
Online Curvature-Aware Replay: Leveraging 2nd Order Information for Online Continual Learning

Edoardo Urettini^{1 2} Antonio Carta¹

Abstract

Online Continual Learning (OCL) models continuously adapt to nonstationary data streams, usually without task information. These settings are complex and many traditional CL methods fail, while online methods (mainly replay-based) suffer from instabilities after the task shift. To address this issue, we formalize replay-based OCL as a second-order online joint optimization with explicit KL-divergence constraints on replay data. We propose Online Curvature-Aware Replay (OCAR) to solve the problem: a method that leverages second-order information of the loss using a K-FAC approximation of the Fisher Information Matrix (FIM) to precondition the gradient. The FIM acts as a stabilizer to prevent forgetting while also accelerating the optimization in non-interfering directions. We show how to adapt the estimation of the FIM to a continual setting, stabilizing second-order optimization for non-iid data, uncovering the role of the Tikhonov damping in the stability-plasticity tradeoff. Empirical results show that OCAR outperforms state-of-the-art methods in continual metrics, achieving higher average accuracy throughout the training process in three different benchmarks.

1. Introduction

Online Continual Learning (OCL) models are trained continuously on a nonstationary data stream. The goal is to obtain a model that is accurate at any point in time, quickly adapts to new data (i.e. plasticity), and does not forget old information (i.e. stability), without any information about task start, end, or identity. In this setting, many standard Continual

Learning (CL) approaches cannot be applied (Aljundi et al., 2019b; Mai et al., 2022). Many specific OCL methods have been developed, nearly all replay-based (Yoo et al., 2024), but several shortcomings still exist: (Soutif-Cormerais et al., 2023) shows that most OCL methods have high forgetting and fail to beat a simple replay baseline on some metrics; (De Lange et al., 2023; Caccia et al., 2022) discovered the stability gap, a sudden drop in performance at task boundaries; (Dohare et al., 2024) even observed loss of plasticity, limiting the ability to learn in time.

These results suggest fundamental failures of the algorithms. Most methods focus on preventing forgetting at the end of a learning task, an approach that does not ensure stable optimization throughout the entire stream (stability gap); furthermore, stability is often optimized at the expense of plasticity. We argue that a proper CL optimizer should seek to maximize both stability and plasticity, assuming that the model is large enough, at every step in time, interpreting it as a continual filtering process (filter intended as in nonstationary time series literature (Durbin & Koopman, 2012)).

This paper proposes Online Curvature-Aware Replay (OCAR), a novel method designed for replay-based online continual learning, aiming at tackling the challenges in both plasticity and stability inherent in this scenario with a continual optimization approach at every step in time. We formalize OCL as the joint optimization on past and new data, with past data approximated using a limited replay buffer, and adding an explicit constraint on the variation of KL-divergence on previous information. The Fisher Information Matrix (FIM) is used to capture the loss function’s curvature, providing both plasticity and stability constraints in the model distribution space. When the KL divergence is used as a metric, the FIM has the additional value of describing the curvature of the parameter space itself, being the Riemannian metric tensor of that space (Amari, 2016). We can directly adapt our gradient to the geometry of the space in stark contrast with traditional CL methods that use the FIM as a penalization term (Kirkpatrick et al., 2017). Kronecker-factored Approximate Curvature (K-FAC) (Martens & Grosse, 2015) is used to efficiently approximate the FIM, with some critical adjustments to make it work in OCL settings.

¹University of Pisa, Pisa, Italy ²Scuola Normale Superiore, Pisa, Italy. Correspondence to: Edoardo Urettini <edoardo.urettini@sns.it>, Antonio Carta <antonio.carta@unipi.it>.

The main contributions of this paper are: the design of Online Curvature-Aware Replay (OCAR) as a combination of replay, second-order optimization, and information geometry; the analysis of the Tikhonov damping and its ratio with the learning rate in the stability-plasticity tradeoff; an improvement of state-of-the-art performance in standard computer vision benchmarks across all the stability and plasticity metrics.

2. Related Work

Continual Stability and OCL: Most continual learning methods assume that stability must be kept at the expense of plasticity (Lange et al., 2022; Masana et al., 2023), the so-called plasticity-stability tradeoff, and therefore are designed to preserve knowledge about previous tasks to mitigate catastrophic forgetting (French, 1999). However, recent evidence in (Lange et al., 2023; Caccia et al., 2022) suggests that even the methods with high "stability" measured at the end of tasks suffer from high instability and forgetting immediately after the task switch, and they recover the lost performance over time. (Kamath et al., 2024) found evidence of the stability gap even in incremental i.i.d. settings. In OCL settings, (Soutif-Cormerais et al., 2023) showed that some state-of-the-art methods are unable to outperform a simple reservoir sampling baseline on some fundamental stability metrics. Furthermore, CL methods also fail at keeping plasticity, and (Dohare et al., 2024) provides evidence of the loss of plasticity in deep continual networks. Overall, the literature suggests that CL methods fail at both stability and plasticity due to instabilities in the learning dynamics. Recently, some methods such as On-Pro (Wei et al., 2023) and OCM (Guo et al., 2022) proposed novel self-supervised auxiliary losses and prototype-based classifiers, two approaches orthogonal to our optimization-based method. Regarding the latest works on OCL with pre-trained models, (Moon et al., 2023) proposed Si-Blurry, a more realistic scenario with stochastic and blurred boundaries, and introduced MVP, a prompt-based tuning method that keeps the backbone frozen. Advancing on this, (Kang et al., 2025) presented MISA, which introduced a forgetting-aware initial session adaptation and a non-parametric logit masking to alleviate forgetting. In our work instead, we focus on the full training, not only finetuning.

Optimization in Continual Learning: Most CL optimization algorithms are designed to prevent forgetting by removing interfering updates. GEM (Lopez-Paz & Ranzato, 2017; Chaudhry et al., 2019a) models interference using the dot product of the task gradients and constrains the model updates to have positive dot products with the gradients of previous tasks. Subsequent work explored orthogonal projection methods (Saha et al., 2021; Farajtabar et al., 2020) that either extend the idea of interfering gradients or project

in the null space of the latent activations. (Mirzadeh et al., 2020) discusses the relationship between the curvature of the first task and the forgetting, proposing a hyperparameter schedule that implicitly regularizes the curvature. More recently, LPR (Yoo et al., 2024) exploits proximal optimization in the L2 space of latent activations, and it is the only projection-based optimizer compatible with replay. (Hess et al., 2023) proposes a combination of GEM and replay as a potential mitigation for the stability gap.

Natural Gradient and FIM in CL: Natural gradients can be used to train neural networks thanks to efficient approximations of the Fisher Information Matrix (FIM) (Martens & Sutskever, 2012), such as the K-FAC (Martens, 2020) and E-KFAC (George et al., 2018). Interestingly, Benzing (2022) showed that K-FAC seems to work better than the full FIM in some empirical settings, which is connected to a form of gradient descent on the neurons. In continual learning, the FIM is typically used to approximate the posterior of the weights with a Laplace approximation. The result is a quadratic regularizer that is combined with the loss on new data, as introduced by EWC (Kirkpatrick et al., 2017) and its several extensions (Chaudhry et al., 2018; Liu et al., 2018; Huszár, 2017). Magistri et al. (2024) proposed to use the FIM only for the final layer, which can be computed efficiently with a closed-form equation. As an alternative, FROMP (Pan et al., 2020) computes a Gaussian Process posterior, which is also used to estimate the importance of samples in the replay buffer. Daxberger et al. (2023) proposes a method that combines EWC, replay, and knowledge distillation. More relevant to our work, NCL (Kao et al., 2021) proposed a modified natural gradient step with a quadratic posterior as in EWC. Here the Fisher is computed only at the boundaries and the method does not support rehearsal, making it difficult to implement in OCL.

Comparison with our work Most CL methods use the FIM to compute a Laplace approximation of the posterior. This is not possible in OCL because the model is never assumed to be at a local minimum, not knowing task boundaries or length. In general, in the CL literature, the FIM is often restricted to its use in quadratic penalties. While quadratic regularizers are easier to use, we show that the use of the Fisher as a gradient preconditioner is a promising direction, improving the optimization path. Many of the limitations found for EWC and similar methods may be due to some suboptimal choices in the use of the FIM, such as the use of the empirical FIM, popular in the CL literature but with different properties from the FIM (Kunstner et al., 2019). Another difference with the literature is that our optimizer is compatible with replay, unlike most projection-based methods. Furthermore, while most methods penalize plasticity indirectly to prevent forgetting, our approach is designed to improve both on learning speed and forgetting.

3. Online Continual Learning

In continual learning (CL), the model learns incrementally from nonstationary data. In most CL settings, we can identify a sequence of tasks $\mathcal{T}_1, \dots, \mathcal{T}_N$, each one with its own distribution. For example, in class-incremental learning (van de Ven & Tolias, 2019), each task has different classes, and tasks are seen sequentially during training. The goal of the model is to learn all the tasks seen during training. Given the model parameters w_t learned at time t , a loss function \mathcal{L} , and a test set for each task $\mathcal{D}_1, \dots, \mathcal{D}_N$, we can evaluate the model with the average task loss $\mathcal{L}^{avg}(t) = \sum_{i=1}^N \mathcal{L}(w_t, \mathcal{D}_i)$.

Online continual learning (OCL) requires some additional constraints and desiderata (Soutif-Cormerais et al., 2023; Yoo et al., 2024; Mai et al., 2022): (*D1-Online training*) at each step, we do not have access to the whole training dataset for the current task, but only a small minibatch, which can be processed for a limited amount of time; (*D2-Anytime inference*) the model should be ready for inference at any point in time; Consequently, (*D3-Continual stability*) the model must be stable at any point in time, instead of only at the task boundaries; (*D4-fast adaptation*) the model must also be able to learn quickly from new data.

There are similarities with the traditional Online Learning (OL) approach: at each time t , a new observation is made from an unknown task, a new loss is experienced, and the final objective is to minimize the cumulative loss on a stream of arbitrary length. The difference in OCL is that the aim is also to minimize forgetting, hence having a model able to perform well also on past data of the stream. To get this, usually, access to some constrained past information (for example, a small collection of past examples) is allowed.

Following the literature, we focus on replay-based methods, which use a limited buffer of previous data for rehearsal by combining new minibatches with samples from the buffer at each iteration. This work assumes no access to knowledge about task identities, boundaries, or length. Each single observation in the stream can be sampled from different distributions/tasks.

4. Online Curvature-Aware Replay

We now present the building process of our method, starting from the standard ER optimization and expanding it to a second-order method, which is then approximated using FIM and followed by some final adjustments to enhance its CL performance.

4.1. The optimization problem

Online continual learning (OCL) is an online learning problem in a nonstationary setting. Common optimizers in ma-

chine learning implicitly assume stationarity, which justifies the gradient estimate from the minibatches (e.g.: ADAM (Kingma & Ba, 2015), SGD). Instead, in OCL the distribution can change at any point in time (non i.i.d.). At each step, the method can only use the current minibatch and, in replay-based methods, an additional minibatch sampled from a small buffer of old data.

