

# CaSpeR: Latent Spectral Regularization for Continual Learning

Emanuele Frascaoli<sup>1,2</sup> Riccardo Benaglia<sup>1,2</sup> Matteo Boschini<sup>1</sup> Luca Moschella<sup>3</sup> Cosimo Fiorini<sup>2</sup>  
Emanuele Rodolà<sup>3</sup> Simone Calderara<sup>1</sup>

## Abstract

While biological intelligence grows organically as new knowledge is gathered throughout life, Artificial Neural Networks forget catastrophically whenever they face a changing training data distribution. Rehearsal-based Continual Learning (CL) approaches have been established as a versatile and reliable solution to overcome this limitation; however, sudden input disruptions and memory constraints are known to alter the consistency of their predictions. We study this phenomenon by investigating the geometric characteristics of the learner’s latent space and find that replayed data points of different classes increasingly mix up, interfering with classification. Hence, we propose a geometric regularizer that enforces weak requirements on the Laplacian spectrum of the latent space, promoting a partitioning behavior. We show that our proposal, called Continual Spectral Regularizer (CaSpeR), can be easily combined with any rehearsal-based CL approach and improves the performance of SOTA methods on standard benchmarks. Finally, we conduct additional analysis to provide insights into CaSpeR’s effects and applicability.

## 1. INTRODUCTION

Within the natural world, intelligent creatures continually learn to adapt their behavior to changing external conditions. In doing so, they seamlessly blend novel notions with previous understanding into a cohesive body of knowledge. On the contrary, ANNs greedily fit the data they are currently trained on. For a model that learns on a changing stream of data, this results in the swift deterioration of previously acquired information – a phenomenon known as *catastrophic forgetting* (McCloskey & Cohen, 1989).

Continual Learning (CL) is a branch of machine learning that designs approaches to help deep models retain previ-

<sup>1</sup>ImageLab University of Modena and Reggio Emilia  
<sup>2</sup>Ammagamma <sup>3</sup>Sapienza University of Rome.

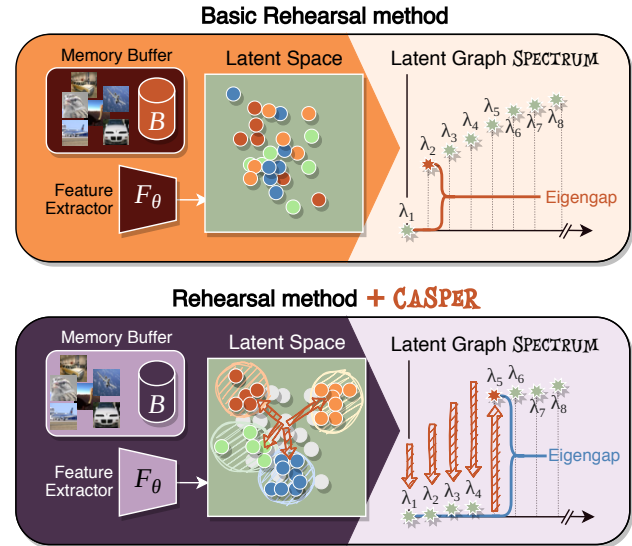


Figure 1. An overview of the proposed CaSpeR regularizer. Rehearsal-based CL methods struggle to separate the latent-space projections of replay data points. Our proposal acts on the spectrum of the latent geometry graph to induce a partitioning behavior by maximizing the *eigengap* for the number of seen classes (*best seen in color*).

ous knowledge while training on new data (De Lange et al., 2021; Parisi et al., 2019). The evaluation of these methods is typically conducted by dividing a classification dataset into disjoint subsets of classes, called *tasks*, letting the model fit one task at a time and, finally, evaluating it on all previously seen classes (van de Ven et al., 2022). Recent literature favors the employment of *rehearsal methods*; namely, CL approaches that address forgetting by retaining a small memory buffer of samples encountered in previous tasks and interleaving them with current training data (Chaudhry et al., 2019; Buzzega et al., 2020a).

On the one hand, rehearsal is a straightforward solution that allows the learner to keep track of the joint distribution of all input classes seen so far. On the other hand, the memory buffer can only accommodate a limited amount of past examples, resulting in overfitting issues (high accuracy on the memory buffer, low accuracy on the test set of the

original tasks). Recent studies characterize this phenomenon in terms of abruptly divergent gradients upon introducing new classes (Caccia et al., 2022; Boschini et al., 2022b) or deteriorating decision surface (Bonicelli et al., 2022). A well-known outcome is the accumulation of a predictive bias in favor of the currently seen classes (Wu et al., 2019; Ahn et al., 2021).

While these works focus their analysis on the overall prediction of the model, we instead consider the changes occurring in its latent space as tasks progress. Specifically, we observe that the learner struggles to separate latent projections of replay examples belonging to different classes. This constitutes a weak spot for the learner, making the downstream classifier prone to interference whenever the input distribution changes and representations are perturbed. Given the Riemannian nature of the latent space of DNNs (Arvanitidis et al., 2018), we naturally revert to spectral geometry to model and constrain its evolution. Spectral geometry is preferred over other geometric tools as it focuses on the latent-space structure without imposing constraints on individual coordinates.

In this work, we introduce a loss term aimed at endowing the model’s latent space with a cohesive structure. Our proposed approach, called **Continual Spectral Regularizer (CaSpeR)**, leverages graph-spectral theory to promote the generation of well-separated latent embeddings, as illustrated Fig. 1. We show that our proposal can be seamlessly combined with any rehearsal-based CL method to improve its classification accuracy and robustness against catastrophic forgetting. Moreover, since CaSpeR does not rely on the availability of annotations for each example, we show that it can be easily applied to semi-supervised scenarios to provide better accuracy and easier convergence. In summary, we make the following contributions:

- We study the interference in rehearsal CL models by investigating the geometry of their latent space. To the best of our knowledge, this is the first attempt at a geometric characterization of catastrophic forgetting;
- We propose Continual Spectral Regularizer: a simple geometrically motivated loss term, inducing the continual learner to produce well-organized latent embeddings;
- We validate our proposal by combining it with several SOTA rehearsal-based CL approaches. Our results show that CaSpeR is effective both in the Class-Incremental and Task-Incremental CL setting (van de Ven et al., 2022) by increasing the geometric consistency of the latent space;
- Finally, we show that CaSpeR can be beneficially applied also to the challenging Continual Semi-

Supervised Learning (CSSL) scenario, producing higher accuracy and easier convergence.

