved along the training trajectory). We can observe that ensembling leads to a significant performance improvement (over 40% on both datasets), and that whenever we increase the number of tasks covered by the ensemble, gains improve. Therefore, in this article, we investigate the use of ensembles for continual learning and provide empirical results that show their benefit for continual learning settings. Importantly, we show that when using a practical and memory efficient ensembling method similar or even better results can be obtained.

Evaluation of continually learned agents typically occurs after learning a task. Recently, another way of evaluating has been proposed, coined *continual evaluation* (Caccia et al., 2022; Lange et al., 2023) or *anytime inference* (Koh et al.,

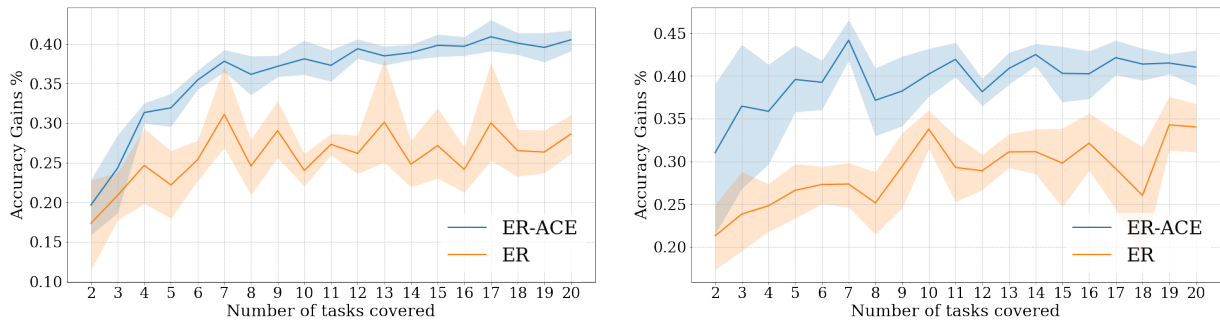


Figure 1: Relative accuracy gains (multiplicative in %), compared to the worst performing ensemble, when naively ensembling 20 models coming from different learning tasks on Split-Cifar100 (Left) and Split-MiniImagenet (Right) for two online continual learning methods, classical replay ER and asymmetric cross entropy loss ER-ACE. Results are reported as a function of *number of covered tasks* which is defined as the number of tasks from which the ensembled models originate. The graph shows that diversity of the ensemble is important.

2022). It aims to evaluate the agent’s performance at any moment during learning. In this setting, Lange et al. (2023) find that continual learning agents suffer from the *stability gap*, where the performance on previous tasks decreases drastically at the start of learning a new task, before returning to normal when continuing training of the new task. This behavior is problematic in many real-world applications where the agent must be applied for inference while it is learning (i.e. for financial market forecasting, or online monitoring tasks). Ensembles can provide improved stability, as they can reduce the variance of the predictions and provide a more robust prediction. By combining multiple models, the errors of one model can be compensated for by the other models in the ensemble, leading to more stable performance.

Lastly, it is known that in class incremental learning, even when using replay, the network is prone to suffer from the task-recency bias, which is a prediction bias towards classes belonging to the last task. This phenomenon has been studied and tackled in many works (Belouadah & Popescu, 2019; Wu et al., 2019; Hou et al., 2019). Ensembling models biased towards different tasks has the potential to reduce the bias compared to the single models; we will analyse this in this paper.

The reasons discussed in the introduction motivate us to investigate the potential of temporal ensembles in online continual learning. Specifically, we believe that they could offer improved stability, reduce task recency bias, and benefit from more functional diversity than in i.i.d learning. By exploring the performance of ensembles of models trained on sequential data, we aim to provide insights into the benefits and limitations of using such models in online continual learning settings. Our contributions are the following

- We show that naively ensembling checkpoints from different continual learning tasks yields strong performance gains for online continual learning. In search for a practical ensembling method, we take inspiration from semi-supervised learning and apply a temporal ensembling method (i.e., an exponential moving average ensemble) as an evaluation model.
- We report the performance increase of the EMA ensemble in combination with several methods, showing a consistent increase in performance (up to 9.3% on Split-MiniImagenet). We also compute continual evaluation metrics and consistently notice a great increase in stability metrics resulting from the use of the EMA ensemble (up to 32.3% on Split-Cifar10). We further observe a reduced task-recency bias of the EMA method.

2 RELATED WORK

2.1 CONTINUAL LEARNING AND ONLINE CONTINUAL LEARNING

Popular continual learning scenarios often assume that data arrive in large batches of i.i.d data, with sharp distribution shifts happening whenever a new batch becomes available. We call this setting boundary-aware continual learning, due to the additional information provided by the arrival of a new task during training. Most of the continual learning literature focus on that setting, either by also providing the task-id at test time (De Lange et al., 2021) (task-incremental learning), or not (Masana et al., 2022) (class-incremental learning). We will focus on class-incremental learning in this paper.

Boundary-free online continual learning (Aljundi et al., 2019b) removes this form of supervision by learning on a stream of small mini-batches. This means that the granularity of the data distribution can be refined. In practice, it is possible to keep tasks that define the granularity of the distribution, while letting the agent assume this distribution could change at any new mini-batch. This is the setting that we experiment on in this article. In MIR (Aljundi et al., 2019a), the authors introduce a selection strategy for replay and select the samples that will be mostly penalised by the current update. In ER-ACE (Caccia et al., 2022), an asymmetric cross-entropy loss is used on current data, using only the logits represented in the current mini-batch, while the classical cross-entropy loss is used on the replay buffer. Zhang et al. (2022) show that a strong baseline (called RAR) in online continual learning is repeatedly training on the available batch by sampling a new replay batch and applying data augmentations. In this article, we propose a simple improvement to these methods based on temporal ensembling which is applied at evaluation time.

Other, more classical class-incremental learning methods can be used in this setting, as long as they don't require knowledge of task-boundaries during training time. In ICaRL (Rebuffi et al., 2017), a selection strategy for storing the buffer sampled is used, along with a nearest mean classifier. In DER (Buzzega et al., 2020), both the samples and the logits of these samples are stored and replayed using data augmentation. Other classes of approaches like EEIL (Castro et al., 2018), perform a balancing step at the end of training on each task, which makes them unsuitable to the boundary-free setting, unless the balancing is done before every evaluation session, in which case it could drastically increase the computation requirements. Other than these methods, we will focus on online continual learning methods in this paper.

2.2 ENSEMBLING IN CONTINUAL LEARNING AND TEMPORAL ENSEMBLING

Aggregating the predictions coming from multiple trained models has been known for a long time as a process that results in increased performance compared to using the predictions of the individual models (Breiman, 1996). The group of trained models used for prediction is referred to as an ensemble of models. Such ensembles have been studied extensively in the literature (Hansen & Salamon, 1990; Perrone & Cooper, 1995; Dietterich, 2000). Several challenges are associated with the creation of such ensembles, and a lot of the requirements of these ensembles seem at first incompatible with the constraints imposed by continual learning. In particular, naively creating an ensemble requires the training of several models that need to be stored in memory and trained independently, thus violating the memory and time constraints of continual learning. Nevertheless, some of these challenges have been already addressed in the literature. Huang et al. (2017) relieve the constraint of having to train separate models by using checkpoints of the same training run and a cyclic learning rate schedule as the ensemble members. Wortsman et al. (2021) train not only one network but a parametric subspace of networks which they can use to create an ensemble.

Wen et al. (2020) develop BatchEnsemble, a memory efficient way to create an ensemble of models by learning a shared weight matrix for all the members, and then a rank one matrix for each of the members. A member is then computed as the result of the hadamard product between the shared matrix and the rank one matrix. They later use this technique in task incremental learning, where they learn one member per task to be used at test time. Doan et al. (2022) lay the grounds of continual learning beyond the use of a single model. They study feasible ways of learning an ensemble of models continually, and compare a variety of ensembling techniques like BatchEnsemble Wen et al. (2020) or Subspace Learning Wortsman et al. (2021). They conclude that ensembling helps in the setting of task-incremental learning, and propose a method that make use of that property to increase the performance in that setting. Compared to our work, they focus more on the effect of adding more models into the ensemble but not so much on the effect of ensembling models coming from different tasks, and they operate in the task-incremental learning setting. Lee et al. (2017) propose *Incremental Moment Matching*, in which they compute the mean of the model weights in the weight space, in turn creating an approximate ensemble. In contrast to our work, they operate in the simpler task-incremental setting, in which task-ID is available at test time.

Temporal ensembling (Samuli & Timo, 2017), is a technique that consists in ensembling the predictions coming from different models on the training trajectory. In the original work, it was done by keeping an exponential moving average of the predictions of the model on the training data, but this technique was later refined in (Tarvainen & Valpola, 2017), where the authors chose to keep a running average of the weights instead of the predictions, and show that this leads to similar or even better performance, while relieving the constraint of having to update the running prediction for each datapoint at every iteration. In both of these works, the resulting ensemble prediction was used to improve the results in semi-supervised learning, where only a small fraction of the sample labels are available. This same Mean-teacher model has also been used successfully in several self-supervised learning works (Grill et al., 2020; Caron et al., 2021). In this article, we study the application of cheap temporal ensembles to the setting of online continual learning.

