



PODNet: Pooled Outputs Distillation for Small-Tasks Incremental Learning

Arthur Douillard^{1,2(✉)}, Matthieu Cord^{2,3}, Charles Ollion¹, Thomas Robert¹,
and Eduardo Valle⁴

¹ Heuritech, Paris, France

{`arthur.douillard, charles.ollion, thomas.robert`}@heuritech.com

² Sorbonne University, Paris, France

`matthieu.cord@sorbonne-universite.fr`

³ Valeo.ai, Paris, France

⁴ University of Campinas, Campinas, Brazil

`dovalle@dca.fee.unicamp.br`

Abstract. Lifelong learning has attracted much attention, but existing works still struggle to fight catastrophic forgetting and accumulate knowledge over long stretches of incremental learning. In this work, we propose PODNet, a model inspired by representation learning. By carefully balancing the compromise between remembering the old classes and learning new ones, PODNet fights catastrophic forgetting, even over very long runs of small incremental tasks – a setting so far unexplored by current works. PODNet innovates on existing art with an efficient spatial-based distillation-loss applied throughout the model and a representation comprising multiple proxy vectors for each class. We validate those innovations thoroughly, comparing PODNet with three state-of-the-art models on three datasets: CIFAR100, ImageNet100, and ImageNet1000. Our results showcase a significant advantage of PODNet over existing art, with accuracy gains of 12.10, 6.51, and 2.85 percentage points, respectively.

Keywords: Incremental-learning · Representation-learning pooling

1 Introduction

Lifelong machine learning [7, 31, 34] focuses on models that accumulate and refine knowledge over large timespans. Incremental learning – the ability to aggregate different learning objectives seen over time into a coherent whole – is paramount to those models. To achieve incremental learning, models must fight *catastrophic forgetting* [7, 31] of previous knowledge. Lifelong and incremental learning have attracted much attention in the past few years, but existing works still struggle

Electronic supplementary material The online version of this chapter (https://doi.org/10.1007/978-3-030-58565-5_6) contains supplementary material, which is available to authorized users.

to preserve acquired knowledge over many cycles of short incremental learning steps¹.

We will focus on image classifiers, which are ordinarily trained once on a fixed set of classes. In *incremental learning*, however, the classifier must learn the classes by steps, in training cycles called *tasks*. At each task, we expose the classifier to a new set of classes. Incremental learning would reduce trivially to ordinary classification if we were allowed to store all training samples, but we are imposed a limited *memory*: a maximum number of samples for previously learned classes. This limitation is motivated by practical applications, in which privacy issues, or storage and computing limitations prevent us from simply retraining the entire model for each new task [21, 22]. Furthermore, incremental learning is different from transfer learning in that we aim to have good performance in both old and new classes.

To overcome catastrophic forgetting, different approaches have been proposed: reusing a limited amount of previous training data [3, 30]; learning to generate the training data [15, 33]; extending the architecture for new phases of data [20, 36]; using a sub-network for each phase [6, 10]; or constraining the model divergence as it evolves [1, 3, 16, 21, 23, 30].

In this work, we propose PODNet, approaching incremental learning as representation learning, with a distillation loss that constrains the evolution of the representation. By carefully balancing the compromise between remembering the old classes and learning new ones, we learn a representation that fights catastrophic forgetting, remaining stable over long runs of small incremental tasks. Our model innovates on existing art with (1) an *efficient spatial-based* distillation-loss applied *throughout the model*; and (2) as a refinement, a representation comprising multiple proxy vectors for each class, resulting in a more flexible representation.

In this paper, we first present the existing state of the art (Sect. 2), which we close by detailing our contributions. We then describe our model (Sect. 3), and evaluate it in an extensive set of experiments (Sect. 4) on CIFAR100, ImageNet100, and ImageNet1000, including ablation studies assessing each contribution, and extensive comparisons with existing methods.

2 Related Work

To approach the problem of incremental learning, consider a single incremental task: one has a classifier already trained over a set of old classes and must adapt it to learn a set of new classes. To perform that single task, we will consider: (1) the data/class representation model; (2) the set of constraints to prevent catastrophic forgetting; (3) the experimental context (including the constraints over the memory for previous training data) for which to design the model.