The code to reproduce our experiments is provided in the supplementary material. It will be made available in a public repository upon acceptance of the paper.

## 2. RELATED WORK

### 2.1. Continual Learning

Continual Learning approaches are designed to complement and assist in-training deep learning models to minimize the incidence of *catastrophic forgetting* (McCloskey & Cohen, 1989) when learning on a changing input distribution. This aim can be pursued through different classes of solutions (De Lange et al., 2021): *architectural methods* explicitly allocate separate portions of the model to separate tasks (Mallya & Lazebnik, 2018; Serra et al., 2018); *regularization methods* rely on a loss term to prevent the model from changing either its structure (Kirkpatrick et al., 2017; Ritter et al., 2018) or its response (Li & Hoiem, 2017; Schwarz et al., 2018); *rehearsal methods* derive from the simple Experience Replay (ER) baseline, which exploits a working memory buffer to stash encountered data-points, and later replays them when they are no longer available on the input stream (Robins, 1995; Chaudhry et al., 2019).

Due to their versatility and effectiveness, current research efforts focus primarily on the latter class of approaches (Aljundi et al., 2019). Recent trends highlight interest in improving several aspects of the basic ER formula, *e.g.*, by introducing better-designed memory sampling strategies (Aljundi et al., 2019; Bang et al., 2021), combining replay with other optimization techniques (Lopez-Paz & Ranzato, 2017; Riemer et al., 2019; Chaudhry et al., 2021) or providing richer replay signals (Buzzega et al., 2020a; Ebrahimi et al., 2021).

One of the most prominent challenges for the enhancement of *rehearsal methods* is the imbalance between stream and replay data. Due to the reduced amount and variety of the latter, a continually learned classifier struggles to produce unified predictions and is instead biased towards the last learned classes (Hou et al., 2019; Wu et al., 2019). To counter this effect, researchers have come up with architectural modifications of the model (Hou et al., 2019; Douillard et al., 2020), purposed alterations to the learning objective of the final classifier (Ahn et al., 2021; Caccia et al., 2022) or the outright removal of it, by applying representation learning instead (Cha et al., 2021; Pham et al., 2021).

Our proposal also aims at reducing the intrinsic bias of *rehearsal methods*, but does so by enforcing a desirable property on the latent space of the model. This is achieved through a geometrically motivated regularization term that

can be easily combined with any existing replay method.

## 2.2. Spectral geometry

Our proposal is built upon the eigendecomposition of the Laplace operator on a graph, thus falling within the broader area of spectral graph theory. In particular, ours can be regarded as an *inverse* spectral technique, as we prescribe the general behavior of some eigenvalues and seek a graph whose Laplacian spectrum matches this behavior.

In the geometry processing area, such approaches take the name of *isospectralization* techniques and have been recently used in diverse applications such as deformable 3D shape matching (Cosmo et al., 2019), shape exploration and reconstruction (Marin et al., 2020), shape modeling (Moschella et al., 2022) and adversarial attacks on shapes (Rampini et al., 2021). Differently from these approaches, we work on a single graph (as opposed to pairs of 3D meshes) and our formulation does not take an input spectrum as a target to be matched precisely. Instead, we pose a weaker requirement: the gap between nearby eigenvalues must be maximized, regardless of its exact value. Since our graph represents a discretization of the latent space of a CL model, this simple regularization has important consequences on its learning process.

## 3. METHOD

Our approach exploits tools from spectral geometry to regularize the model’s latent space to hinder forgetting. In Sec. 3.1 we describe the Continual Learning paradigm; in Sec. 3.2 we present a preliminary experiment, highlighting the problem we want to address; finally, in Sec. 3.3 we illustrate our geometric regularizer.

### 3.1. Continual Learning Setting

Following the CL criterion, the model  $F_\theta$  is exposed incrementally to a stream of tasks  $\tau_i$ , where  $i \in \{1, 2, \dots, T\}$ . The parameters  $\theta$  include both the weights of the feature extractor and the classifier,  $\theta^f$  and  $\theta^c$  respectively. Each task consists of a sequence of images and their corresponding labels  $\tau_i = \{(x_1^i, y_1^i), (x_2^i, y_2^i), \dots, (x_n^i, y_n^i)\}$  and does not contain data belonging to classes already seen in previous tasks, so  $Y^i \cap Y^j = \emptyset$ , with  $i \neq j$  and  $Y^i = \{y_k^i\}_{k=1}^n$ . At each step  $i$ , the model cannot freely access data from previous tasks and is optimized by minimizing a loss function over the current set of examples:

$$\theta^{(i)} = \operatorname{argmin}_{\theta} \ell_{\text{stream}} = \operatorname{argmin}_{\theta} \sum_{j=1}^n \ell(F_\theta(x_j^i), y_j^i), \quad (1)$$

where the parameters are initialized with the ones obtained after training on the previous task  $\theta^{(i-1)}$ . If the model does not include mechanisms to prevent forgetting, the accuracy

on all previous tasks will collapse. Rehearsal-based CL methods preserve a portion of examples from previous tasks and store them in a buffer  $B$ , with fixed size  $m$ . This data is then used by the model in conjunction with a specific loss function  $\ell_b$  to hamper catastrophic forgetting:

$$\theta^{(i)} = \operatorname{argmin}_{\theta} \ell_{\text{stream}} + \ell_b. \quad (2)$$

For instance, Experience Replay (ER) simply employs a cross-entropy loss over a batch of examples from  $B$ :

$$\ell_{\text{er}} \triangleq \text{CrossEntropy}(F_\theta(\mathbf{x}^b), \mathbf{y}^b). \quad (3)$$

There exist different strategies for sampling the task data-points to fill the buffer. These will be explained in Sec. 4, along with detail on the  $\ell_b$  employed by each baseline.