3 PRELIMINARIES

3.1 CONTINUAL EVALUATION AND THE STABILITY GAP

In continual classification, a learning agent learns the parameters $\theta \in \Theta$ of a function $f : (\mathcal{X}, \Theta) \mapsto \mathcal{Y}$ from the image input space \mathcal{X} to the label space \mathcal{Y} . It does so by observing a stream of data $\mathcal{S} = \{(x^1, y^1), (x^2, y^2), \dots, (x^n, y^n)\}$, where $x \in \mathcal{X}$ and $y \in \mathcal{Y}$. Each data tuple is drawn from a time varying distribution $(x^t, y^t) \sim \mathcal{D}_t$. In classical machine learning the training data distribution does not depend on time, but this is added as a constraint in continual learning. In both cases the goal of the agent is to perform well on new samples drawn from the joint distribution \mathcal{D} , which is marginalized over past time. In practice, continual learning is simplified to allow for easier analysis by studying distributions that come from a discrete set and switch from one distribution to another (referred to as tasks $t \in \{t_1, \dots, t_T\}$). In this paper, we will focus on class-incremental learning, where the learner does not have access to the task identifier t at inference time.

While all of the above simplifications make sense, they are still far from the human learning experience, and from fitting the requirements of many real-world applications. In comparison to the above, humans experience continuously time-varying distributions and continual evaluation. In order to address this, [Caccia et al. \(2022\)](#) and [Lange et al. \(2023\)](#) lay the basis and encourage the study of continual evaluation of neural networks. In continual evaluation, the model is continuously evaluated during, instead of after each task. Interestingly, they noticed that the performance on previous tasks often drops at task shifts before coming back to a higher value later in training, this is what they refer to as the *stability gap*.

3.2 CONTINUAL EVALUATION METRICS

In this section, we present various metrics used in the online continual learning setting and that can help measure the stability and more generally, evaluate the performance of the agent over the course of its training. We denote $\mathbf{A}(E_i, f_t)$ the accuracy of f_t (model at current iteration t), on the evaluation task E_i . The most common metric used in this scenario is the *average anytime accuracy*, AAA_t (See Eq. 1), used in many works ([Caccia et al., 2020](#); [2022](#); [Koh et al., 2022](#)). While this metric does not focus on the worst-case performance, it is a nice indicator of the performance of the learning agent over the course of training. It measures the average accuracy on all tasks seen so far, and averages it over all training iterations. In ([Lange et al., 2023](#)), a set of metrics is introduced to measure worst-case performance. These are particularly suited to assess the *stability* of the algorithms. They first define the average minimum accuracy reached by previous tasks when learning task T_k , min-ACC_{T_k} (see Eq. 2). It gives a good idea of the worst case performance of the agent on a given task. Then, the worst-case accuracy, WC-ACC_t (see Eq. 3), combines information from the minimum accuracy on previous tasks and the accuracy on the current task. WC-ACC_t summarizes the trade-off between stability (accuracy on previous task data) and plasticity (accuracy on current task data). This metric is upper-bounded by the average accuracy. Here, t is the current iteration, T_k the current task (at iteration t) and $t_{|T_i|}$ is the iteration at the end of learning T_i . Since WC-ACC_t is upper bounded by the average accuracy, we also report a new metric which is the relative gap between the latter and average accuracy Acc_t as defined in Equation 4, we name it *Relative Accuracy Gap* (RAG) since this measures the relative gap between worst-case accuracy and average accuracy. This metric can then be fairly compared across various methods that have different average accuracy.

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

$$\text{min-ACC}_{T_k} = \frac{1}{k-1} \sum_i^{k-1} \min_n \mathbf{A}(E_i, f_n), \forall n : t_{|T_{i-1}|} < n \leq t \quad (2)$$

$$\text{WC-ACC}_t = \frac{1}{k} \mathbf{A}(E_k, f_t) + \left(1 - \frac{1}{k}\right) \text{min-ACC}_{T_k} \quad (3)$$

$$\text{Acc}_t = \frac{1}{k} \sum_i^k \mathbf{A}(E_i, f_t) \quad \text{and} \quad \text{RAG}_t = \frac{\text{Acc}_t - \text{WC-ACC}_t}{\text{Acc}_t} \quad (4)$$

Method	Accuracy	Memory (Mb)
ER	26.2	4 + 211
ER Naive Ensemble	32.1	1840 + 211
ER + EMA	36.3	8 + 211

Table 1: Comparison on Split-MiniImagenet (20 tasks) of the naive ensembling of checkpoints taken along the training trajectory of a replay method every 10 iterations, against the use of EMA ensemble. For clarity, we divide the memory footprint into the one for the models and the one for the replay buffer (model + buffer).

4 METHOD: EXPONENTIAL MOVING AVERAGE ENSEMBLE (EMA)

In the introduction (see Figure 1) we have shown that ensembles can greatly improve performance, however, they come at a significant increase in memory usage which is in direct conflict with the memory requirements typically imposed on continual learners. Indeed, both continual learning and online learning impose memory constraints since they do not allow retaining more than a fixed amount of data coming from the stream of data. Therefore, in this section, we look into methods to reduce the memory usage, while maintaining the advantages of model ensembling.

Several works have focused on reducing the memory constraint of ensembles [Wen et al. \(2020\)](#); [Wortsman et al. \(2021\)](#); [Tarvainen & Valpola \(2017\)](#), some did so notably by averaging the models in weight space instead of aggregating the predictions in the functional space. While it is not clear under which condition such manipulation of the weights can form a model that performs similarly to the ensemble of the summed members, several works have shown practical working cases. It is possible to perform such a summation ([Wortsman et al., 2021](#)) whenever two models are connected by a linear path of low loss ([Frankle et al., 2020](#)). [Tarvainen & Valpola \(2017\)](#) propose instead to do a weighted sum of an infinite amount of checkpoints by giving older checkpoints less important weight, decreasing exponentially with the distance.

For a potential use of these ensembles in continual learning we are interested in having an ensemble that covers many tasks (see Section 1), and one that is cheap to store and compute. The solution adopted in ([Tarvainen & Valpola, 2017](#)), for semi-supervised learning, fits well to this task since it requires storing only one additional model and is able to ensemble models from all of the previous iterations. Consider a function $f(x) = f(x, \theta^t)$ with learnable weights θ^t (at training iteration t), the exponential moving average (EMA) of it’s weights is defined as:

$$\theta_{ema}^t = \lambda \theta_{ema}^{t-1} + (1 - \lambda) \theta^t, \tag{5}$$

where λ is a user-defined hyperparameter comprised between 0 and 1, which sets the importance of the current model in the running average compared to the one of the previous models, θ_{ema}^{t-1} is the value of the moving average at the previous iteration, and θ^t is the weights of the training model at iteration t (θ^t can be computed with existing online methods, such as ER ([Chaudhry et al., 2019](#)) or MIR ([Aljundi et al., 2019a](#))). This implicit definition can also be written as an explicit sum over all the previous model weights:

$$\theta_{ema}^t = \sum_{i=1}^t (1 - \lambda) \lambda^{t-i} \theta^i + \lambda^t \theta_{ema}^0. \tag{6}$$

This means that the ensemble formed by the sum virtually covers all the previously encountered tasks. However, exponentially less weights will be reserved to older tasks, potentially reducing the effective diversity of the ensemble. Nevertheless, the motivational experiment we conducted on Split-Cifar100 show that after some number of tasks covered by the ensemble, the accuracy gained by covering more and more tasks is less important (sublinear growth). So covering a small number of tasks with the ensemble can be sufficient to get satisfying performance gains. This motivated us to analyze the exponential moving average model in the setting of online continual learning.

In Table 1 we compare Naive Ensembling, discussed in Section 1, with the EMA model when combined with Experience Replay [Chaudhry et al. \(2019\)](#). As can be seen, the EMA model significantly reduces the memory usage (requiring only one additional model). Remarkably, ER+EMA outperforms the Naive Ensemble. This could be caused by the fact that ER+EMA combines many more models, and because of the non-linear weight assignment to the various models (see Appendix Figure 8), where EMA assigns more weight to the last (and better) models in the training trajectory.