First-order optimization: We define our learning process as a sequence of local optimization problems (as in (Martens, 2020), section 7). Unlike in stationary settings, these problems can greatly change from one step to another, preventing the use of more "global" approaches (e.g., learning rate decay and momentum (LeCun et al., 2015)). For example, the previous gradient directions can be of little use after a task change, making momentum detrimental to fast adaptation to new data (Yuan et al., 2016). On the other hand, learning rate decay would forcibly decrease plasticity in time, reaching a point where no new task can be learned at all. Each single step must be meaningful by itself. What is forgotten or not learned could be lost forever.

The Kullback–Leibler (KL) divergence (Thomas & Joy, 2006; LeCun et al., 2015) is used as our objective, aiming to minimize the "distance" between the predicted and the real distribution. The KL is estimated on the current batch of data $\hat{K}L(y_D || f_w(x_D)) = \frac{1}{N} \sum_i^N KL(\hat{p}(y_{D,i} | f_{w^*}(x_{D,i})) || p(y_{D,i} | f_w(x_{D,i})))$. In our notation D represents the batch of data, i the specific sample, y the target variable dependent on the observed x . f_w is the model parametrized by w that is supposed to represent the reality when $w = w^*$. The empirically observed distribution is $\hat{p}(y_{D,i} | f_{w^*}(x_{D,i}))$.

Experience Replay (ER) (Chaudhry et al., 2019b) keeps a small buffer \mathcal{B} of past data and solves a joint optimization on a batch N_t sampled from the unknown real current data distribution and a batch B_t sampled from \mathcal{B} :

$$\begin{aligned} \min_{\delta_t} \quad & \hat{K}L(y_{N_t} || f_{w_t}(x_{N_t})) + \hat{K}L(y_{B_t} || f_{w_t}(x_{B_t})) \\ \text{subject to} \quad & \frac{1}{2} \|\delta\|_2^2 \leq \epsilon, \end{aligned} \quad (1)$$

where $\delta_t = w_t - w_{t-1}$. Here and throughout the paper we write a sum of the two KL divergences for visual ease. Usually, and also in our experiments, the mean is taken, combining the data in a single batch with no distinctions. The KL divergence can be approximated by first-order Taylor expansion:

$$\hat{K}L(y_{D_t} || f_w(x_{D_t})) \approx \hat{K}L(y_{D_t} || f_{w=w_{t-1}}(x_{D_t})) + \vec{\nabla}_{D_t}^T \delta_t$$

where $\vec{\nabla}_{D_t} = \nabla_w \hat{K}L(y_{D_t} || f_{w=w_{t-1}}(x_{D_t}))$ is the gradient with respect to the weights w of the KL between the targets and the model predictions, with ∇ representing the gradient operator and $\vec{\nabla}$ the resulting vector. When computing this on newly observed data $D_t = N_t$, while for buffer

data $D_t = B_t$. After approximation, the solution for problem 1 is $\delta_t^* = -\frac{1}{\lambda}(\vec{\nabla}_{N_t} + \vec{\nabla}_{B_t})$ where λ is the Lagrange multiplier of the constraint. ER actively optimizes on the buffer and the current data distributions with no distinctions and no forgetting constraint. The only stability requirement is $\frac{1}{\lambda}$ (the learning rate), limiting the movement of the weights in all directions, trading stability for plasticity.

Failure of ER with first-order optimization: Unfortunately, the first-order information can be a poor approximation in CL. Often, information about the curvature is necessary to avoid catastrophic forgetting. For example, at task boundaries, the model is often close to a minimum for the previous task. In that case, buffer gradients $\vec{\nabla}_B \approx 0$, while new gradients $\vec{\nabla}_N$ can potentially be much higher, dominating the update direction. This issue partially explains the stability gap (De Lange et al., 2023). Second-order methods can solve this problem by enlarging the "sight" of our optimizer with the local variation of the variation. The second-order Taylor expansion of the KL divergence is:

$$\hat{K}L(y_{D_t} || f_w(x_{D_t})) \approx \hat{K}L(y_{D_t} || f_{w=w_{t-1}}(x_{D_t})) + \vec{\nabla}_{D_t}^T \delta_t + \frac{1}{2} \delta_t^T \mathbf{H}_{D_t} \delta_t$$

where now $\mathbf{H}_{D_t} = \mathcal{H}_w \hat{K}L(y_{D_t} || f_{w=w_{t-1}}(x_{D_t}))$ is the Hessian of the KL computed on data D_t , with \mathcal{H} representing the operator to compute the second-order derivatives and \mathbf{H} the resulting matrix. again, D_t it can be both N_t or B_t . Using this expansion, the problem 1 is solved by

$$\delta_t^* = -(\mathbf{H}_{N_t} + \mathbf{H}_{B_t} + \tau \mathbf{I})^{-1}(\vec{\nabla}_{N_t} + \vec{\nabla}_{B_t}),$$

where the Lagrange multiplier τ , still related to the L_2 regularization imposed on the update, now acts as a Tikhonov damping term (Martens & Sutskever, 2012). Unlike the learning rate in SGD, τ has a different effect on the eigendirections of the Hessian, which depends on their eigenvalues (see Sec. 5).

Stability constraint: By using the Hessian matrix to precondition the gradient, we expect faster learning on both new and past data, in a Newton-like fashion. If we "stretch" the plasticity requirements to include new information in both new and past data, we can consider this approach to directly answer the *fast-adaptation* desiderata. What it still lacks is an explicit stability constraint necessary for non-i.i.d. settings to obtain the *Continual Stability* requirement. Combining experience replay, second-order methods, an explicit constraint for stability on past data, and L_2 regularization on the update, the problem becomes:

$$\begin{aligned} \min_{\delta} \quad & \hat{K}L(y_{N_t} || f_{w_t}(x_{N_t})) + \hat{K}L(y_{B_t} || f_{w_t}(x_{B_t})) \\ \text{subject to} \quad & \hat{K}L(f_{w_{t-1}}(x_{B_t}) || f_{w_t}(x_{B_t})) \leq \rho \\ & \frac{1}{2} \|\delta\|_2^2 \leq \epsilon. \end{aligned}$$

The term $\hat{K}L(f_{w_{t-1}}(x_{B_t}) || f_{w_t}(x_{B_t}))$ measures the variation of the model predicted distribution on buffer data, and its expansion around the pre-update parameters w_{t-1} will have zero zeroth and first-order terms. The remaining term is controlled by the Hessian of the KL-divergence evaluated at $w = w_{t-1}$, which is exactly the Fisher Information Matrix (Ollivier et al., 2017) (Martens, 2020):

$$\mathbf{F}_{B_t,ij} = \left. \frac{\partial^2 \hat{K}L(f_{w_{t-1}}(x_{B_t}) || f_w(x_{B_t}))}{\partial w_i \partial w_j} \right|_{w=w_{t-1}}.$$

The optimization problem becomes

$$\begin{aligned} \min_{\delta} \quad & \vec{\nabla}_{N_t}^T \delta + \frac{1}{2} \delta^T \mathbf{H}_{N_t} \delta + \vec{\nabla}_{B_t}^T \delta + \frac{1}{2} \delta^T \mathbf{H}_{B_t} \delta \\ \text{subject to} \quad & \frac{1}{2} \delta^T \mathbf{F}_{B_t} \delta \leq \rho \\ & \frac{1}{2} \|\delta\|_2^2 \leq \epsilon, \end{aligned} \tag{2}$$

which is solved by

$$\delta_t^* = -(\mathbf{H}_{N_t} + \mathbf{H}_{B_t} + \lambda \mathbf{F}_{B_t} + \tau \mathbf{I})^{-1}(\vec{\nabla}_{N_t} + \vec{\nabla}_{B_t}),$$

where λ controls the strength of the stability constraint, ensuring that the predictions of the model remain stable, while τ is the Tikhonov damping controlling directions of maximum acceleration. As done in many second-order methods (Martens, 2020; Martens & Sutskever, 2012), a learning rate α can be used to uniformly control the step size in all directions, making the final update be $\alpha \delta_t^*$. As we will show, the relation between α and τ has an important impact on the learning dynamics and the stability-plasticity tradeoff.

4.2. Estimations and approximations

Approximating the Hessians: Computing two Hessians and one Fisher Information Matrix would be impractical and inverting them unfeasible. Luckily, in the second-order optimization literature, it has been shown how for some distributions (including the multivariate normal and the multinomial), the FIM is equivalent to a Generalized Gauss-Newton (GGN) matrix, both approximating the Hessian of the KL-divergence (Martens, 2020). They are equal to the Hessian when the model describes well the data (near optimum) and the work of Martens (Martens, 2020) illustrates some general advantages of these approximations over the Hessian also when far from optimum.

Moreover, the use of the FIM has the additional interpretation of describing the curvature of the parameter space itself. In this space, each point represents a distribution, requiring the use of the KL divergence as "distance". The KL-divergence of infinitesimal displacements corresponds to $\frac{1}{2} F_{ij} dw_i dw_j$ giving rise to the Fisher Information as the

metric tensor used in this manifold (Amari, 2016). More than a simple approximation of the Hessian, the FIM can be used to correctly measure distances in the manifold where we are performing the gradient descent if the steps are small enough. Assuming a small learning rate, our method would use the preconditioner not as in a Newton method (that usually requires a learning rate of 1 to directly go to the approximate optimum) but as a metric adaptation.

We believe these two results justify the use of the Fisher information in nonstationary settings as OCL. Following this, we approximate the two Hessian matrices of our solution with the FIM, greatly simplifying the computations. The new optimal update becomes (including the learning rate):

$$\delta_t^* = -\alpha(\mathbf{F}_{N_t} + (1 + \lambda)\mathbf{F}_{B_t} + \tau\mathbf{I})^{-1}(\vec{\nabla}_{N_t} + \vec{\nabla}_{B_t}).$$

This implies the second-order information is now taken with respect to the model prediction instead of the real targets as also $\mathbf{F}_{N_t} = \mathcal{H}_w KL(f_{w_{t-1}}(x_{N_t}) || f_w(x_{N_t}))$. The gradient of our method corresponds to the usual gradient obtained from the loss function when the loss is a derivation of the KL divergence (negative log-likelihood, cross-entropy, etc...). The weighted sum (or mean) of the two FIMs can now be obtained by computing a single FIM on the batch, giving more weight to the buffer data.

There is a deep connection between our method and Natural Gradient (NG) (Amari, 1998). We have to underline that our building process is much different from the one used for the original NG, thought for stationary settings. Moreover, our preconditioner is not really the FIM of the model, but a modified and regularized version stemming from our OCL second-order optimization with the use of replay. We then cannot assume all benefits of NG would apply to our case.