### 3.2. Analysis of changing Latent Space Geometry

We are particularly interested in how the latent space changes in response to the introduction of a novel task on the input stream. For this reason, we compute the graph  $\mathcal{G}$  over the latent-space projection of the replay examples gathered by the CL model after training on  $\tau_i$  ( $i \in \{2, \dots, T\}$ )<sup>1</sup>. In order to measure the sparsity of the latent space w.r.t. classes representations, we compute the Label-Signal Variation  $\sigma$  (Lassance et al., 2021) on the adjacency matrix  $A \in \mathbb{R}^{m \times m}$  of  $\mathcal{G}$ :

$$\sigma \triangleq \sum_{i=1}^m \sum_{j=1}^m \mathbb{1}_{y_i^b = y_j^b} a_{i,j}, \quad (4)$$

where  $\mathbb{1}$  is the indicator function. In Fig. 2a, we evaluate several SOTA rehearsal CL methods and show they exhibit a steadily growing  $\sigma$ , which indicates that examples from distinct classes are increasingly entangled in later tasks. This effect can also be observed qualitatively by considering a TSNE embedding of the points in  $B$  (shown in Fig. 2b for XDER), which suggests that the distances between examples from different classes are reduced in later tasks. We remark that both evaluations improve when our proposed regularizer is applied on top of the evaluated methods.

### 3.3. CaSpeR: Continual Spectral Regularizer

**Motivation.** Our method builds upon the fact that the latent spaces of neural models bear a structure informative of the data space they are trained on (Shao et al., 2018). This structure can be enforced through loss regularizers; e.g., in (Cosmo et al., 2020), a minimum-distortion criterion is applied on the latent space of a VAE for a shape generation task. We follow a similar line of thought and propose adopting a geometric term to regularize the latent representations of a CL model.

<sup>1</sup>Please refer to Sec. 3.3 for a detailed description of this procedure.




Figure 2. Illustrations of how CL alters a model’s latent space. (a) A quantitative evaluation measured as Label-Signal Variation ( $\sigma$ ) within the LGG for buffer data points – *lower is better*; (b) TSNE embedding of the features computed by X-DER for buffered examples in later tasks (top). Interference between classes is visibly reduced if CaSpeR is applied (bottom). All experiments are carried out on Split CIFAR-100, (a) uses buffer size 500, (b) uses 2000 (*best seen in colors*).

### Algorithm 1 CaSpeR Loss Computation

**Input** Memory buffer  $B$  of saved samples

- 1:  $\mathbf{x}^b \leftarrow \text{BalancedSampling}(B)$
- 2:  $\mathbf{z}^b \leftarrow F_{\theta_f}(\mathbf{x}^b)$
- 3:  $\mathbf{A} \leftarrow \text{k-NN}(\mathbf{z}^b)$
- 4:  $\mathbf{D} \leftarrow \text{diag}(\sum_i^b a_{1,i}, \sum_i^b a_{2,i}, \dots, \sum_i^b a_{b,i})$
- 5:  $\mathbf{L} \leftarrow \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$  ▷ Eq. 5
- 6:  $\boldsymbol{\lambda} \leftarrow \text{Eigenvalues}(\mathbf{L})$
- 7:  $\ell_{\text{CaSpeR}} \leftarrow -\lambda_{g+1} + \sum_{j=1}^g \lambda_j$  ▷ Eq. 6

**Output**  $\ell_{\text{CaSpeR}}$

Namely, we root our approach in spectral geometry; our choice is motivated by the pursuit of a **compact representation** characterized by **isometry invariance**. As shown in (Arvanitidis et al., 2018), the latent space of DNNs can be modeled as a Riemannian manifold whose *extrinsic* embedding is encoded in the latent vectors. Being extrinsic, these vectors are simply absolute coordinates encoding only one possible realization of the data manifold, out of its infinitely many possible isometries. Each isometry (e.g.: a rotation by  $45^\circ$ ) would always encode the **same latent space**, but the **latent vectors will change** – this is not desirable, because it may lead to overfitting and lack of generalization. By resorting to spectral geometry, we instead rely on *intrinsic* quantities, that fully encode the latent space and are isometry-invariant.

Our regularizer is based on the graph-theoretic formulation of clustering, where we seek to partition the vertices of  $\mathcal{G}$  into well-separated subgraphs with high internal connectivity. A body of results from spectral graph theory, dating back at least to (Cheeger, 1969; Sinclair & Jerrum, 1989; Shi & Malik, 2000), explain the gap occurring between neigh-

boring Laplacian eigenvalues as a quantitative measure of graph partitioning. Our proposal, called Continual Spectral Regularizer (CaSpeR), draws on these results, but turns the *forward* problem of computing the optimal partitioning of a given graph, into the *inverse* problem of seeking a graph with the desired partitioning.

**Building the LGG.** We take the examples in  $B$  and forward them through the network; their features are used to build a k-NN graph  $\mathcal{G}$ ; following (Lassance et al., 2021), we refer to it as the *latent geometry graph* (LGG).