While the EMA approximate ensemble is commonly used in the literature and has been proven to give good performance ([Tarvainen & Valpola, 2017](#); [Grill et al., 2020](#); [Caron et al., 2021](#)). Other weight summing schemes could be

	w_i	ER	ER-ACE	RAR
-	-	9.9 ± 0.6	16.5 ± 0.7	27.6 ± 1.3
EMA $\lambda = 0.99$	$\frac{w_{i-1}}{\lambda}$	14.0 ± 0.5	19.0 ± 0.3	35.4 ± 1.2
EMA $\lambda = 0.995$	$\frac{w_{i-1}}{\lambda}$	18.0	20.1	36.8
EMA $\lambda = 0.999$	$\frac{w_{i-1}}{\lambda}$	18.3	18.8	31.2
Uniform	1	13.1	12.7	17.3
Linear	i	16.4	16.3	25.4
Logarithmic	$w_{i-1} + \log(i)$	16.6	16.6	26.2
Quadratic	$w_{i-1} + i^2$	18.2	18.8	31.5

Table 2: Comparison of different approximate ensembling methods inspired from the EMA approximate ensembling method. Each ensembling technique was tried on Split-Cifar100 in combination with ER (Chaudhry et al., 2019) ER-ACE (Caccia et al., 2022), and RAR (Zhang et al., 2022). The second column indicates how the weights for the current model are computed at each step. The first row indicates the results when no ensembling technique is used.

tried to compute an approximate ensemble. Such schemes are not necessarily expected to work in classical offline learning since summing models that are far away from each other in the weight space is not guaranteed to work. However, since online learning only performs a few training iterations on the task at hand (compared to offline learning), we can expect the models to be closer from each other and thus allow other summing techniques to work. In order to explore the possibilities of such summing techniques in online continual learning, we compare several weighting techniques. Since we are working in the online learning setting and under the constraint of storing only one additional model, we choose to compute the approximate ensemble following Equation 7,

$$\theta_{ensemble}^t = \frac{1}{\sum_i^t w_i} \sum_i^t w_i \theta^i, \quad (7)$$

and use this ensemble instead of the EMA ensemble. This formulation allows for more freedom in the choice of the weighting scheme, but the formulation used by EMA can also be expressed under this form, we identify the equivalent weight w_i for the EMA ensemble using Equation 6 as being $w_i = \frac{w_{i-1}}{\lambda}$. We update at every iteration the weighted sum of the models and normalize it by the sum of the weights to avoid exploding model weights. In Appendix (Figure 11) we provide a comparison of the weight distribution for the different strategies we tried.

In Table 2, we present the results of combining an ensemble computed with each weighting scheme with three methods from the literature on the Split-Cifar100 dataset (See Section 5 for explanations about the experimental settings). We see that the EMA model gets the best performance overall, especially with $\lambda = 0.995$ ¹. However, we get surprisingly high results with linear, logarithmic, and quadratic weighting, especially in combination with ER, but linear and logarithmic weighting fail to give an advantage when combined with the more advanced methods ER-ACE and RAR. Quadratic weighting obtains the closest results to the EMA methods, we provide a more detailed comparison of Quadratic against EMA method in the Appendix (Figure. 12). Uniform weighting is the scheme that performs worst across the board. which can be understood since it gives equal weight to the first models than to the last models whereas the last models have been trained for a longer time and thus are expected to have better performance.

5 EXPERIMENTS

Datasets. We perform experiments on 3 datasets. **Cifar-10** is a 10-class dataset that contains 60000 images of size 32 by 32 and 3 color channels (Krizhevsky, 2009). **Cifar-100** has the same image dimensions and number of images but with 100 classes. **Mini-Imagenet** (Vinyals et al., 2016) is a 100-class version of **ImageNet** (Russakovsky et al., 2015), that contains 60000 images which are rescaled to 84 by 84. We split these datasets into 5, 20 and 20 tasks respectively, each containing a mutually exclusive set of classes.

Scenario. We present results in the online class-incremental setting. When continual evaluation is performed, we evaluate after each unique mini-batch. All the compared methods are using a replay buffer, with a fixed memory size of 1000 exemplars for Cifar-10, 2000 exemplars for Cifar-100 and 10000 for Mini-Imagenet (as in (Caccia et al., 2022)). Each training mini-batch is formed out of half of exemplars from previous tasks and half from the current task.

Methods. We compare the performance of five replay methods. ER-ACE (Caccia et al., 2022), MIR (Aljundi et al., 2019a), RAR (Zhang et al., 2022), DER (Buzzege et al., 2020) are described in Section 2.1, while ER (Chaudhry

¹These parameters are not the one we use in the main results of Section 6 (we use $\lambda = 0.99$). For a discussion on hyperparameter choices, refer to Section A.1

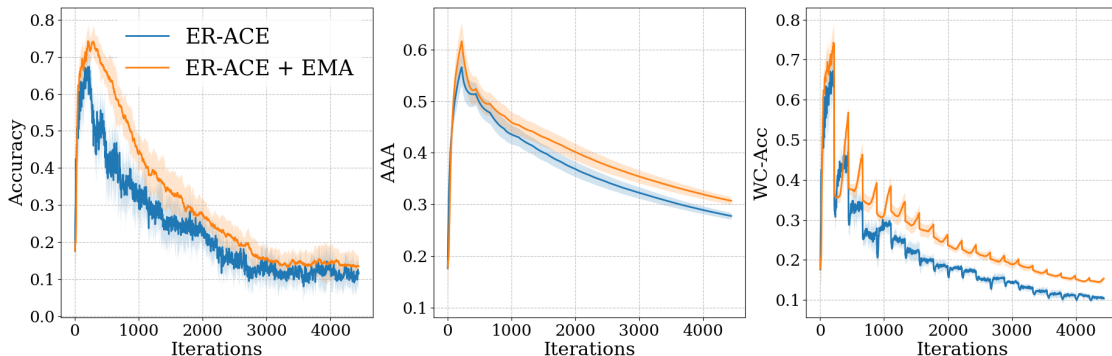


Figure 2: Validation accuracy on task 1 data (Left), Average Anytime Accuracy AAA_t (Middle) and $WC-ACC_t$ (Right) for ER-ACE and its EMA augmented version on Split-Cifar100, using 2000 memory. Mean and standard deviation are computed over 6 runs.

et al., 2019) is the vanilla replay baseline. We display their performance along with the one of their EMA augmented version on the studied datasets. Additionally, we report the results of an *i.i.d* reference method that is allowed the same memory and computational budget as the compared methods, but for which the data arrives in an independent and identically distributed manner (in contrast to the continual learning manner where data arrives split by split). Since some of the methods included in the comparison make use of input transformations (RAR), we also include the results of the *i.i.d*_{w/tr} reference method, which uses the same input transformations.

Training and Implementation Details. For all datasets we use a slim version of Resnet-18 as done in (Lopez-Paz & Ranzato, 2017) and perform 3 passes per mini-batch using Stochastic Gradient Descent with a learning rate of 0.1, and batch size of 32. We run each experiment for six seeds and report the mean and standard deviation. For DER, we stick to the parameters used in the original paper for CIFAR10 ($\alpha = 0.1$ and $\beta = 0.5$). For the EMA ensemble, we chose a momentum parameter of $\lambda = 0.99$. More details on the choice of this parameter can be found in the Appendix Section A.1. We make use of the Avalanche framework (Lomonaco et al., 2021) for all experiments. We make the code available at: https://github.com/AlbinSou/online_ema.

Metrics. For every method, we report the final average accuracy but also the continual evaluation metrics described in Section 3.2, that we computed on a held out validation set after training on each new mini-batch. The validation dataset contains 5% of the total training data. We report both $AAA_{T_{final}}$ and $WC-ACC_{T_{final}}$ in the tables, where T_{final} is the last training iteration. We also report $WC-ACC_t$ at every iteration in the figures. The final value of the *RAG* metric that we define in Equation 4 is reported in percent. Note that we do not make use of this validation data to tune hyperparameters but just to compute the continual evaluation metrics.

6 RESULTS

On **Split-Cifar10** the EMA ensemble offers consistent improvements across all methods (see Table 3), especially for RAR, where it offers a 4.3% improvement in final average accuracy. We hypothesize that the gains of EMA are smaller than on Split-Cifar100 and Split-MiniImagenet because in that case the EMA ensemble weights cover less tasks than in the case of the other two datasets (5 tasks but the same number of training iterations than Split-Cifar100). Nevertheless, the stability metrics are greatly improved also for this dataset.

On **Split-Cifar100** (See Table 3) the EMA ensemble offers considerable improvements for ER, RAR, DER and MIR (from 4.0-7.8%). The improvements are less consequent for ER-ACE, but still important (2.5%). We hypothesize that the smaller gain is due to the smaller task-recency bias of ER-ACE. This is illustrated in Figure 2 and Figure 3. In these two figures on the left, we see that the performance of RAR on task 1 data drops instantly after learning task 1, which means it has been traded for accuracy on task 2, it suffers from the task-recency bias. Whereas for ER-ACE, the performance on task 1 does not drop as much after learning task 1, thus the gap with the EMA augmented version is not as important. This is also reflected in Table 3. In these figures as well, we can see how the use of the EMA model improves the stability, both by looking at the WC-Acc but also at the reduced accuracy variations on a single task.