The Empirical FIM: The Fisher is also defined as the expected value of the squared score:

$$F = \sum_n \mathbb{E}_{y \sim p(y|f_w(x_n))} \left[\nabla_w \log p(y|f_w(x_n)) \cdot \nabla_w \log p(y|f_w(x_n))^T \right].$$

One could argue that a better preconditioner for the gradient would be the Empirical Fisher (EF) matrix, computed using the real target y instead of the one sampled from the predicted distribution, in particular for approximating the Hessian of new data \mathbf{H}_{N_t} . Besides getting us outside the theory about Natural Gradient and Fisher/GGN equivalence, it has been shown it is a questionable choice, even when the model is not a good description of the data (Kunstner et al., 2019). For these reasons, OCL uses the "real" Fisher, unlike other traditional CL approaches (Kirkpatrick et al., 2017).

K-FAC: This approach has theoretical advantages, but it requires the inversion of an extremely large matrix, unfeasible

for large networks. While diagonal approximations are frequently used (Kirkpatrick et al., 2017; Kingma & Ba, 2015), given the particular challenges of OCL a more informative approximation is needed. We rely on the Kronecker-factored Approximate Curvature (K-FAC) in its block-diagonal version to approximate the FIM (Martens & Grosse, 2015):

$$\tilde{F} = \text{diag}(\bar{A}_{0,0} \otimes G_{1,1}, \dots, \bar{A}_{l-1,l-1} \otimes G_{l,l}),$$

where each block corresponds to an approximation of the covariance matrix of the score of a specific layer, obtained by the Kronecker product of $\bar{A}_{i,i} = \mathbb{E}[\bar{a}_i \bar{a}_i^T]$ and $G_{i,i} = \mathbb{E}[g_i g_i^T]$. \bar{a}_i is the vector of the layer activations (with an additional 1 appended for the bias) and g_i the gradient of the prediction with respect to the output of the layer before the activation function. The expected values of both are computed with an Exponential Moving Average (EMA) of past values. The efficiency of this method relies on the property of the Kronecker product $(A \otimes B)^{-1} = A^{-1} \otimes B^{-1}$ allowing us to invert much smaller matrices for each layer, ignoring layer interactions.

4.3. Nuts and Bolts for OCL

To make the method work in practice, we found some adjustments are needed.

Hyperparameter optimization and τ scheduling: Usually, hyperparameter selection is done only on the first K tasks of the stream, but it must generalize to longer streams during training. Our method uses three hyperparameters: the learning rate α , the Tikhonov damping strength τ , and the parameter γ for the EMA used for Kronecker factors estimation. In hyperparameter selection, we found it beneficial to search for a value for the increase of τ instead of α itself. τ is then initialized at the same value of the learning rate and increased by the selected value at each optimization step improving long-term stability.

Estimate of K-FAC factors at boundaries: In class-incremental settings, the shape of the classifier will grow over time as new classes are observed. This means the K-FAC factor $G_{l,l}$ of the last layer will change shape, also breaking the relations with the previously observed gradients. To avoid keeping track of errors, if $G_{l,l}$ changes shape, the EMA of this factor is reset. This is not done for the other factor of the last layer $\bar{A}_{l-1,l-1}$. We assume the model has consistent representations, and that the majority of instability is happening on the classifier.

Scheduling of λ : The stability constraint of problem 2 is estimated on the current batch extracted from the buffer, but it should represent the whole buffer. As new experiences are encountered in time, the information content of the buffer will grow, with less redundancy. At the limit, we can get a buffer where each example represent a different class or domain. Being λ the Lagrange multiplier, it is directly

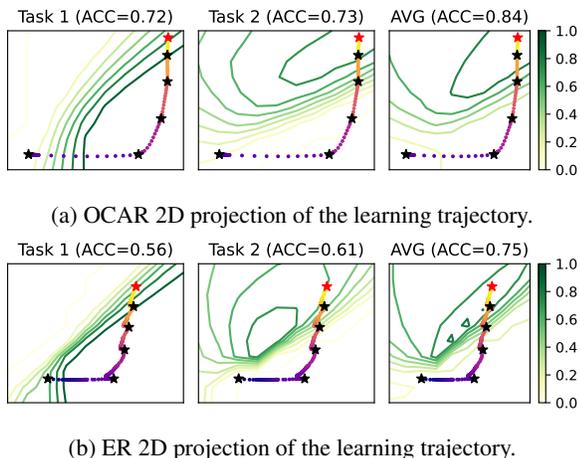


Figure 1. 2D projections of the training trajectories for ER and OCAR on Split MNIST (5 Tasks). Loss surface on the first task (left), second task (middle), and the average loss on all the 5 tasks (right). The black stars highlight the task boundaries. More details on the 2D projections and additional plots are available in the Appendix.

connected with the constraint ρ . By decreasing ρ to have a stronger constraint and save buffer information, we get an increase in λ . For class-incremental problems, λ increases with the number of different classes encountered. In domain-incremental settings, it grows in time (as done with τ).

The composition of all these pieces together forms the **Online Curvature-Aware Replay (OCAR)**: at each step, the method is the optimal step of a second-order replay-based optimization problem, approximated with the FIM made computationally viable via K-FAC, with some fundamental adjustments to make it work in practice. The algorithm can be found in the Appendix A.

5. Continual Stability in OCAR

In this section, we provide a qualitative analysis in a simplified setting, showing how OCAR results in a smoother continual optimization compared to ER and how α , τ , and their ratio can be used to control stability and plasticity.

Loss landscape and model trajectories: We can visualize the improvements in the optimization trajectory of OCAR in a simple continual learning setting. We train a small feed-forward network with ER and OCAR on Split MNIST (5 Tasks). Given the small size of the model, we can store the entire training history of the model, which allows us to plot 2D projections of the model trajectory in the loss landscape (Figure 1) (details in Appendix F).

Looking at their task-wise and joint loss surfaces in figures 1a and 1b, OCAR shows a much smoother model trajectory

across all analyzed loss landscapes (Task 1, Task 2, and Average loss), which results in consistent improvements over time (plasticity) and mild forgetting (stability). Second-order information moves the optimization directly toward the next minimum (black stars in the plots), in fewer steps. On the other hand, ER always suffers from instability at the task boundaries (right after black stars in the plots) which results in an abrupt deviation from the optimal path. The learning curves on the first task (Fig. 10b) and the average of all tasks (Fig. 10a), available in Appendix G, confirm the result. OCAR maintains a smoother learning path without experiencing any stability gap on the first task and ending the stream with higher overall accuracy. ER instead, while still performing well in the basic MNIST setting, suffers from much more instability during training.

Role of the Eigenvectors and Hyperparameters: One approach to understand OCAR’s effect is to study how the eigenvalues of the matrix $\alpha(F_N + (1 + \lambda)F_B + \tau I)^{-1}$ are related to the hyperparameters α , λ , and τ . We can diagonalize $\bar{F} = F_N + (1 + \lambda)F_B = Q\Sigma Q^T$, where Q is a unitary matrix whose columns are the eigenvectors and Σ a diagonal matrix with the eigenvalues σ_i . Therefore, we find that the eigenvalues of $\alpha(F_N + (1 + \lambda)F_B + \tau I)^{-1}$ are $\sigma_i^* = \frac{\alpha}{\sigma_i + \tau}$. In the new coordinate system defined by the eigenvectors Q , we can interpret OCAR as slowing or accelerating directions depending on their curvature. Directions with $\sigma_i^* \ll 1$ ($\sigma_i^* \gg 1$) correspond to directions with high (low) curvature for some tasks. The learning rate α and the Tikhonov damping strength τ rescale these eigenvalues (as shown in Figure 9 in Appendix). In particular, a learning rate $\alpha < 1$ decreases the step size. Conversely, τ mitigates the acceleration caused by small eigenvalues. When $\sigma_i \rightarrow 0$, $\sigma_i^* \rightarrow \frac{\alpha}{\tau}$, limiting the maximum acceleration in each direction. When $\sigma_i \gg \tau$, we have $\sigma_i^* \approx \frac{\alpha}{\sigma_i}$, which is approximately independent of τ .

Empirically, we find that the spectrum spans several orders of magnitude, with the smallest eigenvalues close to zero and a small set of very large values around $(10^4, 10^6)$ (consistently with the common intuition behind methods such as EWC, which expect few important parameters).

Stability-Plasticity tradeoff of OCAR hyperparameters: An interpretation of OCAR hyperparameters in the stability-plasticity tradeoff can be given. While higher λ gives more importance to the FIM of the buffer (stability) and higher values of α allow larger learning steps (plasticity), the role of τ and $\frac{\alpha}{\tau}$ is less intuitive. First, higher values of τ are needed for the introduction of new classes, that creates instability due to the FIM being the variance of gradients (See the Appendix C for a detailed explanation). Figure 2 shows the results of a grid search on α and $\frac{\alpha}{\tau}$ (Complete figure in Appendix fig 11). We can make some empirically-based speculations: (1) the ratio $\frac{\alpha}{\tau}$ controls the learning

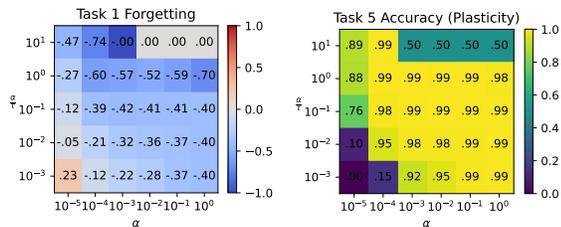


Figure 2. Grid search over α and $\frac{\alpha}{\tau}$: (left) forgetting on the first task, (right) plasticity measured as the accuracy on the final task. Metrics are computed on the test stream at the end of training.

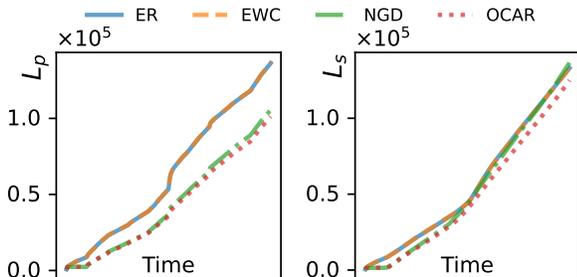


Figure 3. *Left*: L_p Cumulative loss of single batches. *Right*: L_s Cumulative loss measured on all previous 5 data of the stream.

plasticity. As a result, bottom-left elements in Figure 2 show less plasticity; (2) $\frac{\alpha}{\tau}$ controls the effective step size in the eigendirections with small eigenvalues, problematic at task boundaries. If the ratio is too low, OCAR may use updates that are too large, resulting in instability (top-right); (3) Decreasing $\frac{\alpha}{\tau}$ consistently increases forgetting; (4) After a sufficient number of steps after a task drift, the FIM becomes more stable, and $\frac{\alpha}{\tau}$ becomes less important in the learning dynamics. At this point, small α will result in conservative updates, with low plasticity and high stability. In the figure, we notice that for equal values of the ratio, models may have a similar accuracy but a different stability-plasticity tradeoff, depending on their α (low=stability, high=plasticity).