**Spectral Regularizer.** Let us denote by  $\mathbf{A}$  the adjacency matrix of  $\mathcal{G}$ , we calculate its degree matrix  $\mathbf{D}$  and we compute its normalized Laplacian as:

$$\mathbf{L} = \mathbf{I} - \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}, \quad (5)$$

where  $\mathbf{I}$  is the identity matrix. Finally, we compute the eigenvalues  $\boldsymbol{\lambda}$  of  $\mathbf{L}$  and sort them by ascending order. Let  $g$  be the number of different classes within the buffer, we calculate our regularizing loss as:

$$\ell_{\text{CaSpeR}} \triangleq -\lambda_{g+1} + \sum_{j=1}^g \lambda_j. \quad (6)$$

The proposed loss term is weighted through the hyperparameter  $\rho$  and added to the stream classification loss. Overall, our model optimizes the following objective:

$$\arg \min_{\theta} \ell_{\text{stream}} + \ell_b + \rho \ell_{\text{CaSpeR}}. \quad (7)$$

Through Eq. 6, we increase the eigengap  $\lambda_{g+1} - \lambda_g$  while minimizing the first  $g$  eigenvalues – the intuition being that the number of eigenvalues close to zero corresponds to the number of loosely connected partitions within the

graph (Lee et al., 2014). Therefore, our loss indirectly encourages the points in the buffer to be clustered without strict supervision. We refer the reader to Algorithm 1 for a step-by-step summary of the outlined procedure.

**Efficient Batch Operation.** While seemingly straightforward, the operation of CaSpeR entails the cumbersome task of constructing the entire LGG  $\mathcal{G}$  at each forward step. Indeed, accurately mapping the model’s ever-changing latent space requires processing all available replay examples in the buffer  $B$ , which is typically orders of magnitude larger than a batch of examples on the input stream.

To avoid a slow training procedure with high memory requirements, we propose an efficient approximation of our initial objective. Instead of operating on  $\mathcal{G}$  directly, we sample a randomly chosen sub-graph  $\mathcal{G}_p \subset \mathcal{G}$  spanning only  $p$  out of the  $g$  classes represented in the memory buffer. As  $\mathcal{G}_p$  still includes a conspicuous amount of nodes, we resort to an additional sub-sampling and extract  $\mathcal{G}_p^t \subset \mathcal{G}_p$ , a smaller graph with  $t$  exemplars for each class.

By repeating these random samplings in each forward step, we optimize a Monte Carlo approximation of Eq. 6:

$$\ell_{\text{CaSpeR}}^* \triangleq \mathbb{E}_{\mathcal{G}_p \subset \mathcal{G}} \left[ \mathbb{E}_{\mathcal{G}_p^t \subset \mathcal{G}_p} \left[ -\lambda_{p+1}^{\mathcal{G}_p^t} + \sum_{j=1}^p \lambda_j^{\mathcal{G}_p^t} \right] \right], \quad (8)$$

where the  $\lambda_p^{\mathcal{G}_p^t}$  denote the eigenvalues of the Laplacian of  $\mathcal{G}_p^t$ . It must be noted that enforce the eigengap at  $p$ , as we know by construction that each  $\mathcal{G}_p^t$  comprises samples from  $p$  communities within  $\mathcal{G}$ .

## 4. EVALUATION

### 4.1. Evaluation protocol

**Settings.** To assess the effectiveness of the proposed method, we consider both split incremental classification protocols formalized in (van de Ven et al., 2022): *Task-Incremental Learning* (Task-IL), where the task information is given during training and evaluation; and *Class Incremental Learning* (Class-IL), where the model learns to make predictions in the absence of task information. On the one hand, Class-IL is recognized as a more realistic and challenging benchmark (Farquhar & Gal, 2018; Aljundi et al., 2019); on the other, Task-IL is especially relevant for the quantification of forgetting, as it is unaffected by data imbalance biases (Wu et al., 2019; Boschini et al., 2022a).

**Benchmarked models.** To evaluate the benefit of our regularizer, we apply it to the following state-of-the-art rehearsal-based methods:

- **Experience Replay with Asymmetric Cross-Entropy (ER-ACE)** (Caccia et al., 2022): starting

from classic Experience Replay, the authors obtain a significant performance gain by freezing the previous task heads of the classifier while computing the loss on the streaming data;

- **Incremental Classifier and Representation Learning (iCaRL)** (Rebuffi et al., 2017): this method seeks to learn the best representation of data that fits a nearest-neighbor classifier w.r.t. class prototypes stored in the buffer;
- **Dark Experience Replay (DER++)** (Buzzega et al., 2020a): another variant of ER, which combines the standard classification replay with a distillation loss;
- **eXtended-DER (X-DER)** (Boschini et al., 2022a): a method which improves DER++ by addressing its shortcomings and focusing on organically accommodating future knowledge<sup>2</sup>;
- **Pooled Outputs Distillation Network (POD-Net)** (Douillard et al., 2020): the authors extend iCaRL’s classification method: their model learns multiple representations for each class and adopts two additional distillation losses.

These approaches adopt different strategies for the construction of their memory buffer: X-DER, iCaRL and PODNet use a class-balanced offline sampling strategy; ER-ACE and DER++ use *reservoir* sampling (Vitter, 1985), which might lead to uneven class representation within the stored examples. Since CaSpeR relies on the availability of a minimum amount of samples per class, we adjust the latter sampling strategy to enforce equity, as done in (Buzzega et al., 2020b).

To have a better understanding of the results, we include the performance of the upper bound (Joint), obtained by training on all classes together in a standard offline manner, and the lower bound (Finetune) obtained by training on each task sequentially without any method to prevent forgetting.

**Datasets.** We conduct all the experiments on two commonly used image datasets, splitting the classes from the main dataset into separate disjoint sets used to sequentially train the evaluated models.

- **Split CIFAR-100:** CIFAR100 (Krizhevsky et al., 2009) contains 100 classes with 500 images per class, where each image has a dimension of  $32 \times 32$ . We split the dataset into 10 subsets of 10 classes each;
- **Split miniImageNet:** miniImageNet (Vinyals et al., 2016) is a subset of the ImageNet dataset where each image is resized to  $84 \times 84$ . We use the 20 tasks per 5 classes protocol.