On **Split-MiniImagenet**, we see the largest performance improvements. This time, RAR sees similar gains than ER (9.3% and 10%), while ER-ACE gains are also more important than in the case of Split-Cifar100. This is coherent with the observations that we had in the motivational experiments (Figure 1) that the gains from ensembling are slightly

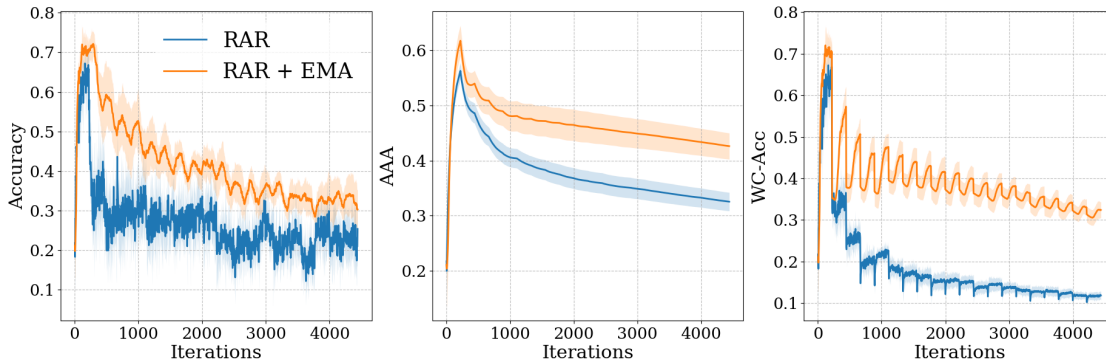


Figure 3: Split-Cifar100, validation accuracy on task 1 data (Left), Average Anytime Accuracy AAA_t (Middle) and WC-ACC (Right), for RAR and its EMA augmented version, using 2000 memory. Mean and standard deviation are computed over 6 runs.

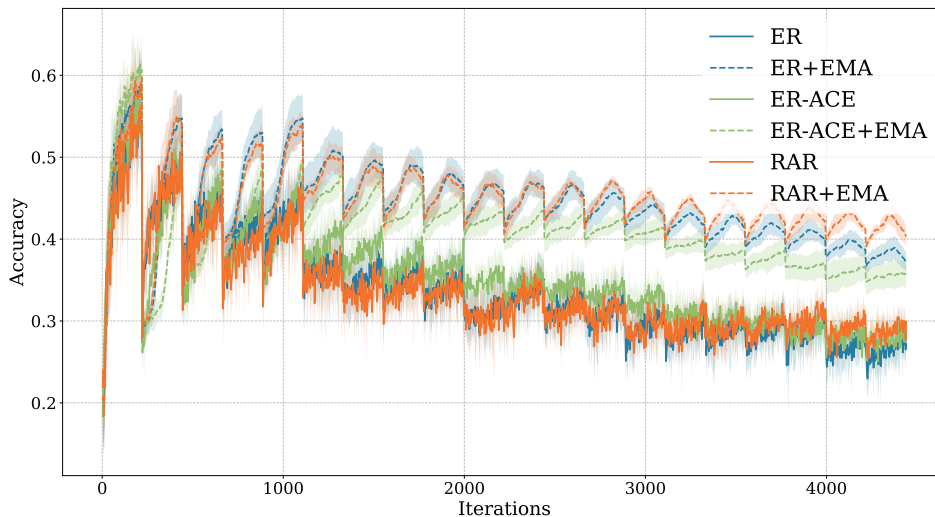


Figure 4: Comparison of the average accuracy on the validation data of three methods trained on Split-MiniImagenet (20 tasks), and their EMA augmented version. EMA performance is indicated with the dotted lines. We report the mean and standard deviation over 6 runs.

more important in the case of Split-MiniImagenet. This might also be due to the memory size used that is different from the one used for Split-Cifar100. In Figure 4, we show a comparison of three methods and their EMA augmented version. We notice that in that case, the use of ER-ACE hurts the performance of the EMA model, which performs worse than just using ER and EMA. Also, for ER and RAR, we notice various bumps in the validation accuracy² of the EMA model that are not present in the current model accuracy. We believe these bumps occur when the previous task bias is compensated by bias towards the current task. The location and width of these bumps depend on the λ parameter chosen for the exponential moving average (for more analysis see Appendix 7).

Effect on the task-recency bias: In Figure 6, we display the task confusion matrices of RAR, for the final training model, the final EMA model, and the EMA model taken at the tip of the bump (selected using a hold-out validation memory mechanism). The matrix shows the number of test images from a particular task (y-axis) that are classified as being from another task (x-axis). For the last training model, a lot of samples are predicted to be in the last task as indicated by the last column, which shows an important task-recency bias. For the final EMA model, the task-recency

²More generally we observe these bumps on all the studied datasets and report them for Split-MiniImagenet and Split-Cifar100 in the Appendix Figure 4 and Figure 7

Method	Split-Cifar10				Split-Cifar100			
	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow	RAG ^{val} \downarrow	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow	RAG ^{val} \downarrow
<i>i.i.d</i>	65.2 \pm 1.8				23.4 \pm 0.7			
<i>i.i.d</i> ^{+EMA}	67.3 \pm 1.3				25.8 \pm 0.6			
	+2.1				+2.4			
<i>i.i.d</i> _{w/tr}	71.2 \pm 1.4				32.3 \pm 1.0			
<i>i.i.d</i> _{w/tr} ^{+EMA}	75.5 \pm 1.3				37.2 \pm 0.8			
	+4.3				+4.9			
<i>ER</i>	37.5 \pm 1.6	57.3 \pm 0.6	26.0 \pm 1.0	31.0 \pm 5.6	9.9 \pm 0.6	22.5 \pm 0.8	5.8 \pm 0.4	43.0 \pm 3.4
<i>ER</i> ^{+EMA}	38.8 \pm 1.3	59.9 \pm 0.7	35.1 \pm 1.1	9.9 \pm 1.5	14.0 \pm 0.5	29.2 \pm 0.9	13.1 \pm 0.7	5.6 \pm 1.6
	+1.2	+2.6	+9.1	-21.1	+4.1	+6.7	+7.3	-37.4
<i>MIR</i>	40.2 \pm 2.8	54.0 \pm 0.4	17.0 \pm 0.6	57.1 \pm 4.4	10.6 \pm 0.7	22.8 \pm 0.7	6.3 \pm 0.4	39.6 \pm 3.0
<i>MIR</i> ^{+EMA}	42.7 \pm 2.1	55.9 \pm 4.2	34.5 \pm 1.5	19.7 \pm 4.8	14.9 \pm 0.4	28.8 \pm 0.9	14.3 \pm 0.6	5.4 \pm 1.4
	+2.5	+1.9	+7.5	-37.4	+4.3	+6.0	+8.0	-34.2
<i>ER-ACE</i>	50.2 \pm 1.2	62.7 \pm 1.4	34.6 \pm 2.0	29.9 \pm 4.8	16.5 \pm 0.7	27.7 \pm 0.6	10.3 \pm 0.4	38.7 \pm 3.3
<i>ER-ACE</i> ^{+EMA}	51.5 \pm 1.6	64.6 \pm 1.4	48.7 \pm 1.9	4.6 \pm 0.8	19.0 \pm 0.3	30.7 \pm 0.6	15.3 \pm 0.6	20.9 \pm 3.5
	+1.3	+1.9	+14.1	-25.2	+2.5	+3.0	+5.0	-17.8
<i>DER</i> ₊₊	51.1 \pm 3.0	56.4 \pm 1.6	17.9 \pm 0.6	64.5 \pm 2.9	18.2 \pm 0.7	24.7 \pm 1.0	4.5 \pm 0.6	75.8 \pm 3.2
<i>DER</i> ₊₊ ^{+EMA}	53.1 \pm 3.7	59.5 \pm 1.7	36.4 \pm 1.1	30.6 \pm 3.5	23.2 \pm 1.2	35.0 \pm 1.3	19.7 \pm 0.9	16.7 \pm 2.3
	+2.0	+3.2	+18.5	-33.9	+5	+10.3	+15.2	-59.1
<i>RAR</i>	63.0 \pm 1.4	67.3 \pm 0.9	28.2 \pm 2.9	55.1 \pm 4.1	27.6 \pm 1.3	32.5 \pm 1.5	11.8 \pm 0.9	57.0 \pm 1.8
<i>RAR</i> ^{+EMA}	67.3 \pm 0.8	72.4 \pm 0.8	60.5 \pm 1.1	10.1 \pm 1.5	35.4 \pm 1.2	42.6 \pm 2.1	32.4 \pm 1.7	7.9 \pm 0.9
	+4.3	+4.9	+32.3	-40.0	+7.8	+10.1	+20.6	-49.9

Table 3: Comparison of the final average accuracy on the test set and continual evaluation metrics on the validation set for various methods with their EMA model augmented version, on Split-Cifar10 (Left) and Split-Cifar100 (Right).