6. Experiments

6.1. Comparison Between EWC, NGD, and OCAR in a Convex Setting

First, we compare OCAR with alternative uses of the Fisher Information, that are not commonly considered for OCL, in a small-scale convex setting. In CL, EWC (Kirkpatrick et al., 2017) generated a cascade of derived methods (Chaudhry et al., 2018; Liu et al., 2018; Huszár, 2017) based on the idea to add a quadratic regularization term to the loss, penalizing the movement of the parameters from an optimal configuration, weighted by the FIM. This approach finds some limi-

tations in OCL when no task boundaries are provided and it’s not possible to select the previous task’s ”best” weights (Mai et al., 2022), breaking the fundamental assumption behind the Laplace approximation. Additionally, regularizing the loss does not directly speed up non-important directions, a non-optimal approach in OCL. On the other hand, outside CL, the Natural Gradient Descent (NGD) (Amari, 1998) uses the FIM as a preconditioner for the gradient, slowing it down in the direction of high curvature and accelerating it in others. We believe NGD can be well-suited for OCL problems. The problem is that raw NGD is derived for i.i.d. settings. OCAR, on the other hand, is an adaptation for non-i.i.d. problems. To underline the differences, we tested the three approaches in combination with ER in an online stream of 10 small convex tasks (all details in Appendix D). We measure the cumulative loss experienced on single batches during the training $L_p = \sum_t L(y_t, \hat{y}_t)$ to measure the ability to adapt to current data and the cumulative loss experienced on all previous data $L_s = \sum_t L(y_{0:t}, \hat{y}_{0:t})$ to measure the stability of the model. The results in figure 3 shows that, while NG is much more adaptable than EWC, it is slightly less stable. EWC performance in OCL is very similar to the one of basic ER, a result aligned with (Mai et al., 2022). OCAR, thanks to its explicit memory constraint, and its dynamic hyperparameters is able to improve both on the speed and on the stability, showing a slight optimization superiority already in this very basic setting.

6.2. Literature benchmarks and SOTA comparison

To ensure reproducibility and a fair comparison we use the same code and experimental setup in (Soutif-Cormerais et al., 2023), later used by LPR (Yoo et al., 2024), an ICML24 paper, that represents the SOTA in OCL and our main ”competitor”. We use Avalanche (Carta et al., 2023) and nngemetry (George, 2021). In line with the literature Split-CIFAR100 (20 tasks) and Split-TinyImageNet (20 tasks), are used as task incremental benchmarks. Following (Yoo et al., 2024) we experimented also on Online CLEAR (10 tasks) (Lin et al., 2021), a domain incremental scenario fundamentally different from the other two datasets. All the methods use a reservoir sampling buffer with 2000 samples for Split-CIFAR100 and online CLEAR and 4000 for Split-TinyImageNet. More details about the experimental setting are available in the Appendix B with specific information for CLEAR in Appendix E. Our entire code for the experiments can be found at https://github.com/edo-urettini/CL_stability.

We compare our approach with a large set of CL baselines (Table 1), evaluated on 4 fundamental OCL metrics: **WC-Acc** (De Lange et al., 2023): the worst-case accuracy is a metric of the stability of the model, measuring a trade-off between the model accuracy and the minimum accuracy among all tasks encountered. It is measured as a weighted

Online Curvature-Aware Replay

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.d	35.3 \pm 1.5	-	-	45.8 \pm 0.6	26.5 \pm 0.6	-	-	34.3 \pm 0.5
ER(Chaudhry et al., 2019b)	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
GDumb(Prabhu et al., 2020)	18.5 \pm 0.5	-	-	-	13.1 \pm 0.4	-	-	-
AGEM(Chaudhry et al., 2019a)	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 + LwF(Li & Hoiem, 2017)	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(Aljundi et al., 2019a)	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
RAR(Kumari et al., 2022)	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
DER++ (Buzzega et al., 2020)	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
ER-ACE (Caccia et al., 2022)	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
SCR(Mai et al., 2021)	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
OnPro(Wei et al., 2023)	31.7 \pm 1.2	36.6 \pm 2.5	12.2 \pm 1.1	-	17.1 \pm 1.5	24.2 \pm 0.4	8.00 \pm 0.8	-
OCM(Guo et al., 2022)	30.9 \pm 0.7	33.3 \pm 1.9	14.9 \pm 0.4	-	20.6 \pm 0.6	24.8 \pm 1.1	10.9 \pm 0.5	-
LPR (Yoo et al., 2024)	33.3 \pm 0.6	42.5 \pm 0.5	19.3 \pm 0.3	-	23.1 \pm 0.2	34.9 \pm 0.4	16.2 \pm 0.2	-
OCAR (ours)	34.9 \pm 0.6	48.2 \pm 1.2	25.0 \pm 1.1	46.2 \pm 0.6	21.7 \pm 1.0	38.3 \pm 1.4	17.4 \pm 0.6	38.3 \pm 0.6
OCAR-DER++ (ours)	34.3 \pm 1.1	46.8 \pm 1.7	25.4 \pm 0.8	46.0 \pm 0.8	-	-	-	-
OCAR-ACE (ours)	<u>35.6</u> \pm 1.2	<u>48.7</u> \pm 1.7	<u>26.5</u> \pm 0.4	44.1 \pm 0.7	<u>25.6</u> \pm 0.4	<u>39.8</u> \pm 2.0	<u>21.5</u> \pm 0.9	34.7 \pm 0.3

Table 1. Results on Split CIFAR100 (20 Tasks) and Split Tiny ImageNet (20 Tasks). Best in bold for the base methods. Best underlined for all methods included OCAR+OTHER.

sum between the accuracy of the current model on the current task and the *min-ACC*. *Min-ACC* is the average of the absolute minimum accuracy that has been reached on previous tasks after they have been learned. Hence, the metric is worse the stronger the stability gap is. **AAA** (Caccia et al., 2022): The Average Anytime Accuracy is a metric developed specifically for OCL, measuring the mean accuracy of the model in time on all the encountered tasks. If we evaluate the model on all encountered tasks at each time step (as shown in the plots in Appendix G), the AAA is the mean of these evaluations. It is a measure of the model’s overall accuracy history. **Acc** (Lopez-Paz & Ranzato, 2017): The final average accuracy is a snapshot of the accuracy of the model at the end of the entire stream. If again we evaluate at each step the model on all encountered tasks, *Acc* would be equal to the last of these evaluations. **Probed Acc** (Davari et al., 2022): The metric is the final accuracy after freezing the feature extractor and retraining only the linear classifier on all the training data. It is a measure of the representation quality of the model.

6.3. Results

Using the same code and setup, all results of other methods are taken from (Soutif-Cormerais et al., 2023) and (Yoo et al., 2024) (for LPR), except for OnPro (Wei et al., 2023) and OCM (Guo et al., 2022), where we reused the original code and run the experiments following the setup in (Soutif-Cormerais et al., 2023; Yoo et al., 2024). Everything is evaluated on 5 runs (except LPR in 10).

Continual metrics: When evaluated at every step in time, OCAR obtains the best AAA and WC-Acc among the methods in both benchmarks. This confirms OCAR as a robust, continual optimization approach with high accuracy and

stability at every point in time. Both metrics have shown significant improvement, with a jump of several points.

End of training accuracy: While in Split-Cifar100 OCAR is able to achieve the best *Acc* among base models, ER-ACE takes the crown in Tinyimagenet (Table 1). However, we must point out that *Acc* is a very poor evaluation metric for the OCL performance. Since it is evaluated at the very final iteration, it badly represents the whole training and can be affected by noise or overly stable methods. We believe that ER-ACE and other methods have artificially lower plasticity, learning less than OCAR from each task, and that this approach can give an advantage at the very end of the stream as we approach the “learning limits” of the model. Of course, this works only for a specific stream length, while OCAR showed increased performance on the overall stream (AAA and WC-Acc).

Probed Accuracy: OCAR obtains the best results on linear probing in both benchmarks, underlying the optimization improvements also on the feature extractor. This is a significant result when compared to the i.i.d. case: while the classifier suffers from forgetting, OCAR can learn a better feature representation if compared to the i.i.d. case, confirming how it learns efficiently in nonstationary settings.

Integration with other methods: Being ER-ACE the method with the best accuracy on Tinyimagenet, we try a combination with it to see if OCAR can improve the results. Even if ER-ACE uses a slightly modified loss, the integration with OCAR works very well. The combination OCAR-ACE beats all other methods (including base OCAR) on *Acc*, AAA and WC-Acc. What is surprising is that this combination is able to overcome even the i.i.d. accuracy for Split-Cifar100 while being very near to it for Split-TinyImagenet. Both ER-ACE and OCAR-ACE show a

lower probing accuracy, which suggests again that ER-ACE may be tuned to prefer stability over plasticity, learning less transferable features despite the high accuracy. Possibly, the robustness of the classifier given by ER-ACE slightly interferes with deeper feature learning. Following this experiment, we also tried the combination with DER++. Unfortunately, we have found OCAR-DER to be much less stable (failing optimization on Tinyimagent), which we conjecture to be caused by the entropy regularization implicit in the DER loss, which results in a loss that is "too different" from the KL divergence and breaks the assumption for the use of the Fisher.

Online CLEAR The Online CLEAR benchmark is fundamentally different: no classes are added from task to task, but all classes evolve in time, in a domain incremental fashion. This puts much more importance on forward and backward transfer and less on catastrophic forgetting due to the similarities between tasks. For this reason, the final accuracy is higher than AAA (see figure 8 in Appendix). This setting, being more similar to the standard iid one, makes the estimation and the stability of the FIM much easier, making OCAR remarkably better than the previous SOTA LPR (see table 2). This confirms OCAR robustness also on domain incremental settings, underlying its optimization improvements.

Method	Online CLEAR (10 Tasks)		
	Acc \uparrow	AAA ^{val} \uparrow	WC-Acc ^{val} \uparrow
ER	63.1 \pm 0.7	58.9 \pm 0.8	47.7 \pm 1.6
LPR	65.2 \pm 0.9	63.5 \pm 1.0	62.6 \pm 0.7
OCAR(Ours)	75.3 \pm 0.8	73.9 \pm 0.5	70.3 \pm 0.5

Table 2. Results on Online CLEAR (10 Tasks) domain incremental setting. 2000 Buffer size. Best in bold.

Method	Training Time (seconds)
ER	14
ER + LWF	15
MIR	31
ER-ACE	17
DER	17
RAR	72
SCR	131
LPR	213
OCAR (Ours)	38

Table 3. Training Time for the First Task on Split-CIFAR-100.

Final Comment: OCAR showed remarkable performance across all continual metrics, improving both on task-

incremental and domain-incremental the previous SOTA results, including the ICML24 paper LPR (Yoo et al., 2024), using the same code, setting, and benchmarks. All of this is done efficiently: in table 3 the first-task training times are presented, showing how OCAR even improves on speed with respect to some previous SOTA approaches (Appendix B for more information). The method showed the possibility of being combined with ER-ACE, obtaining even stronger results.