<sup>2</sup>Specifically, we use the more effective baseline based on a Regular Polytope Classifier (Pernici et al., 2021)

Table 1. Class-IL results –  $\bar{A}_F$  ( $\bar{F}_F^*$ ) – for SOTA rehearsal CL methods, with and without CaSpeR.

Class-IL	Split CIFAR-100		Split <i>miniImageNet</i>	
Joint (UB)	63.11 (–)		52.76 (–)	
Finetune (LB)	8.38 (100.00)		3.87 (100.00)	
<b>Buffer Size</b>	500	2000	2000	5000
ER-ACE	34.99 (51.41)	46.63 (28.78)	22.03 (49.04)	27.26 (29.99)
+ CaSpeR	36.70+1.71 (46.61)	47.85+1.22 (27.73)	23.36+1.33 (47.90)	29.15+1.89 (28.36)
iCaRL	39.80 (32.73)	40.54 (32.61)	19.42 (36.89)	20.17 (33.23)
+ CaSpeR	40.57+0.77 (32.31)	41.83+1.29 (25.55)	20.46+1.04 (35.90)	21.45+1.28 (32.26)
DER++	28.01 (57.56)	42.27 (34.94)	20.88 (74.48)	28.55 (61.03)
+ CaSpeR	32.16+4.15 (53.41)	46.34+4.07 (30.08)	22.61+1.73 (71.01)	29.72+1.17 (57.60)
X-DER	35.89 (44.54)	46.37 (23.57)	24.80 (44.69)	30.98 (30.12)
+ CaSpeR	38.23+2.34 (43.90)	50.39+4.02 (17.65)	26.24+1.44 (41.72)	31.55+0.57 (28.71)
PODNet	28.16 (58.49)	32.12 (46.73)	16.82 (52.32)	20.81 (46.50)
+ CaSpeR	31.40+3.24 (48.50)	36.97+4.85 (39.00)	18.09+1.27 (50.33)	22.45+1.64 (46.08)

**Metrics.** We mainly quantify the performance of the compared models in terms of *Final Average Accuracy* ( $\bar{A}_F$ ), i.e., the average classification accuracy of the model at the end of the overall training process:

$$\bar{A}_F \triangleq \frac{1}{T} \sum_{i=1}^T a_i^T, \quad (9)$$

where  $a_i^j$  is the accuracy of the model at the end of task  $j$  calculated on the test set of task  $\tau_i$  and reported in percentage value. To quantify the severity of the performance degradation that occurs as a result of catastrophic forgetting, we propose a novel measure called *Final Average Adjusted Forgetting* ( $\bar{F}_F^*$ ), which we define as follows:

$$\bar{F}_F^* = \frac{1}{T-1} \sum_{i=1}^{T-1} \left[ \frac{a_i^* - a_i^T}{a_i^*} \right]^+, \quad (10)$$

where  $a_i^* = \max_{t \in \{i, \dots, T-1\}} a_i^t, \forall i \in \{1, \dots, T-1\}$ .

$\bar{F}_F^*$  is bounded in  $[0, 100]$ , where the upper bound is given by a method that retains no accuracy on previous tasks (as is the case for the Finetune baseline). This measure derives from the widely employed Forgetting metric (Chaudhry et al., 2018), which tends to be more forgiving of those methods that do not properly learn the current task (Buzzega et al., 2020a).

**Hyperparameter selection.** To ensure a fair evaluation, we train all the models with the same batch size and the same number of epochs. Moreover, we employ the same backbone for all experiments on the same dataset. In particular, we use Resnet18 (He et al., 2016) for Split CIFAR-100 and EfficientNet-B2 (Tan & Le, 2019) for Split *miniImageNet*. The best hyperparameters for each model-dataset configuration are found via grid search.

We refer the reader to the Appendix for additional details.

## 4.2. Experimental results

We report a breakdown of the results of our evaluation in Tab. 1 (Class-IL) and 2 (Task-IL). At first glance, CaSpeR leads to a steady improvement in  $\bar{A}_F$  across all evaluated methods and settings. However, some interesting additional trends emerge upon closer examination.

Firstly, we notice that the improvement in accuracy does not always grow with the memory buffer size. This is in contrast with the typical behavior of replay regularization terms (Cha et al., 2021; Chaudhry et al., 2019). We believe such a tendency to be the result of our distinctively geometric approach: as spectral properties of graphs are understood to be robust w.r.t. to coarsening (Jin et al., 2020), CaSpeR does not need a large pool of data to be effective.

On Split Cifar-100, CaSpeR is able to improve the performance of existing methods solidly. In Task-IL, the gain is lower than in Class-IL: existing methods are already strong in evaluating each task individually (Task-IL), so the benefits of our regularizer are more evident when all tasks are mixed (Class-IL). The methods that most capitalize on CaSpeR feature space are DER++, PODNet and X-Der.

On Split *miniImageNet*, we see a steady but reduced upgrade of  $\bar{A}_F$  over the baselines. This suggests that our approach struggles when dealing with the protracted amount of tasks in this dataset. However,  $\bar{F}_F^*$  results in both CL settings attest to CaSpeR’s ability to yield the consistent improvements.

As the reader can notice, iCaRL appears to be an outlier: the improvement of our regularization is minor in every scenario. We motivate this effect to the intrinsic nature of iCaRL: by leveraging a *nearest-mean-of-exemplars* classifier, its feature space is already effectively separating points of different classes prior to the application of CaSpeR.

Table 2. Task-IL results –  $\bar{A}_F$  ( $\bar{F}_F^*$ ) – for SOTA rehearsal CL methods, with and without CaSpeR.