Method	Split-MiniImagenet			
	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow	RAG ^{val} \downarrow
<i>i.i.d</i>	29.9 \pm 2.2			
<i>i.i.d</i> ^{+EMA}	34.9 \pm 1.6			
	+5.0			
<i>i.i.d</i> _{w/tr}	32.3 \pm 1.7			
<i>i.i.d</i> _{w/tr} ^{+EMA}	39.9 \pm 1.1			
	+7.6			
<i>ER</i>	26.2 \pm 0.2	33.9 \pm 0.7	11.0 \pm 0.9	59.6 \pm 2.2
<i>ER</i> ^{+EMA}	36.3 \pm 1.1	44.3 \pm 0.9	34.2 \pm 0.6	7.9 \pm 1.4
	+10.1	+10.4	+13.2	-51
<i>MIR</i>	27.3 \pm 1.7	33.9 \pm 0.4	9.6 \pm 0.9	66.3 \pm 2.3
<i>MIR</i> ^{+EMA}	36.1 \pm 1.2	43.5 \pm 0.8	34.0 \pm 1.4	8.6 \pm 2.0
	+8.8	+9.6	+24.4	-57.7
<i>ER-ACE</i>	27.4 \pm 1.7	35.3 \pm 0.5	12.9 \pm 0.9	54.2 \pm 2.6
<i>ER-ACE</i> ^{+EMA}	34.5 \pm 0.8	41.0 \pm 0.8	19.8 \pm 1.1	44.6 \pm 1.8
	+7.0	+5.7	+6.9	-9.6
<i>DER</i> ₊₊	18.4 \pm 1.8	20.3 \pm 2.6	4.0 \pm 0.4	78.5 \pm 1.5
<i>DER</i> ₊₊ ^{+EMA}	23.3 \pm 2.3	29.4 \pm 2.9	17.8 \pm 3.3	25.0 \pm 6.9
	+4.9	+9.1	+13.8	-53.5
<i>RAR</i>	29.1 \pm 0.8	33.8 \pm 0.8	11.4 \pm 0.5	61.6 \pm 2.1
<i>RAR</i> ^{+EMA}	38.4 \pm 0.8	44.9 \pm 0.4	35.6 \pm 0.4	11.6 \pm 1.7
	+9.3	+11.1	+24.2	-60.0

Table 4: Comparison of the final average accuracy on the test set and continual evaluation metrics on the validation set for various methods along with their EMA model augmented version, on Split-MiniImagenet.

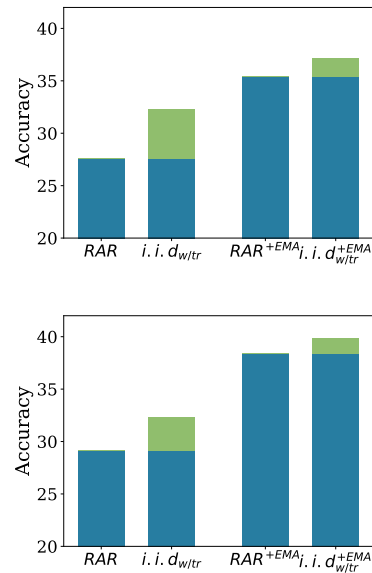


Figure 5: Comparison of previous state-of-the-art method in online continual learning *RAR* against the reference method *i.i.d*_{w/tr} on Split-Cifar100 (Top) using 2000 memory and Split-MiniImagenet using 10000 memory (Bottom). The performance gap is indicated in green, and is greatly reduced by the use of EMA.

bias is also present though slightly diminished, but for the best selected EMA model, it is almost absent, confirming our hypothesis about the origin of the bump³.

³Note that all our results are based on the finishing point of training, which does typically not coincide with the bump.

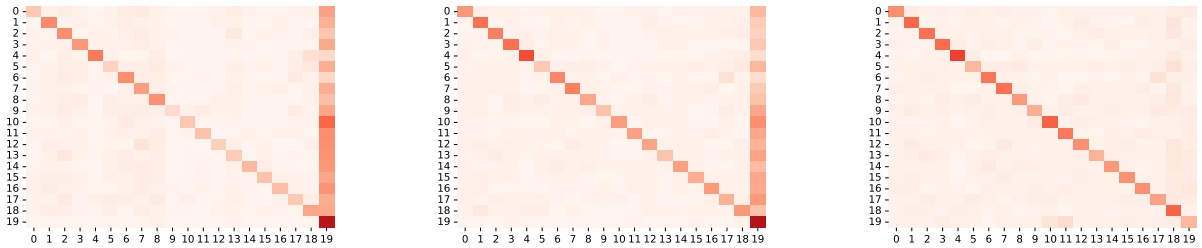


Figure 6: Task Confusion matrices computed on the test set after training of the last task for RAR on Split-MiniImagenet (20 tasks), final training model (Left), RAR+EMA model (Middle), and RAR+EMA model taken at the tip of the bumps observed in Figure 4 (Right). Note the drop in task-recency bias from RAR (a) to RAR+EMA (b) as a consequence of ensembling.

Comparison with the gains in online i.i.d setting: When applied on the reference methods $i.i.d$ and $i.i.d_{w/tr}$, the use of EMA model at evaluation also improves the results by a good margin, showing that the gains obtained in continual learning are not only due to continual-learning based improvements (like reducing the task-recency bias), but also on more general online-learning improvements. However, we observe in general higher gains in continual learning than in the i.i.d setting, except on Split-Cifar10, where they are equivalent, showing that the adaptation of the momentum parameter to the distribution drift speed is essential to get continual learning related gains. To illustrate the higher gains in continual learning, we highlight the gap in final accuracy between previous state-of-the-art method (RAR), and the $i.i.d_{w/tr}$ baseline, along with the EMA augmented versions (See Figure 5). We see that for two of the studied datasets, the gap between RAR and $i.i.d_{w/tr}$ is reduced by the use of the EMA model. For Split-Cifar100, an initial 4.7% gap is reduced to a 1.8% gap, while for Split-MiniImagenet, an initial 3.2% gap is reduced to a 1.5% gap.

Effect on the stability metrics: Finally, for all methods and on all datasets, the AAA and the $WC-Acc$ are greatly improved by the use of EMA, which shows that aside from raising the accuracy, the EMA model offers an important stability boost. We illustrate this effect by showing one single task accuracy curve along with the $WC-Acc$ curve during training in Figure 2 and Figure 3. We see that in both cases both the fluctuations due to small-batch training and the bigger fluctuations due to task shift are reduced by the use of the EMA ensemble. We also present a detailed analysis of stability at the level of a single task shift in Appendix (Figure 10). In general, the biggest increase in $WC-Acc$ also correspond to the biggest decrease in *Relative Accuracy Gap* (RAG), and is significant, confirming that the increase in $WC-Acc$ is not due to an increase in average accuracy, but due to a better stability.

7 CONCLUSION

We investigate the effect of employing temporal ensembling methods, such as EMA, in online continual learning. This direction is particularly interesting because of the distinctive nature of this combination. In online continual learning, temporal ensembles offer the potential to combine models from various training tasks, leading to novel dynamics that cannot be achieved in classical offline learning, where every model ensemble is trained on the same distribution. In the experiments, we show that temporal ensembles can greatly improve continual learning performance and stability. To circumvent the increased memory requirements for the usage of ensembles, we propose to use a memory efficient ensembling solution for online continual learning. We report results using this method in combination with other state-of-the-art methods and conclude that this method consistently increases the final performance and overall stability of several replay methods, closing in on the performance that can be reached in the $i.i.d$ setting. Most surprisingly, we do this without affecting the training process but just by ensembling models from the training trajectory.

We hope that this work inspires the design of more robust ensembling methods for continual learning. In particular, it would be important to find a similarly efficient method that allows to decorrelate the number of training iteration from the number of tasks, since it could then be applied for arbitrary number of iterations per task while similarly covering as many previous tasks as possible. Another future direction could focus on impacting the training using such kind of ensembles, for instance by combining it with distillation ⁴.

Acknowledgements: We acknowledge the support of the Grant PID2019-104174GB-I00 funded by MCIN/AEI/ 10.13039/501100011033 and Grant PID2021-128178OB-I00 funded by MCIN/AEI/ 10.13039/501100011033 and by ERDF A way of making Europe, Ramón y Cajal fellowship Grant RYC2019-027020-I funded by MCIN/AEI/

⁴We provide comments about the application of distillation in combination with EMA model in the Appendix (Section A.4).

10.13039/501100011033 and by ERDF A way of making Europe, and the CERCA Programme of Generalitat de Catalunya. Antonio Carta was partially supported by the H2020 project TAILOR (952215) Connectivity Fund.