7. Conclusion

In this paper, we revisit replay-based OCL as a sequence of second-order optimizations with hard stability constraints and information geometry rooting. The resulting method OCAR shows consistent improvements in plasticity and continual stability both in class incremental and domain incremental scenarios, with clear hyperparameter interpretation in the stability-plasticity tradeoff. The combination of OCAR and ER-ACE obtained even stronger results, being also capable of reaching similar final accuracies as the i.i.d. training in the class incremental setting. Future research directions include the comparison of different approximations of the curvature (e.g. George et al. (2018)), alternative derivations for the optimization problem (e.g. Benzing (2022)), and feasible dynamic adaptation of α and τ . We believe OCAR can be the starting point for further improvements towards a deeper understanding of continual learning dynamics and the design of a flexible and general continual optimizer.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

Acknowledgements

This paper has been partially supported by the CoEvolution project, funded by EU Horizon 2020 under GA n 101168559.

References

Aljundi, R., Belilovsky, E., Tuytelaars, T., Charlin, L., Caccia, M., Lin, M., and Page-Caccia, L. Online continual learning with maximal interfered retrieval. *Advances in neural information processing systems*, 32, 2019a.

Aljundi, R., Kelchtermans, K., and Tuytelaars, T. Task-free continual learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,

- pp. 11254–11263, 2019b.
- Amari, S.-I. Natural gradient works efficiently in learning. *Neural computation*, 10(2):251–276, 1998.
- Amari, S.-i. *Information geometry and its applications*, volume 194. Springer, 2016.
- Benzing, F. Gradient descent on neurons and its link to approximate second-order optimization. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvári, C., Niu, G., and Sabato, S. (eds.), *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pp. 1817–1853. PMLR, 2022. URL <https://proceedings.mlr.press/v162/benzing22a.html>.
- Buzzega, P., Boschini, M., Porrello, A., Abati, D., and Calderara, S. Dark experience for general continual learning: a strong, simple baseline. *Advances in neural information processing systems*, 33:15920–15930, 2020.
- Caccia, L., Aljundi, R., Asadi, N., Tuytelaars, T., Pineau, J., and Belilovsky, E. New insights on reducing abrupt representation change in online continual learning. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL <https://openreview.net/forum?id=N8MaByOzUfb>.
- Carta, A., Pellegrini, L., Cossu, A., Hemati, H., and Lomonaco, V. Avalanche: A pytorch library for deep continual learning. *Journal of Machine Learning Research*, 24(363):1–6, 2023.
- Chaudhry, A., Dokania, P. K., Ajanthan, T., and Torr, P. H. S. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In Ferrari, V., Hebert, M., Sminchisescu, C., and Weiss, Y. (eds.), *Computer Vision - ECCV 2018 - 15th European Conference, Munich, Germany, September 8-14, 2018, Proceedings, Part XI*, volume 11215 of *Lecture Notes in Computer Science*, pp. 556–572. Springer, 2018. doi: 10.1007/978-3-030-01252-6_33. URL https://doi.org/10.1007/978-3-030-01252-6_33.
- Chaudhry, A., Ranzato, M., Rohrbach, M., and Elhoseiny, M. Efficient lifelong learning with A-GEM. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019a. URL https://openreview.net/forum?id=Hkf2_sC5FX.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P., Torr, P., and Ranzato, M. Continual learning with tiny episodic memories. In *Workshop on Multi-Task and Lifelong Reinforcement Learning*, 2019b.
- Davari, M., Asadi, N., Mudur, S., Aljundi, R., and Belilovsky, E. Probing representation forgetting in supervised and unsupervised continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16712–16721, 2022.
- Daxberger, E., Swaroop, S., Osawa, K., Yokota, R., Turner, R. E., Hernández-Lobato, J. M., and Khan, M. E. Improving continual learning by accurate gradient reconstructions of the past. *Transactions on Machine Learning Research*, 2023.
- De Lange, M., van de Ven, G., and Tuytelaars, T. Continual evaluation for lifelong learning: Identifying the stability gap, 2023.
- Dohare, S., Hernandez-Garcia, J. F., Lan, Q., Rahman, P., Mahmood, A. R., and Sutton, R. S. Loss of plasticity in deep continual learning. *Nat.*, 632(8026):768–774, 2024. doi: 10.1038/S41586-024-07711-7. URL <https://doi.org/10.1038/s41586-024-07711-7>.
- Durbin, J. and Koopman, S. J. *Time series analysis by state space methods*, volume 38. OUP Oxford, 2012.
- Farajtabar, M., Azizan, N., Mott, A., and Li, A. Orthogonal gradient descent for continual learning. In Chiappa, S. and Calandra, R. (eds.), *The 23rd International Conference on Artificial Intelligence and Statistics, AISTATS 2020, 26-28 August 2020, Online [Palermo, Sicily, Italy]*, volume 108 of *Proceedings of Machine Learning Research*, pp. 3762–3773. PMLR, 2020. URL <http://proceedings.mlr.press/v108/farajtabar20a.html>.
- French, R. M. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 3(4):128–135, 1999.
- George, T. NNGeometry: Easy and Fast Fisher Information Matrices and Neural Tangent Kernels in PyTorch, February 2021. URL <https://doi.org/10.5281/zenodo.4532597>.
- George, T., Laurent, C., Bouthillier, X., Ballas, N., and Vincent, P. Fast approximate natural gradient descent in a kronecker factored eigenbasis. In Bengio, S., Wallach, H. M., Larochelle, H., Grauman, K., Cesa-Bianchi, N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pp. 9573–9583, 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/48000647b315f6f00f913caa757a70b3-Abstract.html>.

- Guo, Y., Liu, B., and Zhao, D. Online continual learning through mutual information maximization. In *International conference on machine learning*, pp. 8109–8126. PMLR, 2022.
- Hess, T., Tuytelaars, T., and van de Ven, G. M. Two complementary perspectives to continual learning: Ask not only what to optimize, but also how. *CoRR*, abs/2311.04898, 2023. doi: 10.48550/ARXIV.2311.04898. URL <https://doi.org/10.48550/arXiv.2311.04898>.
- Huszár, F. On quadratic penalties in elastic weight consolidation. *CoRR*, abs/1712.03847, 2017. URL <http://arxiv.org/abs/1712.03847>.
- Kamath, S., Soutif-Cormerais, A., van de Weijer, J., and Raducanu, B. The expanding scope of the stability gap: Unveiling its presence in joint incremental learning of homogeneous tasks. *CoRR*, abs/2406.05114, 2024. doi: 10.48550/ARXIV.2406.05114. URL <https://doi.org/10.48550/arXiv.2406.05114>.
- Kang, Z., Wang, L., Zhang, X., and Alahari, K. Advancing prompt-based methods for replay-independent general continual learning. *arXiv preprint arXiv:2503.00677*, 2025.
- Kao, T., Jensen, K. T., van de Ven, G., Bernacchia, A., and Hennequin, G. Natural continual learning: success is a journey, not (just) a destination. In Ranzato, M., Beygelzimer, A., Dauphin, Y. N., Liang, P., and Vaughan, J. W. (eds.), *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual*, pp. 28067–28079, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/ec5aa0b7846082a2415f0902f0da88f2-Abstract.html>.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences*, 114(13):3521–3526, 2017.
- Kumari, L., Wang, S., Zhou, T., and Bilmes, J. A. Retrospective adversarial replay for continual learning. *Advances in neural information processing systems*, 35:28530–28544, 2022.
- Kunstner, F., Hennig, P., and Balles, L. Limitations of the empirical fisher approximation for natural gradient descent. *Advances in neural information processing systems*, 32, 2019.
- Lange, M. D., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G. G., and Tuytelaars, T. A continual learning survey: Defying forgetting in classification tasks. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(7):3366–3385, 2022. doi: 10.1109/TPAMI.2021.3057446. URL <https://doi.org/10.1109/TPAMI.2021.3057446>.
- Lange, M. D., van de Ven, G. M., and Tuytelaars, T. Continual evaluation for lifelong learning: Identifying the stability gap. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023. URL <https://openreview.net/forum?id=Zy350cRstc6>.
- LeCun, Y., Bengio, Y., and Hinton, G. Deep learning. *nature*, 521(7553):436–444, 2015.
- Li, Z. and Hoiem, D. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12):2935–2947, 2017.
- Lin, Z., Shi, J., Pathak, D., and Ramanan, D. The clear benchmark: Continual learning on real-world imagery. In *Thirty-fifth conference on neural information processing systems datasets and benchmarks track (round 2)*, 2021.
- Liu, X., Masana, M., Herranz, L., Van de Weijer, J., Lopez, A. M., and Bagdanov, A. D. Rotate your networks: Better weight consolidation and less catastrophic forgetting. In *2018 24th International Conference on Pattern Recognition (ICPR)*, pp. 2262–2268. IEEE, 2018.
- Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. In Guyon, I., von Luxburg, U., Bengio, S., Wallach, H. M., Fergus, R., Vishwanathan, S. V. N., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pp. 6467–6476, 2017. URL <https://proceedings.neurips.cc/paper/2017/hash/f87522788a2be2d171666752f97ddeb-Abstract.html>.
- Magistri, S., Trinci, T., Soutif-Cormerais, A., van de Weijer, J., and Bagdanov, A. D. Elastic feature consolidation for cold start exemplar-free incremental learning. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=7D9X2cFnt1>.