Task-IL	Split CIFAR-100		Split miniImageNet	
Joint (UB)	88.81 (–)		87.39 (–)	
Finetune (LB)	30.10 (62.84)		24.05 (67.37)	
Buffer Size	500	2000	2000	5000
ER-ACE	73.86 (10.73)	80.69 (5.37)	69.05 (13.72)	72.78 (8.93)
+ CaSpeR	75.14+1.28 (4.91)	81.57+0.88 (4.93)	69.59+0.54 (13.05)	74.14+1.36 (8.12)
iCaRL	78.38 (5.38)	78.47 (4.91)	70.35 (3.92)	70.44 (2.68)
+ CaSpeR	79.31+0.93 (4.61)	79.43+0.96 (3.41)	71.19+0.84 (3.67)	71.93+1.49 (3.65)
DER++	70.55 (11.12)	78.60 (5.96)	69.78 (13.37)	73.81 (8.59)
+ CaSpeR	73.25+2.70 (9.49)	80.78+2.18 (3.04)	70.97+1.19 (11.75)	75.18+1.37 (7.93)
X-DER	77.28 (2.43)	82.55 (0.92)	74.32 (4.95)	77.70 (3.71)
+ CaSpeR	78.26+0.98 (5.47)	83.77+1.22 (0.27)	75.99+1.67 (3.88)	78.71+1.01 (2.32)
PODNet	67.37 (19.76)	69.63 (15.16)	60.60 (14.00)	66.15 (10.71)
+ CaSpeR	70.81+3.44 (15.26)	71.90+2.27 (11.32)	64.84+4.24 (10.01)	70.85+4.70 (7.99)

 Table 3. Class-IL  $\bar{A}_F$  values of k-NN classifiers trained on top of the latent representations of replay data points. Results on Split CIFAR-100 for Buffer Size 2000.

k-NN Clsf (Class-IL)	w/o CaSpeR		w/ CaSpeR	
	5-NN	11-NN	5-NN	11-NN
ER-ACE	43.73	44.41	46.75+3.02	47.29+2.88
iCaRL	34.86	37.78	36.00+1.14	38.33+0.55
DER++	44.21	44.24	45.75+1.54	46.00+1.76
X-DER	43.44	44.62	49.47+6.03	49.49+4.87
PODNet	21.11	22.60	27.88+6.77	28.94+6.34

## 5. MODEL ANALYSIS

### 5.1. k-NN classification

To further verify whether CaSpeR successfully separates the latent embeddings for examples of different classes, we evaluate the accuracy of k-NN-classifiers (Wu et al., 2018) trained on top of the latent representations produced by the methods of Sec. 4. In Tab. 3, we report the results for 5-NN and 11-NN classifiers using the final buffer  $B$  as a support set. We observe that CaSpeR also shows its steady beneficial effect on top of this classification approach, further confirming that it is instrumental in disentangling the representations of different classes.

### 5.2. Latent Space Consistency

To provide further insights into the dynamics of the latent space on the evaluated models, we study the emergence of distortions in the LGG. Given a continual learning model, we are interested in a comparison between  $\mathcal{G}_5$  and  $\mathcal{G}_{10}$ , the LGGs produced after training on  $\tau_5$  and  $\tau_{10}$  respectively, computed on the test set of tasks  $\tau_1, \dots, \tau_5$ .

The comparison between  $\mathcal{G}_5$  and  $\mathcal{G}_{10}$  can be better understood in terms of the node-to-node bijection  $T : \mathcal{G}_5 \rightarrow \mathcal{G}_{10}$ , which can be represented as a functional map matrix  $C$  (Ovsjanikov et al., 2012) with elements

$$c_{i,j} \triangleq \langle \phi_i^{\mathcal{G}_5}, \phi_j^{\mathcal{G}_{10}} \circ T \rangle, \quad (11)$$

where  $\phi_i^{\mathcal{G}_5}$  is the  $i$ -th Laplacian eigenvector of  $\mathcal{G}_5$  (similarly for  $\mathcal{G}_{10}$ ), and  $\circ$  denotes the standard function composition. In other words, the matrix  $C$  encodes the similarity between the Laplacian eigenspaces of the two graphs. In an ideal scenario where the latent space is subject to no modification between  $\tau_5$  and  $\tau_{10}$  w.r.t. previously learned classes,  $T$  is an *isomorphism* and  $C$  is a diagonal matrix (Ovsjanikov et al., 2012). In a practical scenario,  $T$  is only approximately isomorphic and, the better the approximation, the more  $C$  is sparse and funnel-shaped.

In Fig. 3, we report  $C^{|\cdot|} \triangleq \text{abs}(C)$  for ER-ACE, DER++, iCaRL and X-DER on Split CIFAR-100, both with and without CaSpeR. It can be observed that the methods that benefit the most from our proposal (ER-ACE, X-DER) display a tighter functional map matrix. This indicates that the partitioning behavior promoted by CaSpeR leads to reduced interference, as the portion of the LGG that refers to previously learned classes remains geometrically consistent in later tasks. On the other hand, in line with the considerations made in Sec. 4.2, the improvement is only marginal for iCaRL. Its different training regime, which is less discriminative in nature, seemingly induces a limited amount of change on the structure of the latent space.

To quantify the similarity of each  $C^{|\cdot|}$  matrix to the identity, we also report its off-diagonal energy, computed as follows (Rodolà et al., 2017):

$$ODE \triangleq \frac{1}{\|C\|_F^2} \sum_i \sum_{j \neq i} c_{i,j}^2, \quad (12)$$




Figure 3. For several rehearsal methods with and without CaSpeR, the functional map magnitude matrices  $C^{l-1}$  between the LGGs  $\mathcal{G}_5$  and  $\mathcal{G}_{10}$ , computed on the test set of  $\tau_1, \dots, \tau_5$  after training up to  $\tau_5$  and  $\tau_{10}$  respectively (Split CIFAR-100 - buffer size 2000). The closer  $C^{l-1}$  to the diagonal, the less geometric distortion between  $\mathcal{G}_5$  and  $\mathcal{G}_{10}$ . We report the first 25 rows and columns of  $C^{l-1}$ , focusing on smooth (low-frequency) correspondences (Ovsjanikov et al., 2012), and apply a  $C^{l-1} > 0.15$  threshold to increase clarity.

where  $\|\cdot\|_F$  indicates the Frobenius norm. CaSpeR produces a clear decrease in  $OD_E$ , signifying an increase in the diagonality of the functional matrices.