REFERENCES

- 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 (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019a. URL <https://proceedings.neurips.cc/paper/2019/file/15825aee15eb335cc13f9b559f166ee8-Paper.pdf>.
- Rahaf Aljundi, Klaas Kelchtermans, and Tinne Tuytelaars. Task-free continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 11254–11263, 2019b.
- Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 583–592, 2019.
- Eden Belouadah, Adrian Popescu, and Ioannis Kanellos. A comprehensive study of class incremental learning algorithms for visual tasks. *Neural Networks*, 135:38–54, 2021.
- Leo Breiman. Bagging predictors. *Machine learning*, 24:123–140, 1996.
- Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and SIMONE CALDERARA. Dark experience for general continual learning: a strong, simple baseline. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 15920–15930. Curran Associates, Inc., 2020. URL <https://proceedings.neurips.cc/paper/2020/file/b704ea2c39778f07c617f6b7ce480e9e-Paper.pdf>.
- 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*, 2022. URL <https://openreview.net/forum?id=N8MaByOzUfb>.
- Massimo Caccia, Pau Rodriguez, Oleksiy Ostapenko, Fabrice Normandin, Min Lin, Lucas Page-Caccia, Issam Hadj Laradji, Irina Rish, Alexandre Lacoste, David Vázquez, and Laurent Charlin. Online fast adaptation and knowledge accumulation (osaka): a new approach to continual learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 16532–16545. Curran Associates, Inc., 2020.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9650–9660, 2021.
- Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 233–248, 2018.
- Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K Dokania, Philip HS Torr, and Marc Aurelio Ranzato. Continual learning with tiny episodic memories. *ICML Workshop: Multi-Task and Lifelong Reinforcement Learning*, 2019.
- 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.
- Thomas G Dietterich. Ensemble methods in machine learning. In *Multiple Classifier Systems: First International Workshop, MCS 2000 Cagliari, Italy, June 21–23, 2000 Proceedings 1*, pp. 1–15. Springer, 2000.
- Thang Doan, Seyed Iman Mirzadeh, and Mehrdad Farajtabar. Continual learning beyond a single model, 2022. URL <https://arxiv.org/abs/2202.09826>.
- Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Linear mode connectivity and the lottery ticket hypothesis. In *International Conference on Machine Learning*, pp. 3259–3269. PMLR, 2020.
- Ian Goodfellow, Yoshua Bengio, Aaron Courville, and Yoshua Bengio. *Deep learning*, volume 1. MIT Press, 2016.

- Ian J Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. An empirical investigation of catastrophic forgetting in gradient-based neural networks. In *Proc. Int. Conf. Learn. Repres.*, 2014.
- Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.
- Lars Kai Hansen and Peter Salamon. Neural network ensembles. *IEEE transactions on pattern analysis and machine intelligence*, 12(10):993–1001, 1990.
- Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 831–839, 2019.
- Gao Huang, Yixuan Li, Geoff Pleiss, Zhuang Liu, John E Hopcroft, and Kilian Q Weinberger. Snapshot ensembles: Train 1, get m for free. *arXiv preprint arXiv:1704.00109*, 2017.
- James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- 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. URL https://openreview.net/forum?id=nrGGfMby_qK.
- Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.
- 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. URL <https://openreview.net/forum?id=Zy350cRstc6>.
- Sang-Woo Lee, Jin-Hwa Kim, Jaehyun Jun, Jung-Woo Ha, and Byoung-Tak Zhang. Overcoming catastrophic forgetting by incremental moment matching. *Advances in neural information processing systems*, 30, 2017.
- Vincenzo Lomonaco, Lorenzo Pellegrini, Andrea Cossu, Antonio Carta, Gabriele Graffieti, Tyler L. Hayes, Matthias De Lange, Marc Masana, Jary Pomponi, Gido van de Ven, Martin Mundt, Qi She, Keiland Cooper, Jeremy Forest, Eden Belouadah, Simone Calderara, German I. Parisi, Fabio Cuzzolin, Andreas Tolia, Simone Scardapane, Luca Antiga, Subutai Amhad, Adrian Popescu, Christopher Kanan, Joost van de Weijer, Tinne Tuytelaars, Davide Bacciu, and Davide Maltoni. Avalanche: an end-to-end library for continual learning. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, 2nd Continual Learning in Computer Vision Workshop, 2021.
- David Lopez-Paz and Marc' Aurelio Ranzato. Gradient episodic memory for continual learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- 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*, 2022.
- Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. 1989.
- Michael P Perrone and Leon N Cooper. When networks disagree: Ensemble methods for hybrid neural networks. In *How We Learn; How We Remember: Toward An Understanding Of Brain And Neural Systems: Selected Papers of Leon N Cooper*, pp. 342–358. World Scientific, 1995.
- 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*, pp. 2001–2010, 2017.

- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Laine Samuli and Aila Timo. Temporal ensembling for semi-supervised learning. In *International Conference on Learning Representations (ICLR)*, volume 4, pp. 6, 2017.
- Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017.
- Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. In *NeurIPS Continual Learning Workshop*, 2018.
- Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. *Advances in neural information processing systems*, 29, 2016.
- Yeming Wen, Dustin Tran, and Jimmy Ba. Batchensemble: an alternative approach to efficient ensemble and lifelong learning. *arXiv preprint arXiv:2002.06715*, 2020.
- Mitchell Wortsman, Maxwell C Horton, Carlos Guestrin, Ali Farhadi, and Mohammad Rastegari. Learning neural network subspaces. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 11217–11227. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/wortsman21a.html>.
- Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019.
- Yaqian Zhang, Bernhard Pfahringer, Eibe Frank, Albert Bifet, Nick Jin Sean Lim, and Alvin Jia. A simple but strong baseline for online continual learning: Repeated augmented rehearsal. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neural Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=bhvUOhnsGz>.

A APPENDIX

A.1 DETAILS ABOUT HYPERPARAMETER CHOICE FOR EMA MODEL

Hyperparameter choice. In all of our experiments, we chose a λ of 0.99 for the EMA model, and additionally warm up the EMA model in the first few iterations by setting this parameter to 0.9 in the beginning. This parameter tells us about the horizon of the EMA model. Since we were interested in an ensemble that covers several experiences, we chose an horizon parameter that is high enough so that the EMA model gives non-negligible weights to previous tasks. However, this process has a flaw since it requires to know how many iterations the learner is going to perform on a given task, which is not suppose to happen in continual learning or in online learning. Nevertheless, we believe more practical solutions can be applied to tune this parameter but we did not investigate them. We think the choice of this parameter should depend on the speed at which the input distribution changes, which might be captured by some other mechanism, we leave this direction to future works. In our experiments, we saw that a relatively wide range of parameters between 0.95 and 0.995 were working correctly and always boosting the accuracy in more or less significant ways. We present results for varying λ on Split-Cifar100 in Figure 7. We see in that figure that increasing λ until the accuracy bump discussed in Section 6 disappears is overall beneficial for the final average accuracy, but this might not be possible to do in practice when not knowing the amount of data to be received from each class. Also, higher λ values, while beneficial for the final average accuracy, affect negatively the accuracy for earlier tasks. One work around to that problem would be to tune λ progressively according to the number of already seen classes vs the number of new coming classes.

The accuracy bumps that we observe in various figures when looking at the average accuracy indicate that the gains that we report in the tables could be bigger if we would retain the model at the tip of the bump. However, this requires the use of a hold-out validation memory and an additional copy of the model on disk. In Figure 6, we retain the model using the above technique of validation memory and observe that these bumps correspond to a lower task-recency bias in the saved EMA model.

Trade-off. This choice of lambda also highlights a problem of this method that makes it difficult to export to non-online learning. It is its dependency on the number of training iterations. Indeed, choosing a higher λ increases the

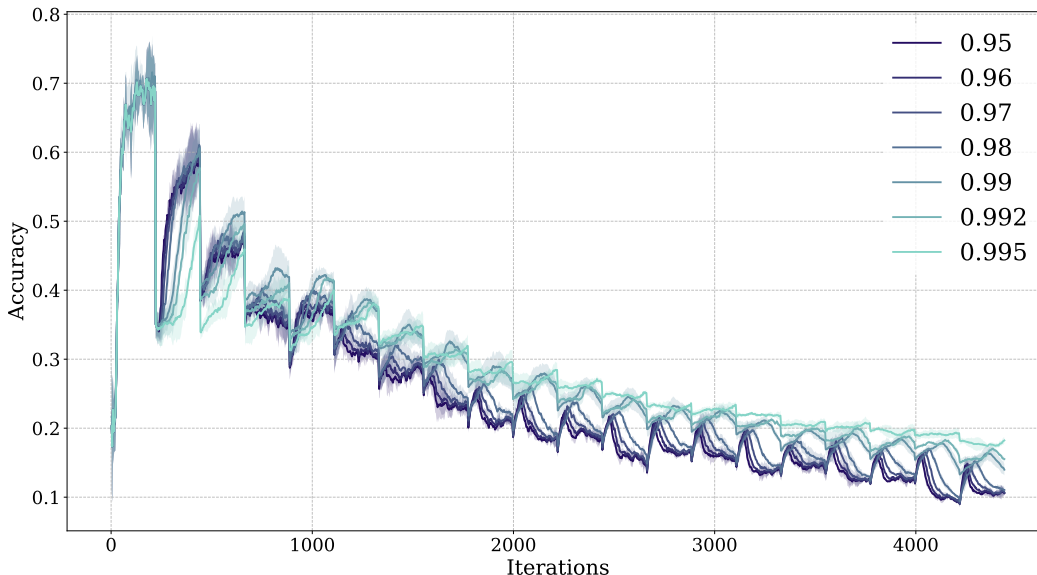


Figure 7: Comparison of the results obtained for various λ values of the EMA model on Split-Cifar100 (20 tasks) for ER. Each curve depicts the average accuracy of the EMA model using a different λ on all tasks seen so far. The value we chose in our experiments (0.99) is not optimal but nor do we have a way of choosing the optimal one. Mean and standard deviation are computed over 3 seeds.