- Mai, Z., Li, R., Kim, H., and Sanner, S. 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*, pp. 3589–3599, 2021.
- Mai, Z., Li, R., Jeong, J., Quispe, D., Kim, H., and Sanner, S. Online continual learning in image classification: An empirical survey. *Neurocomputing*, 469:28–51, 2022. doi: 10.1016/J.NEUCOM.2021.10.021. URL <https://doi.org/10.1016/j.neucom.2021.10.021>.
- Martens, J. New insights and perspectives on the natural gradient method. *Journal of Machine Learning Research*, 21(146):1–76, 2020.
- Martens, J. and Grosse, R. Optimizing neural networks with kronecker-factored approximate curvature. In *International conference on machine learning*, pp. 2408–2417. PMLR, 2015.
- Martens, J. and Sutskever, I. Training deep and recurrent networks with hessian-free optimization. In *Neural Networks: Tricks of the Trade: Second Edition*, pp. 479–535. Springer, 2012.
- Masana, M., Liu, X., Twardowski, B., Menta, M., Bagdanov, A. D., and van de Weijer, J. Class-incremental learning: Survey and performance evaluation on image classification. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(5):5513–5533, 2023. doi: 10.1109/TPAMI.2022.3213473. URL <https://doi.org/10.1109/TPAMI.2022.3213473>.
- Mirzadeh, S., Farajtabar, M., Pascanu, R., and Ghasemzadeh, H. Understanding the role of training regimes in continual learning. In Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020. URL <https://proceedings.neurips.cc/paper/2020/hash/518a38cc9a0173d0b2dc088166981cf8-Abstract.html>.
- Mirzadeh, S., Farajtabar, M., Görür, D., Pascanu, R., and Ghasemzadeh, H. Linear mode connectivity in multi-task and continual learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL https://openreview.net/forum?id=Fmg_fQYUejf.
- Moon, J.-Y., Park, K.-H., Kim, J. U., and Park, G.-M. Online class incremental learning on stochastic blurry task boundary via mask and visual prompt tuning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 11731–11741, 2023.
- Ollivier, Y., Arnold, L., Auger, A., and Hansen, N. Information-geometric optimization algorithms: A unifying picture via invariance principles. *Journal of Machine Learning Research*, 18(18):1–65, 2017.
- Pan, P., Swaroop, S., Immer, A., Eschenhagen, R., Turner, R., and Khan, M. E. E. Continual deep learning by functional regularisation of memorable past. *Advances in neural information processing systems*, 33:4453–4464, 2020.
- Prabhu, A., Torr, P. H., and Dokania, P. K. 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*, pp. 524–540. Springer, 2020.
- Saha, G., Garg, I., and Roy, K. Gradient projection memory for continual learning. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net, 2021. URL <https://openreview.net/forum?id=3AOj0RCNC2>.
- Soutif-Cormerais, A., Carta, A., Cossu, A., Hurtado, J., Lomonaco, V., van de Weijer, J., and Hemati, H. A comprehensive empirical evaluation on online continual learning. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023 - Workshops, Paris, France, October 2-6, 2023*, pp. 3510–3520. IEEE, 2023. doi: 10.1109/ICCVW60793.2023.00378. URL <https://doi.org/10.1109/ICCVW60793.2023.00378>.
- Thomas, M. and Joy, A. T. *Elements of information theory*. Wiley-Interscience, 2006.
- van de Ven, G. M. and Tolias, A. S. Three scenarios for continual learning. *CoRR*, abs/1904.07734, 2019. URL <http://arxiv.org/abs/1904.07734>.
- Wei, Y., Ye, J., Huang, Z., Zhang, J., and Shan, H. Online prototype learning for online continual learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 18764–18774, 2023.
- Yoo, J., Liu, Y., Wood, F., and Pleiss, G. Layerwise proximal replay: a proximal point method for online continual learning. In *Proceedings of the 41st International Conference on Machine Learning, ICML’24*. JMLR.org, 2024.
- Yuan, K., Ying, B., and Sayed, A. H. On the influence of momentum acceleration on online learning. *Journal of Machine Learning Research*, 17(192):1–66, 2016.

A. Algorithm

Algorithm 1. Online Curvature-Aware Replay (OCAR)

Input: network parameters w , learning rate α , per batch gradient steps count S , Tikhonov increase $\Delta\tau$, EMA parameter α_{EMA} .

Output: trained network parameters w .

```

1: Initialize replay buffer:  $\mathcal{B} \leftarrow \{\}$ .
2:  $\tau \leftarrow \alpha$ 
3:  $\lambda \leftarrow 1$ 
4: for  $t \in \{1, \dots, \infty\}$  do
5:   Obtain new data batch  $N_t$ .
6:   Sample buffer data batch  $B_t$ 
7:   for  $s \in \{1, \dots, S\}$  do
8:     Compute loss  $\mathcal{L}(w)$  using  $N_t$  and  $B_t$ .
9:     Compute loss gradient  $\nabla\mathcal{L}(w)$ .
10:     $\tau \leftarrow \tau + \Delta\tau$ 
11:    if Class Incremental then
12:      Update known classes list with  $N_t$ 
13:      Increase  $\lambda$  with new classes
14:    else
15:       $\lambda \leftarrow \lambda + \Delta\tau$ 
16:    end if
17:    if  $s = 1$  then
18:      Compute K-FAC factors  $A$  and  $G$  with  $N_t$  and  $B_t$  ( $B_t$  influence weighted by  $\lambda$ )
19:      for  $l \in 1, \dots, L$  do
20:         $A_{EMA,l} \leftarrow (1 - \alpha_{EMA})A_{EMA,l} + \alpha_{EMA}A_l$ 
21:         $G_{EMA,l} \leftarrow (1 - \alpha_{EMA})G_{EMA,l} + \alpha_{EMA}G_l$ 
22:        if  $l = L$  and  $L$  changed shape then
23:           $G_{EMA,l} \leftarrow G_l$ 
24:        end if
25:      end for
26:       $F_{EMA} \leftarrow A_{EMA}, G_{EMA}$ 
27:       $F_{INV} \leftarrow (F_{EMA} + \tau\mathbf{I})^{-1}$ 
28:    end if
29:     $\tilde{\nabla}\mathcal{L}(w) \leftarrow F_{INV}\nabla\mathcal{L}(w)$ 
30:     $w \leftarrow w - \alpha\tilde{\nabla}\mathcal{L}(w)$ 
31:  end for
32:   $\mathcal{B} \leftarrow \text{Reservoir.update}(\mathcal{B}, N_t, \text{maxsize})$ 
33: end for

```

The K-FAC computations are done after weighting with λ the data related to the buffer.

B. Experimental Setup

Code The entire code used to perform the full experiments (CIFAR100, TinyImageNet e Clear) can be found in the anonymous repository at https://github.com/edo-urettini/CL_stability. The anonymization process could have removed some important parts of the code but this is highly unlikely. The code was built on a combination of the repository for the OCL survey (Soutif-Cormerais et al., 2023) and the nngometry repository (George, 2021).

Hardware The experiments were performed in a Linux cluster equipped with Nvidia Tesla V100 16GB GPUs and Intel Xeon Gold 6140M CPUs. The training was not parallelized between the GPUs.

Main training setup All the experiments have been performed as in (Soutif-Cormerais et al., 2023). We only report the most important aspects here, referring to the original paper for all the additional details. The datasets used for the main experiments are Split-CIFAR100 and Split-TinyImageNet. Both are class-incremental benchmarks, divided in 20 tasks, with a number of new classes encountered at each new experience (5 in Split-CIFAR100 and 10 in Split-TinyImageNet). The model used for the main experiments is a Slim-ResNet18 trained using SGD. SGD can help the model adapt faster when new data are encountered after a task boundary, while ADAM (Kingma & Ba, 2015) will need some time to update its running statistics. The batches are tiny, with 10 examples from the current portion of the stream and 10 from the replay buffer. The replay buffer for Split-CIFAR100 was kept at 2000 samples, while for Split-TinyImageNet at 4000. The metrics computed at the end of the training (cumulative accuracy and linear probing accuracy) are computed on a test stream. The continual metrics (average anytime accuracy and worst-case accuracy) are computed after the training on each batch on a held-out validation set. The only additional information OCAR has access to is the number of classes on each experience (5 for one and 10 for the other). This information is only used to increase the weight λ for the computation of the Fisher Information on the buffer data. If the buffer data contains n classes and the number of classes observed in the current portion of the stream is k , then the FIM of the buffer will weight $\frac{n}{k}$. While this information is useful, if it is not accessible, other weighting procedures can be performed, or, it can simply be estimated by watching how many classes are observed in the first few steps of the training. Usually few steps are enough to observe all the different classes present in the task.

The same setup has been used in (Yoo et al., 2024), with the only difference they used 10 seeds for their experiments instead of only 5. There is also a possible difference in the hyperparameter selection. Given the equality of the training and validation setup, we compared our methods’ performances directly with the results of (Yoo et al., 2024) and (Soutif-Cormerais et al., 2023). This minimizes the probability of errors or bugs and uses the results obtained in the best conditions.

Training time We report the training time to complete a single task on Split-CIFAR100 of different methods. Since OCAR and other OCL methods do not increase their computational cost over time, we report the training times on the first task to avoid unnecessary retraining of all the methods. We train on the first task without the continual evaluation to remove the evaluation overhead and provide a fair comparison between the methods.

Method	Training Time (seconds)
ER	14
ER + LWF	15
MIR	31
ER-ACE	17
DER	17
RAR	72
SCR	131
LPR	213
OCAR(Ours)	38

Table 4. Training Time for the First Task on Split-CIFAR-100.

OCAR is about 3 times slower than standard ER, which does not perform any additional operations apart from the basic SGD loop. It is also faster than other sophisticated methods such as MIR, SCR, and LPR. We underline that in our tests LPR resulted much slower than what the authors found in their original paper (Yoo et al., 2024). We did not identify the cause of this. It can be something related to our hardware or the implementation details of our code.

Hyperparameters selection As with everything else, the hyper-optimization is performed as in (Soutif-Cormerais et al., 2023). The best configuration is selected as the best performing in the cumulative accuracy metrics on the validation set at the end of the fourth experience. This approach is suboptimal when working on nonstationary data. We notice that in our model, this shorter stream would select a combination of parameters prone to having high plasticity and lower stability. For this reason, instead of selecting a fixed Tikhonov regularizer τ , we instead selected a speed of increase for τ . In this way, we found a much more robust approach. Differently from (Soutif-Cormerais et al., 2023), all our hyperparameters for OCAR have been selected after trying 100 combinations using the tree-structured Parzen estimator algorithm, instead of 200 trials used in the original survey. This can give our method a small disadvantage, but we found it to be sufficient to get good

results. The only exception was OCAR-ACE on Split-TinyImageNet where 200 trials were needed due to the particular complexity of using a loss function not derived from a KL divergence on a complex dataset as TinyImageNet. We want to highlight that in the LPR paper (Yoo et al., 2024), it seems the hyperparameters selection is done by training on all the experiences and selecting the best final accuracy. It is possible this can result in an advantage, avoiding the problem of the short and partial stream used in (Soutif-Cormerais et al., 2023) and for our method.