### 5.3. Continual Semi-supervised Learning

In Sec. 4.2, we shed light on some interesting properties of CaSpeR, *i.e.*, its ability to operate well in a low-data regime and its role in facilitating the convergence of underperforming baselines. Both issues naturally emerge in the Continual Semi-Supervised Learning (CSSL) setting (Boschini et al., 2022b), a recently-proposed CL experimental benchmark, where only a fraction of the examples on the input stream are associated with an annotation.

In a supervised CL setting, we apply CaSpeR to buffer data points, thus encouraging the separation of all previously encountered classes in the latent space. However, we remark that our proposed approach does not have strict supervision requirements, as it does not need the labels attached to each node in the LGG, but rather just the total amount of classes  $g$  that must be clustered (Eq. 6).

In Tab. 4, we report the results of an experiment on Split CIFAR-100 in the CSSL setting with only 0.8% or 5% annotated labels. Typical CL methods operating in this scenario are forced to discard a consistent amount of data (ER-ACE), leading to majorly reduced performance w.r.t. the fully-supervised case, or to use the in-training model to annotate unlabeled samples (*pseudo-labeling*, PsER-ACE), but might backfire if the provided supervision does not suffice for the learner to produce reliable responses (as is the case with 0.8% labels).

To allow for the exploitation of unlabeled exemplars, we also apply CaSpeR on data points from the input stream, by taking  $k$  equal to the number of classes in a given task. We show that this leads to an overall improvement of the tested models and – particularly – counteracts the failure case where PseudoER-ACE is applied on top of a few annotated data.

Table 4. Class-IL  $\bar{A}_F$  values on Split CIFAR-100, with reduced amount of annotations (CSSL). Buffer size 2000.  $\dagger$  indicates results taken from (Boschini et al., 2022b).

CSSL	w/o CaSpeR		w/ CaSpeR	
	Labels %	0.8%	5%	0.8%
ER-ACE	8.46	11.87	8.55+0.09	14.16+2.29
PsER-ACE	2.31	16.35	9.69+7.38	17.42+1.07
CCIC	11.5 $\dagger$	19.5 $\dagger$	12.22+0.72	20.32+0.82

This indicates that CaSpeR manages to limit the impact of the noisy labels produced by *pseudo-labeling*. Finally, we show that CaSpeR can be easily applied to CCIC (Boschini et al., 2022b) – a CSSL method that leverages both labeled and unlabeled data – to improve its  $\bar{A}_F$ .

## 6. CONCLUSION

In this work, we investigate how the latent space of a CL model changes throughout training. We find that latent-space projections of past exemplars are relentlessly drawn closer together, possibly interfering and paving the way for catastrophic forgetting.

Drawing on spectral graph theory, we propose Continual Spectral Regularizer (CaSpeR): a regularizer that encourages the clustering of data points in the latent space. We show that our approach can be easily combined with any rehearsal-based CL approach, improving the performance of SOTA methods on standard benchmarks.

Furthermore, we analyze the effects of CaSpeR showing that the regularized latent space correctly separates examples from different classes and is subject to fewer distortions. Finally, we verify that our proposed approach is also applicable with partial supervision, improving the accuracy of Continual Semi-Supervised Learning baselines and facilitating their convergence in a low-label regime.



## References

- Ahn, H., Kwak, J., Lim, S., Bang, H., Kim, H., and Moon, T. SS-IL: Separated Softmax for Incremental Learning. In *IEEE International Conference on Computer Vision*, 2021.
- Aljundi, R., Lin, M., Goujaud, B., and Bengio, Y. Gradient Based Sample Selection for Online Continual Learning. In *Advances in Neural Information Processing Systems*, 2019.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. In *International Conference on Learning Representations Workshop*, 2018.
- Bang, J., Kim, H., Yoo, Y., Ha, J.-W., and Choi, J. Rainbow memory: Continual learning with a memory of diverse samples. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2021.
- Bonicelli, L., Boschini, M., Porrello, A., Spampinato, C., and Calderara, S. On the Effectiveness of Lipschitz-Driven Rehearsal in Continual Learning. In *Advances in Neural Information Processing Systems*, 2022.
- Boschini, M., Bonicelli, L., Buzzega, P., Porrello, A., and Calderara, S. Class-incremental continual learning into the extended der-verse. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022a.
- Boschini, M., Buzzega, P., Bonicelli, L., Porrello, A., and Calderara, S. Continual semi-supervised learning through contrastive interpolation consistency. *Pattern Recognition Letters*, 2022b.
- Buzzega, P., Boschini, M., Porrello, A., Abati, D., and Calderara, S. Dark Experience for General Continual Learning: a Strong, Simple Baseline. In *Advances in Neural Information Processing Systems*, 2020a.
- Buzzega, P., Boschini, M., Porrello, A., and Calderara, S. Rethinking Experience Replay: a Bag of Tricks for Continual Learning. In *International Conference on Pattern Recognition*, 2020b.
- 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 *International Conference on Learning Representations Workshop*, 2022.
- Cha, H., Lee, J., and Shin, J. Co2l: Contrastive continual learning. In *IEEE International Conference on Computer Vision*, 2021.
- Chaudhry, A., Dokania, P. K., Ajanthan, T., and Torr, P. H. Riemannian walk for incremental learning: Understanding forgetting and intransigence. In *Proceedings of the European Conference on Computer Vision*, 2018.
- Chaudhry, A., Rohrbach, M., Elhoseiny, M., Ajanthan, T., Dokania, P. K., Torr, P. H., and Ranzato, M. On tiny episodic memories in continual learning. In *International Conference on Machine Learning Workshop*, 2019.
- Chaudhry, A., Gordo, A., Dokania, P., Torr, P., and Lopez-Paz, D. Using hindsight to anchor past knowledge in continual learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2021.
- Cheeger, J. A lower bound for the smallest eigenvalue of the laplacian. In *Problems in analysis*. Princeton University Press, 1969.
- Cosmo, L., Panine, M., Rampini, A., Ovsjanikov, M., Bronstein, M. M., and Rodolà, E. Isospectralization, or how to hear shape, style, and correspondence. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019.
- Cosmo, L., Norelli, A., Halimi, O., Kimmel, R., and Rodolà, E. Limp: Learning latent shape representations with metric preservation priors. In *Proceedings of the European Conference on Computer Vision*, 2020.
- De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., and Tuytelaars, T. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
- Douillard, A., Cord, M., Ollion, C., Robert, T., and Valle, E. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Proceedings of the European Conference on Computer Vision*, 2020.
- Ebrahimi, S., Petryk, S., Gokul, A., Gan, W., Gonzalez, J. E., Rohrbach, M., and Darrell, T. Remembering for the right reasons: Explanations reduce catastrophic forgetting. *Applied AI Letters*, 2021.
- Farquhar, S. and Gal, Y. Towards Robust Evaluations of Continual Learning. In *International Conference on Machine Learning Workshop*, 2018.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2016.
- Hou, S., Pan, X., Loy, C. C., Wang, Z., and Lin, D. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019.