time horizon of the method, but also has negative effects for several reasons. First of all, it is not clear how high of a λ can be chosen before the performance of the ensemble decreases since giving too much weights to single models that are too far in the weight space from the last model might break the assumption that summing models in the weight space leads to a model with accuracy equivalent to the ensemble of these two models [Wortsman et al. \(2021\)](#). This assumption is believed to be true as long as two models are connected by a linear path of low loss, and it is possible that two models that are far away from each other do not respect this constraint. Secondly, it would be a nice addition to be able to sum models obtained after each training batch, or even every training task (instead of each training iteration, since several iterations can be performed on the same training batch in online learning). However for the same reason invoking linear connectivity properties, it is not clear that this would work any better than increasing the momentum value in the computation of the exponential moving average.

Task Covering. For the chosen parameter value, on Split-Cifar100, we can compute the total amount of weight assigned to each task. From Eq. 6, we deduce the weight for a single member in the ensemble θ^i is $(1 - \lambda)\lambda^{t-i}$ where t is the iteration of the current training model. We can sum this over the iterations covering one task. If we place ourselves at the end of one task, with $\lambda = 0.99$ (value used in our experiments) we get 88% of the weights in the last task, 9% of the weight in the previous task, and 1% in the second last task. This indicates that a majority of the weights are assigned to later tasks (in particular from last task to second last task). However, we argue here that as the mean can be highly influenced by outliers, the exponential moving average is nothing else than a weighted mean, and can be influenced by outliers as well. Since the distance in the weight space between models from different tasks is more important than the one between models from the same task, it is possible that earlier models influence more strongly the current moving average than what we could believe by looking at the weights value.

A.2 DETAILS ABOUT THE MOTIVATIONAL EXPERIMENT

Figure 1 shows the comparison between ensembles of models that cover a different number of tasks. To perform this comparison, we first train a model continually using a dedicated continual learning method (ER and ER-ACE in that case), along the way, we save model checkpoints every 10 iterations, leaving about 460 model checkpoints at the end of learning the task (22 models per task in the case of Split-Cifar100 20 tasks). Then, for each number of tasks covered, we select a subset of models that cover no more than that number of tasks, and sample 20 models from that subset that we join into an ensemble for which we record the test performance. We sample 10 of such ensembles for every

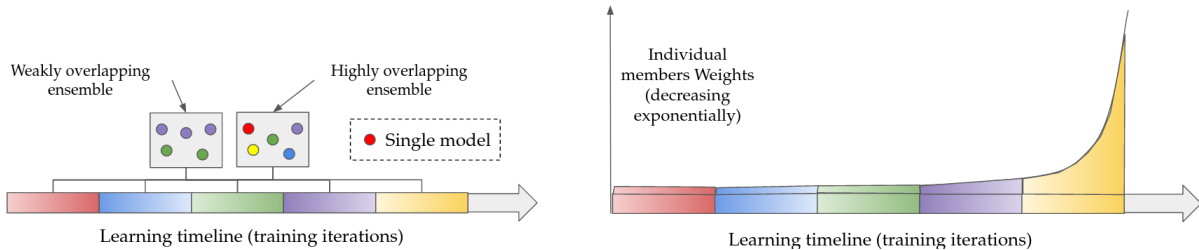


Figure 8: Schema of the motivational experiment ensemble (Left) and of the Exponential Moving Average ensemble (Right). For the EMA ensemble, a continuum of models is inserted in the ensemble which only occupies as much space as one additional model in the memory. Weights of previous model checkpoints decrease exponentially. This procedure permits to cover a spectrum of different tasks in online continual learning.

x-axis value to get mean and confidence interval that we report in the figure. We always add a model in the ensemble which is the last training checkpoint, so that we can always give an accuracy number to later classes when sampling models. The ensembling of models coming from different tasks is done in the following way. We first compute the output probabilities that each model gives to the input. We then compute the per-class mean probability, it is possible that one class is predicted different number of times than another class by the ensemble, so we take that into account and divide the sum of the probabilities for that class by the number of models in the ensemble that can predict that class.

A.3 SUPPLEMENTARY MATERIAL ABOUT THE STABILITY GAP

Here we show several continual evaluation curves (see Figure 9). All of these curves were created by evaluating the model on a hold-out validation set which is 5% of the training set size. The evaluation is performed after training for 3 iterations on each mini-batch before dropping it, as required by the constraints of online learning.

In Figure 10, we compare the accuracy of RAR (Zhang et al., 2022) and RAR+EMA on the data of two subsequent tasks during the task shift. We see that the current training model instantly loses accuracy on previous task at the task shift before regaining part of this accuracy later in training. This results in an important *stability gap* since the accuracy on previous task reaches a low point during training before going back to "normal". This observation is coherent with the one made in Lange et al. (2023). The orange curves indicate the performance of the EMA model on the same tasks. We see that the EMA models takes longer to get good performance on new task but ends up getting better performance than the training model. EMA models also displays improved stability (performance on previous task is smoothly going from initial to final performance).

A.4 ATTEMPTS AT USING EMA MODEL FOR DISTILLATION

Tarvainen & Valpola (2017) initially introduced the EMA model in order to use it as a teacher model in semi-supervised learning. They find that distilling knowledge from the EMA model to the student is beneficial in that setting. We likewise tried to apply distillation by using the EMA model as a teacher, however we found that gains obtained by this method (when present) were not robust enough to be reported. In Table 5, we report the student and teacher (EMA) performance when using, and not using, Mean-Teacher distillation (See Eq. 8). We found that some improvements could be observed in combination with simple ER on Split-Cifar100 both in terms of student and teacher performance, however these improvements did not generalize to Split-MiniImagenet, neither do they combine well with a stronger method like RAR. Notably, in the case of ER_{+MTD} for Split-MiniImagenet and RAR_{+MTD} for both datasets, we observe that the final accuracy of the student is only slightly modified by the distillation process (sometimes increased, sometimes decreased), but the accuracy of the teacher is decreased consequently to its application, this indicates that the distillation process might reduce the diversity of the EMA ensemble by pulling models from the training trajectory closer from one another.

$$L_{CE}(f_{\theta}(x), y) + \alpha L_{CE}(f_{\theta}(x), f_{\theta_{EMA}(x)}) \quad (8)$$

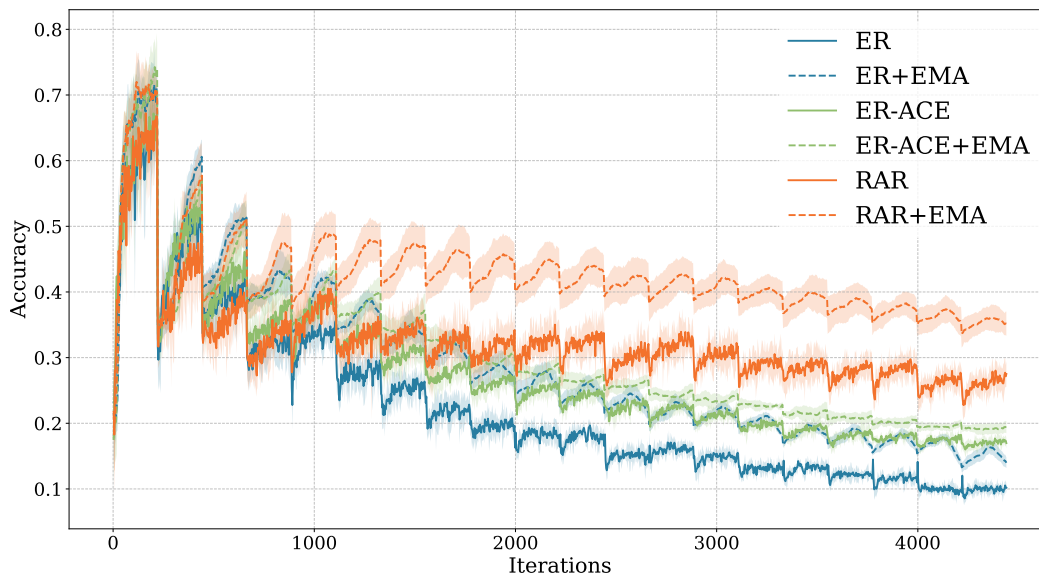


Figure 9: Comparison of the average accuracy on the validation data of three methods trained on Split-Cifar100 (20 tasks), and their EMA augmented version. EMA performance is indicated with the dotted lines. We report the Mean and standard deviation over 6 runs