¹ Code is available at: github.com/arthurdouillard/incremental_learning_pytorch.

Data/Class Representation Model. Representation learning was already implicitly present in iCaRL [30]: it introduced the Nearest Mean Exemplars (NME) strategy which averages the outputs of the deep convolutional network to create a single proxy feature vector per class that are then used by a nearest-neighbor classifier predict the final classes. Hou et al. [13] adopted this method and also introduced another, named CNN, which uses the output class probabilities to classify incoming samples, freezing (during training) the classifier weights associated with old classes, and then fine-tuning them on an under-sampled dataset.

Hou et al. [13], in the method called here UCIR, made representation learning explicit, by noticing that the limited memory imposed a severe imbalance on the training samples available for the old and for the new classes. To overcome that difficulty, they designed a metric-learning model instead of a classification model. That strategy is often used in few-shot learning [8] because of its robustness to few data. Because classical metric architectures require special training sampling (e.g., semi-hard sampling for triplets), Hou et al. chose instead to redesign the classifier’s last layer of their model to use the cosine similarity [25].

Model Constraints to Prevent Catastrophic Forgetting. Constraining the model’s evolution to prevent forgetting is a fruitful idea proposed by several methods [1, 3, 16, 21, 23, 30]. Preventing the model’s parameters from diverging too much forces it to remember the old classes, but care must be taken to still allow it to learn the new ones. We call this balance the *rigidity-plasticity trade-off*.

Existing art on knowledge distillation/compression [12] was an important source of inspiration for constraints on models. The goal is to distill a large trained model (called teacher) into a new smaller model (called student). The distillation loss forces the features of the student to approach those of its teacher. In our case, the student is the current model and the teacher—with same capacity—is its version at the previous task. Zagoruyko and Komodakis [17] investigated attention-based distillation for image classifiers, by pooling the intermediate features of convolutional networks into attention maps, then used in their distillation losses. Li and Hoiem [21]—and several authors after them [3, 30, 35]—used a binary cross-entropy between the output probabilities by the models. Hou et al. [13], used instead *Less-Forget*, a cosine-similarity constraint on the flat feature embeddings after the global average pooling. Dhar et al. [5] proposed to constrain the gradient-based attentions generated by GradCam [32], a visualization method. Wu et al. [35] proposed BiC, an algorithm oriented towards large-scale datasets, which employs a small linear model learned on validation data to recalibrate the output probabilities before applying a distillation loss.

Experimental Context. A critical component of incremental learning is the convention used for the memory storing samples of previous data. An usual convention is to consider a fixed amount of samples allowed in that memory, as illustrated in Fig. 1.




Fig. 1. Training protocol for incremental learning. At each training task we learn a new set of classes, and the model must retain knowledge about *all* classes. The model is allowed a *limited* memory of samples of old classes.

Still, there are two experimental protocols for such fixed-sample convention: we may either use the memory budget at will (M_{total}), or add a constraint on the number of samples per class for the old classes (M_{per}). When $M_{\text{total}} = M_{\text{per}} \times \#$ of classes, both settings have equivalent *final* memory size, but the latter, that we adopt, is much more challenging since early tasks cannot benefit from the full memory size. *The granularity of the increments* is another critical element: with a fixed number of classes, increasing the number of tasks decreases the number of classes per task. More tasks imply stronger forgetting of the earliest classes, and pushing that number creates a challenging protocol, so far unexplored by existing art. Hou et al. evaluate at most 10 tasks on CIFAR100, while we propose as much as 50 tasks.

Finally, to score the experiments, Rebuffi et al. [30] proposed a global metric that they called **average incremental accuracy**, taking into account the entire history of the run, averaging the accuracy at the end of each task (including the first).

Contributions. As seen, associating representation learning to model constraints is a particularly fruitful idea for incremental learning, but requires carefully balancing the goals of rigidity (to avoid catastrophic forgetting) and plasticity (to learn new classes).