- Jin, Y., Loukas, A., and JaJa, J. Graph coarsening with preserved spectral properties. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- 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*, 2017.
- Krizhevsky, A. et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- Lassance, C., Gripon, V., and Ortega, A. Representing deep neural networks latent space geometries with graphs. *MDPI Algorithms*, 2021.
- Lee, J. R., Gharan, S. O., and Trevisan, L. Multiway spectral partitioning and higher-order cheeger inequalities. *Journal of the ACM*, 2014.
- Li, Z. and Hoiem, D. Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.
- Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. In *Advances in Neural Information Processing Systems*, 2017.
- Mallya, A. and Lazebnik, S. Packnet: Adding multiple tasks to a single network by iterative pruning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.
- Marin, R., Rampini, A., Castellani, U., Rodolà, E., Ovsjanikov, M., and Melzi, S. Instant recovery of shape from spectrum via latent space connections. In Struc, V. and Fernández, F. G. (eds.), *International Conference on 3D Vision*, 2020.
- McCloskey, M. and Cohen, N. J. Catastrophic interference in connectionist networks: The sequential learning problem. *Psychology of learning and motivation*, 1989.
- Moschella, L., Melzi, S., Cosmo, L., Maggioli, F., Litany, O., Ovsjanikov, M., Guibas, L. J., and Rodolà, E. Learning spectral unions of partial deformable 3d shapes. *Computer Graphics Forum*, 2022.
- Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., and Guibas, L. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 2012.
- Parisi, G. I., Kemker, R., Part, J. L., Kanan, C., and Wermter, S. Continual lifelong learning with neural networks: A review. *Neural Networks*, 2019.
- Pernici, F., Bruni, M., Baecchi, C., Turchini, F., and Del Bimbo, A. Class-incremental learning with pre-allocated fixed classifiers. In *International Conference on Pattern Recognition*, 2021.
- Pham, Q., Liu, C., and Hoi, S. Dualnet: Continual learning, fast and slow. In *Advances in Neural Information Processing Systems*, 2021.
- Rampini, A., Pestarini, F., Cosmo, L., Melzi, S., and Rodolà, E. Universal spectral adversarial attacks for deformable shapes. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2021.
- Rebuffi, S.-A., Kolesnikov, A., Sperl, G., and Lampert, C. H. iCaRL: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2017.
- Riemer, M., Cases, I., Ajemian, R., Liu, M., Rish, I., Tu, Y., and Tesauro, G. Learning to Learn without Forgetting by Maximizing Transfer and Minimizing Interference. In *International Conference on Learning Representations Workshop*, 2019.
- Ritter, H., Botev, A., and Barber, D. Online structured laplace approximations for overcoming catastrophic forgetting. *Advances in Neural Information Processing Systems*, 2018.
- Robins, A. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 1995.
- Rodolà, E., Cosmo, L., Bronstein, M. M., Torsello, A., and Cremers, D. Partial functional correspondence. In *Computer Graphics Forum*, 2017.
- Schwarz, J., Czarnecki, W., Luketina, J., Grabska-Barwinska, A., Teh, Y. W., Pascanu, R., and Hadsell, R. Progress & compress: A scalable framework for continual learning. In *International Conference on Machine Learning*, 2018.
- Serra, J., Suris, D., Miron, M., and Karatzoglou, A. Overcoming Catastrophic Forgetting with Hard Attention to the Task. In *International Conference on Machine Learning*, 2018.
- Shao, H., Kumar, A., and Fletcher, P. T. The riemannian geometry of deep generative models. In *IEEE International Conference on Computer Vision and Pattern Recognition Workshops*, 2018.
- Shi, J. and Malik, J. Normalized cuts and image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2000.

- Sinclair, A. and Jerrum, M. Approximate counting, uniform generation and rapidly mixing markov chains. *Information and Computation*, 1989.
- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, 2019.
- van de Ven, G. M., Tuytelaars, T., and Tolias, A. S. Three types of incremental learning. *Nature Machine Intelligence*, 2022.
- Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al. Matching networks for one shot learning. In *Advances in Neural Information Processing Systems*, 2016.
- Vitter, J. S. Random sampling with a reservoir. *ACM Transactions on Mathematical Software*, 1985.
- Wu, Y., Chen, Y., Wang, L., Ye, Y., Liu, Z., Guo, Y., and Fu, Y. Large scale incremental learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2019.
- Wu, Z., Xiong, Y., Yu, S. X., and Lin, D. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2018.