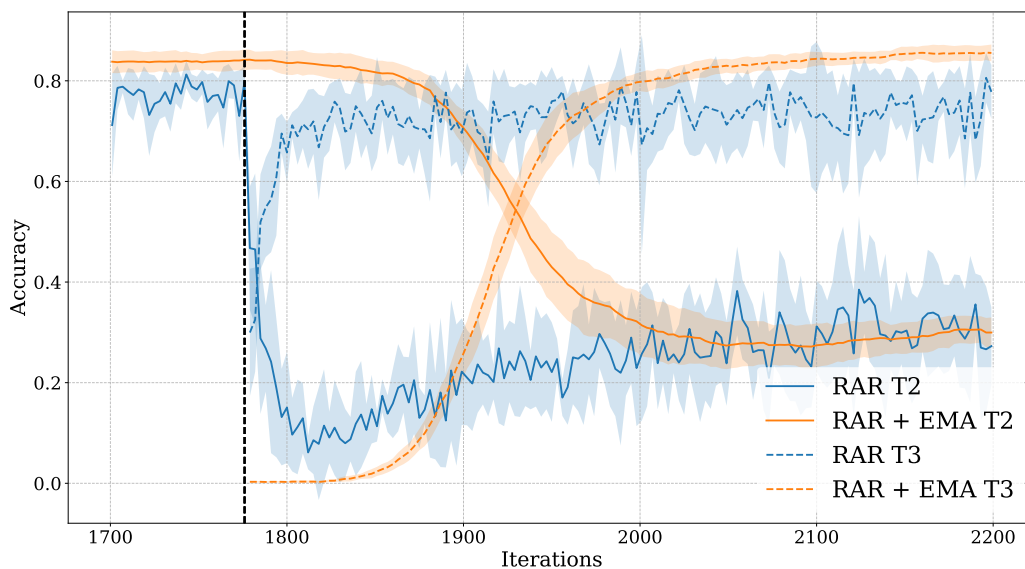


Figure 10: Split-Cifar10, validation accuracy on task 2 data and task 3 data at task shift (indicated with dotted black line). In blue, the accuracy for the training model, in orange, the accuracy for the EMA model. Solid lines indicate accuracy on task 2 data while dotted lines indicate accuracy on task 3 data.

Method	Split Cifar100 Acc	Split MiniImagenet Acc
ER	9.9 ± 0.6	26.2 ± 0.2
ER (EMA)	14.0 ± 0.5	36.3 ± 1.1
ER_{+MTD}	14.7 ± 0.5	27.1 ± 2.1
ER_{+MTD} (EMA)	19.0 ± 0.4	33.5 ± 1.4
RAR	27.6 ± 1.3	29.1 ± 0.8
RAR (EMA)	35.4 ± 1.2	38.4 ± 0.8
RAR_{+MTD}	27.0 ± 1.0	27.1 ± 2.5
RAR_{+MTD} (EMA)	32.7 ± 0.9	33.4 ± 2.5

Table 5: Mean Teacher distillation results for teacher and non teacher on Split-Cifar100 (20 Tasks) (Left) and Split-MiniImagenet (Right). Mean and standard deviation are reported over 6 seeds.

B COMPARISON OF MEMORY USAGE WITH AND WITHOUT THE USE OF EMA MODEL

We include a comparison of memory overhead induced by the use of EMA across all datasets (see Table 6). Using EMA model at evaluation requires to store an additional model. In terms of relative memory usage increase, it is more interesting to use EMA model for larger datasets for which exemplars require more memory to store (MiniImagenet) than for small datasets like Cifar. The methods we compare in this article do not differ significantly from each other in terms of memory usage since they all use replay with the same number of exemplars, this is why we omit them from the comparison.

Method Memory	Split Cifar10		Split Cifar100		Split MiniImagenet	
	Exemplars	Model	Exemplars	Model	Exemplars	Model
w/o EMA	3 Mb	4 Mb	6 Mb	4 Mb	211 Mb	4 Mb
w/ EMA	3 Mb	8 Mb	6 Mb	8 Mb	211 Mb	8 Mb

Table 6: Memory usage for the considered datasets and for the exemplars used in our experiments (1000, 2000 and 10000 for the three datasets respectively).

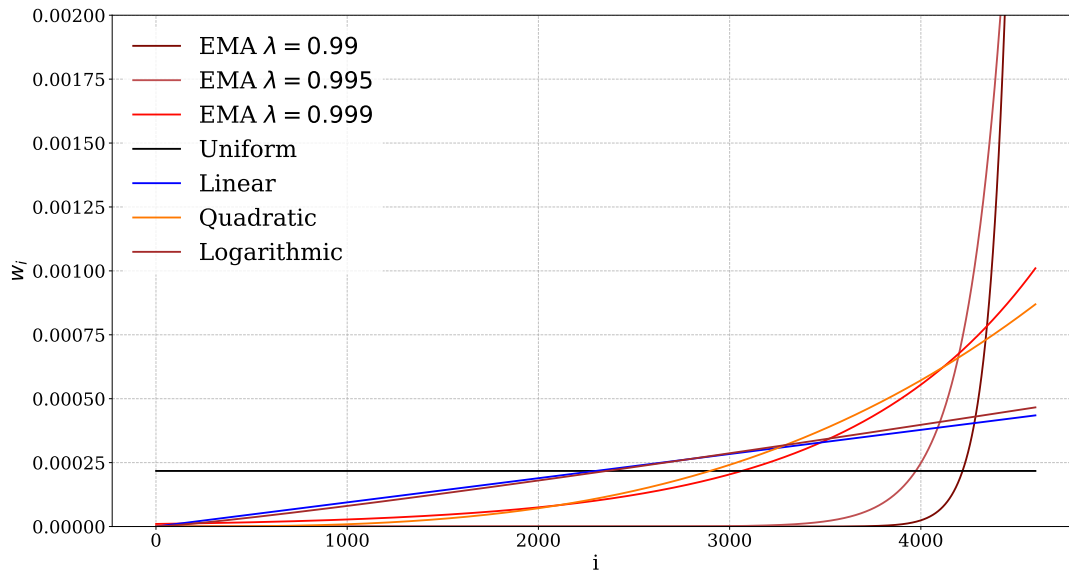


Figure 11: Curve showing the evolution of different weighting schemes (w_i) under the form described in Equation 7. We compare the performance of these weighting schemes in Section 4. We observe here that EMA with a high lambda leads to a weighting that looks similar to quadratic weighting. To draw these curves we place ourselves at the last training iteration ($\theta_{ensemble}^{last}$)

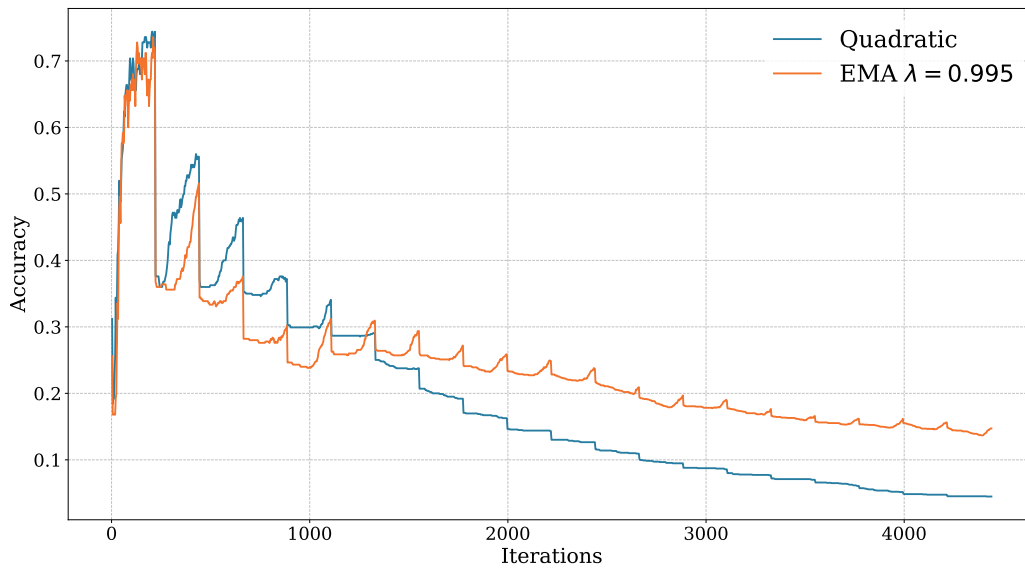


Figure 12: WC-ACC_t for Quadratic vs EMA with $\lambda = 0.995$ in combination with ER on Split-Cifar100 dataset. We see that the quadratic weighting scheme gives more interesting stability (in terms of WC-ACC_t) in the first few tasks, but then fails short compared to the EMA ensemble.