C. Fisher Matrix Computations

The Fisher Information Matrix can be computed both as the variance of the score (the gradient of the log-likelihood) or as the negative expected value of the curvature of the log-likelihood. For a classification problem we assume the model f_w is predicting a probability vector probability vector $\mathbf{p} = [p_1, \dots, p_K]$ where each p_k is the probability that $y_k = 1$, such that $\mathbf{p} = f(x; w)$. The distribution assumed is then a categorical distribution. On a single example x the FIM can be computed as:

$$\begin{aligned}
 F(w) &= \mathbf{E}_{\mathbf{y} \sim \text{Cat}(\mathbf{y}|x;w)}[\nabla_w^2 \log \text{Cat}(\mathbf{y}|x;w)] = \mathbf{E}_{\mathbf{y} \sim \text{Cat}(\mathbf{y}|x;w)}[\nabla_w^2 \sum_k y_k \log f_w(x)_k] = \\
 &= \mathbf{E}_{\mathbf{y} \sim \text{Cat}(\mathbf{y}|x_n;w)} \left[\begin{array}{c} \sum_k \frac{y_k}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \\ \vdots \\ \sum_k \frac{y_k}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_H} \end{array} \begin{array}{c} \left[\sum_k \frac{y_k}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \right]^T \\ \vdots \\ \left[\sum_k \frac{y_k}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_H} \right]^T \end{array} \right] = \\
 &= \mathbf{E}_{\mathbf{y} \sim \text{Cat}(\mathbf{y}|x_n;w)} \left[\begin{array}{c} \sum_k \frac{y_k^2}{f_w^2(x)_k} \left(\frac{\partial f_w(x)_k}{\partial w_1} \right)^2 \quad \sum_k \frac{y_k^2}{f_w^2(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \frac{\partial f_w(x)_k}{\partial w_2} \cdots \\ \sum_k \frac{y_k}{f_w^2(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \frac{\partial f_w(x)_k}{\partial w_2} \\ \vdots \end{array} \right] = \\
 &= \left[\begin{array}{c} \sum_k \frac{1}{f_w(x)_k} \left(\frac{\partial f_w(x)_k}{\partial w_1} \right)^2 \quad \sum_k \frac{1}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \frac{\partial f_w(x)_k}{\partial w_2} \cdots \\ \sum_k \frac{1}{f_w(x)_k} \frac{\partial f_w(x)_k}{\partial w_1} \frac{\partial f_w(x)_k}{\partial w_2} \\ \vdots \end{array} \right]
 \end{aligned}$$

Note how each element is a sum of the contribution from each class, weighted by the inverse predicted probability for that class. The Fisher computed for a gradient composed of the sum (or mean) of gradients from multiple examples is the sum (or mean) of the individual Fishers. This is due to the required assumption to compute the Fisher: being at the optimum of the log-likelihood optimization, assuming the predicted parameters are the true ones. A consequence of this assumption is that the correlations between gradients of different examples will be zero: $\mathbf{V}[\nabla(\mathbf{y}_i|x_i;w) + \nabla(\mathbf{y}_j|x_j;w)] = \mathbf{V}[\nabla(\mathbf{y}_i|x_i;w)] + \mathbf{V}[\nabla(\mathbf{y}_j|x_j;w)]$. Then, our λ to give more weight to the buffer data is easily implemented. For the same assumption about the optimum, the variance of the score is computed as the expected value of the squared gradients: the squared expected value would be equal to zero.

Last Layer FIM Focusing on the FIM of the last layer and ignoring all the correlations between the weights for a moment, we can compute the diagonal element of a classifier’s weight. Assuming the loss \mathcal{L} is a standard cross entropy loss:

$$\frac{\partial \mathcal{L}}{\partial w_{j,i}} = h_j(p_i - y_i)$$

where p_i is the prediction for the class i and y_i an indicator function for the real value. The variance of the score for this weight is:

$$\mathbf{E} \left[\left(\frac{\partial \mathcal{L}}{\partial w_{j,i}} \right)^2 \right] = \mathbf{E}[h_j^2(p_i - y_i)^2] = h_j^2 p_i(1 - p_i)$$

If we transform the partial derivative with the inverse variance:

$$\mathbf{E} \left[\left(\frac{\partial \mathcal{L}}{\partial w_{j,i}} \right)^2 \right]^{-1} \frac{\partial \mathcal{L}}{\partial w_{j,i}} = \frac{1}{h_j} \frac{p_i - y_i}{p_i(1 - p_i)}$$

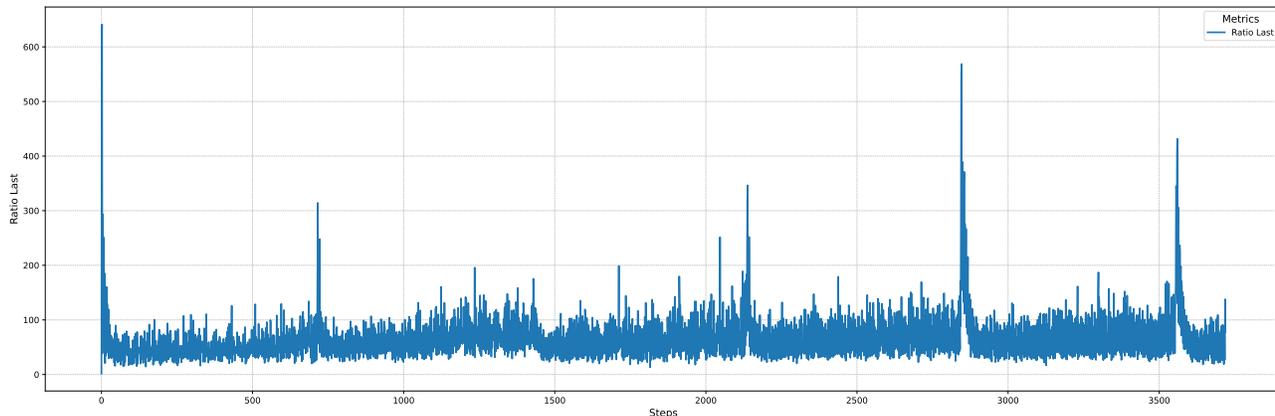


Figure 4. Ratio between the norm of the gradient after being transformed with OCAR and the norm of the original gradient when a small τ is used

In class-incremental settings, at task boundaries, we can assume the probability predicted for the new classes will be pretty low due to the random initialization of the weights. In this case, the gradient of the weights connected to this new unit will be high while the variance of the gradient will be very low, greatly accelerating the gradient, and contributing to possible instabilities due to the assumption of the Fisher of predicting the "true" distribution while, particularly at task boundaries, this is not true. Clearly this analysis ignores all the complex cross-correlations between the weights. To visualize the complete effect, we plot in Figure 4 the ratio between the norm of the gradient after being transformed with OCAR and the norm of the original cross-entropy gradient. We show only a subset of the training on Split-CIFAR100, using a small τ ($1r/100$). All task boundaries are clearly visible, showing the great acceleration the FIM provides when new classes arrive. This underlines the importance of the Tikhonov regularization (τ) in class-incremental settings, where the Fisher can exacerbate some instabilities when the predictions of the model are completely off.

D. Qualitative Experiment

The code of the qualitative experiment is included in our repository. We used a basic linear model with 10 inputs and a single output trained to solve a simple linear regression problem with an MSE loss. 10 tasks with 1000 samples per task are randomly generated. Each task has a different multivariate normal from which input data are sampled and different real weights that should be estimated by the model. The data are accessed as a stream, with a single pass per batch. The Standard Vitter algorithm is used to keep a buffer of old data. Each method uses a batch composed of 10 new data and 10 data sampled from the buffer.

Natural Gradient is applied directly by estimating the full Fisher information of the model with an EMA of the FIM computed on the single batches. This FIM after the EMA is used to precondition the gradient.

OCAR uses the exact same approach but uses a λ to give more weight to old data and uses a scheduling of both λ and τ (the Tikhonov regularize) to increase them both in time. EWC is more tricky to implement in its basic form for online problems. Our approach is very similar to what is done in (Mai et al., 2022), an online extension of the EWC++ strategy (Chaudhry et al., 2018). The difference is that we use also replay data for computing the gradient. Namely, the loss is computed on both old and new data, but the Fisher (being a penalization) is computed only on old data (an EMA is kept also in this case for estimation). Given that no task boundaries are accessible, the penalization of the movement of the weights is done with respect to the weights at the previous step in time. In this way, a more regularized descent should be followed, penalizing the displacement from step to step using the Fisher Information of old data. The loss is a basic MSE but for analysis purposes, the figure D shows the cumulative loss (simply the loss experienced in each batch accumulated in time) and the cumulative loss on all previous data (at each step the loss on all previous data is computed after the optimization step and accumulated in time). The first measure shows a general capability of fast adaptation while the second an ability to find an optimum stable for the entire previous stream.

A hyperparameter selection is performed to select the best setting on the sum of the final value of both these stability and

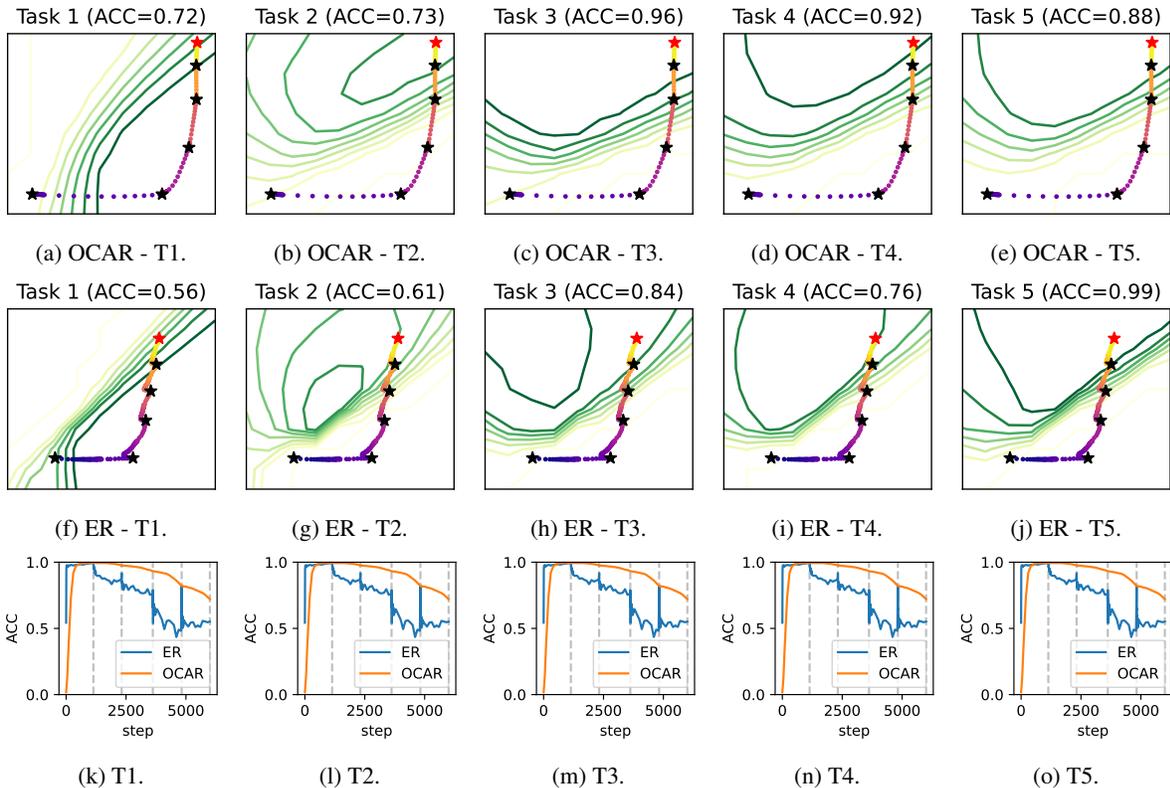


Figure 5. 2D projections of the training trajectories for ER and OCAR on Split MNIST (5 Tasks). The black stars highlight the task boundaries, the red star the final model. We also show learning curves on each task separately.

plasticity measures. Then, on the same data, using the best hyperparameters, the results are averaged across 10 random seeds to test for training stability with different optimization paths.

E. Online CLEAR

Following the very recent work of (Yoo et al., 2024), we tested our method also on the Online CLEAR benchmark (Lin et al., 2021), a domain incremental learning scenario. This scenario is fundamentally different from the class-incremental, with the same classes that undergo some sort of evolution. The CL aspect is then less impactful, with much more forward/backward transfer and less catastrophic forgetting due to the similarity of the tasks in different domains. We followed the same settings as in LPR paper (Yoo et al., 2024), that are very similar to the standard approach used for our main experiments, but with some differences: the use of the full ResNet18 instead of the Slim version and 10 gradient steps per batch. Unfortunately, we encountered some bugs when their exact code was used in ours, requiring some slight modifications. Not being able to ensure the exact same conditions, we rerun all the experiments. We decided to compare the baseline of Experience Replay, the very recent LPR approach, and our method. ER-ACE is tailored for task-incremental settings, so we avoid its use. In this scenario, we increase λ not with the number of classes in the buffer (as in class incremental), but we increase it in time as new batches arrive. LPR was tested using 100 samples from the buffer to estimate its preconditioner.

F. Learning Trajectory on Split MNIST

The model is a small feedforward network with two hidden layers of 100 units and ReLU activations. The model is trained online with 3 passes for each mini-batch with a small replay buffer of 100 elements, corresponding to 10 samples per class by the end of training.

To plot the 2D learning projections in Figure 5, we follow a procedure similar to (Mirzadeh et al., 2021).

We consider the 2D plane that intersects the model’s initialization w_0^* , the model after the first task w_1^* , and the final model w_5^* . Many other possibilities were considered before this choice (e.g. using different tasks or random directions), but they all resulted in qualitatively similar plots. The coordinates system is obtained by orthonormalizing the directions $u = w_1^* - w_0^*$ and $v = w_5^* - w_0^*$, obtaining $\bar{u} = \frac{u}{\|u\|}$ and $\bar{v} = \frac{v - \cos(u,v)u}{\|v - \cos(u,v)u\|}$. Given a model w , its coordinates in the 2D space are the unique $\langle x, y \rangle$ such that $w = x\bar{u} + y\bar{v} + w_0^*$. Notice that each method has different values for w_1^* and w_5^* , and therefore different 2D planes and coordinate systems are chosen for each method. Since the coordinates are not meaningful and cannot be compared across plots, we do not show them in Figure 5.

G. Additional Figures

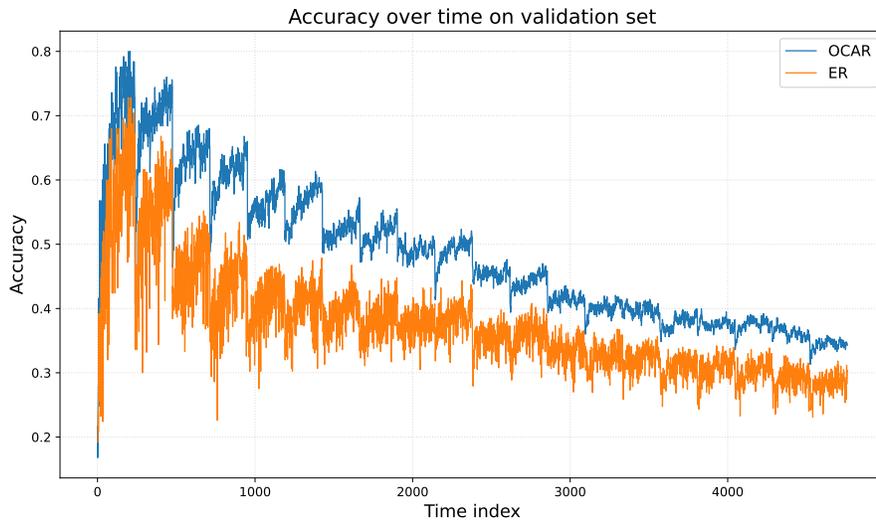


Figure 6. Accuracy over time on the validation set for Split-CIFAR100 experiment.

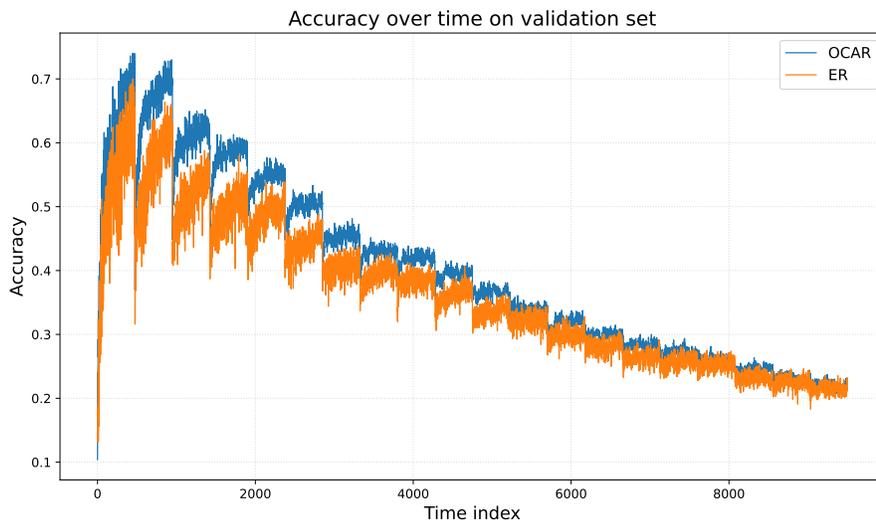


Figure 7. Accuracy over time on the validation set for Split-TinyImagenet experiment.

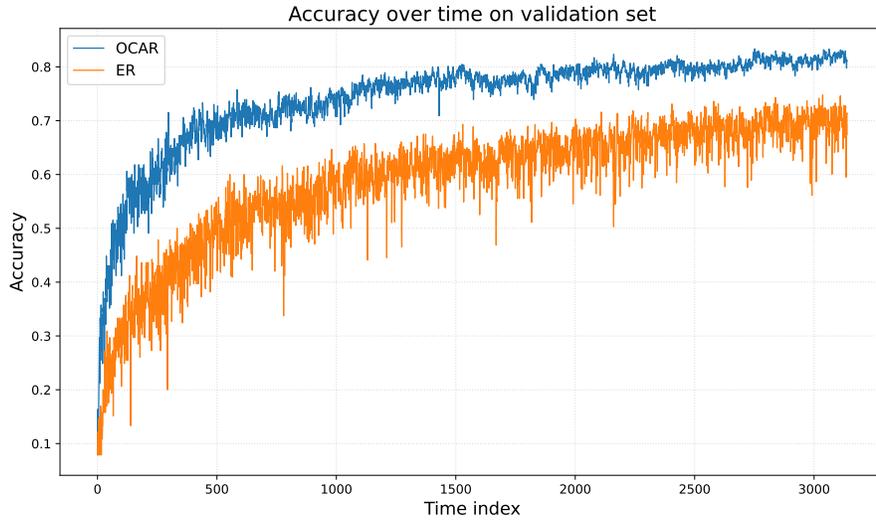


Figure 8. Accuracy over time on the validation set for Online CLEAR experiment.

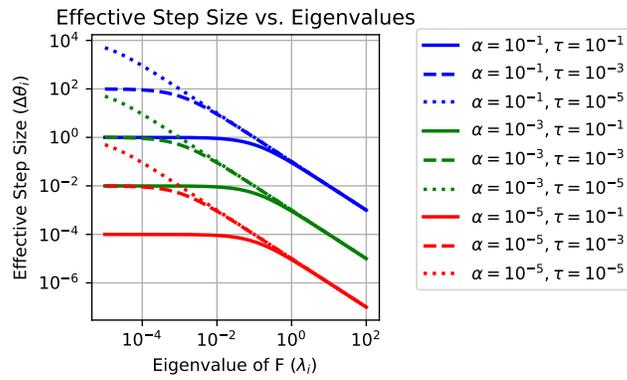


Figure 9. Effective step size against the FIM eigenvalues. The step size in the directions with small eigenvalues is regularized via τ , while large eigenvalues are unaffected by it. The learning rate α affects all the directions equally.

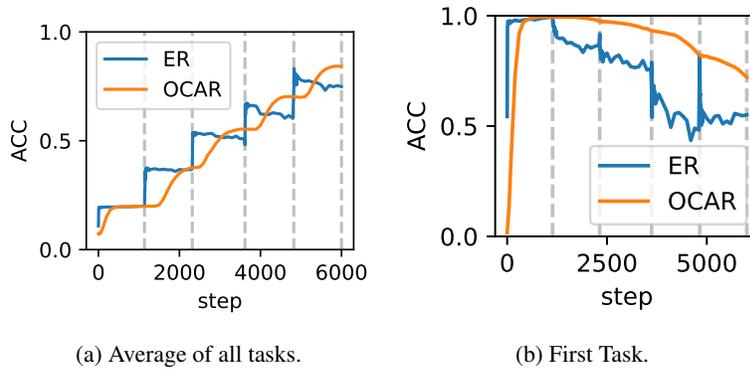


Figure 10. Learning curves on the first task (10b) and the average accuracy for all tasks (10a).

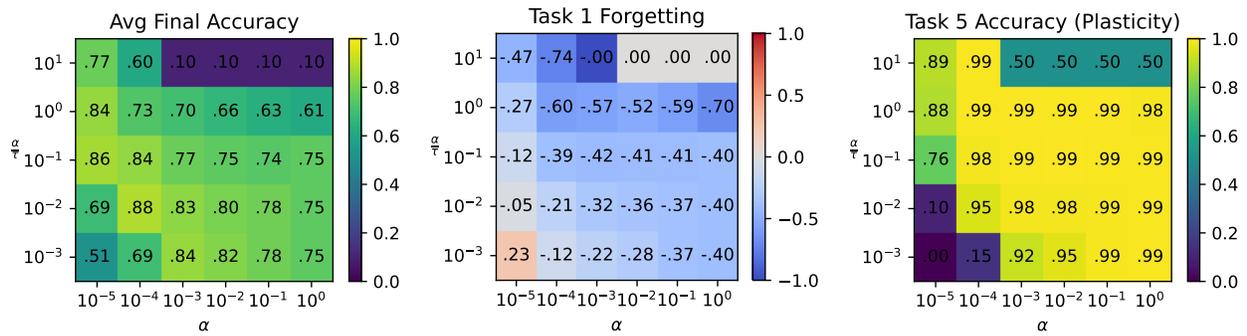


Figure 11. Grid search over α and $\frac{\alpha}{\tau}$: (left) average accuracy, (middle) forgetting on the first task, (right) plasticity measured as the accuracy on the final task. Metrics are computed on the test stream at the end of training.