Employing a distillation-based loss to constrain the evolution of the representation has also resulted in leading results [5, 13, 35, 37]. Our model improves existing art by employing a *novel and efficient spatial-based* distillation loss, which we are able to apply *throughout the model*.

Implicit or explicit proxy vectors representing each class inside the models have lead to state of the art results [13, 30]. Our model extends that idea allowing for *multiple proxy vectors* per class, resulting in a more flexible representation.

3 Model

Formally, we learn the model in T tasks, task t comprising a set of new classes C_N^t , and a set of old classes C_O^t , and aiming at classifying all seen classes $C_O^t \cup C_N^t$.




Fig. 2. Different possible poolings. The output from a convolutional layer $\mathbf{h}_{\ell,c,w,h}^t$ may be pooled (summed over) one or more axes. The resulting loss considers only the pooled activations instead of the individual components, allowing more plasticity across the pooled axes.

Between tasks, the new set C_O^t will be set to $C_O^{t-1} \cup C_N^{t-1}$, but the amount of training samples from C_O^t (called *memory*) is constrained to exactly M_{per} samples per class, while all training samples in the dataset are allowed for the classes in C_N^t , as shown in Fig. 1. The resulting imbalance, if unmanaged, leads to *catastrophic forgetting* [7, 31], i.e., learning the new classes at the cost of forgetting the old ones.

Our base model is a deep convolutional network $\hat{\mathbf{y}} = g(f(\mathbf{x}))$, where \mathbf{x} is the input image, \mathbf{y} is the output vector of class probabilities, $\mathbf{h} = f(\mathbf{x})$ is the “feature extraction” part of the network (all layers up to the next-to-last), $\hat{\mathbf{y}} = g(\mathbf{h})$ is the final classification layer, and \mathbf{h} is the final embedding of the network before classification (Fig. 3). The superscript t denotes the model learned at task t : f^t , g^t , \mathbf{h}^t , etc.

3.1 POD: Pooled Outputs Distillation Loss

Constraining the evolution of the weights is crucial to reduce forgetting. Each new task t learns a new (student) model, whose weights are not only initialized with those of the previous (teacher) model, but also constrained by a distillation loss. That loss must be carefully balanced to prevent forgetting (rigidity), while allowing the learning of new classes (plasticity).

To this goal, we propose a set of constraints we call **Pooled Outputs Distillation (POD)**, applied not only over the final embedding output by $\mathbf{h}^t = f^t(\mathbf{x})$, but also over the output of its intermediate layers $\mathbf{h}_\ell^t = f_\ell^t(\mathbf{x})$ (where by notation overloading $f_\ell^t(\mathbf{x}) \equiv f_\ell^t \circ \dots \circ f_1^t(\mathbf{x})$, and thus $f^t(\mathbf{x}) \equiv f_L^t \dots \circ f_\ell^t \circ \dots \circ f_1^t(\mathbf{x})$).

The convolutional layers of the network output tensors \mathbf{h}_ℓ^t with components $\mathbf{h}_{\ell,c,w,h}^t$, where c stands for channel (filter), and $w \times h$ for column and row of the spatial coordinates. The loss used by POD may pool (sum over) one or several of those indexes, more aggressive poolings (Fig. 2) providing more freedom, and thus, plasticity: the lowest possible plasticity imposes an exact similarity between the previous and current model while higher plasticity relaxes the similarity definition.

Pooling is an important operation in Computer Vision, with a strong theoretical motivation. In the past, pooling has been introduced to obtain invariant representations [19, 24]. Here, the justification is similar, but the goal is different: as we will see, the pooled indexes are aggregated in the proposed loss, allowing

plasticity. Instead of the model acquiring invariance to the input image, the desired loss acquires invariance to model evolution, and thus, representation. The proposed pooling-based formalism has two advantages: first, it organizes disparately proposed distillation losses into a neat, general formalism. Second, as we will see, it allowed us to propose novel distillation losses, with better plasticity-rigidity compromises. Those topics are explored next.

Pooling of Convolutional Outputs. As explained before, POD constrains the output of each intermediate convolutional layer $\mathbf{h}_{\ell,c,w,h}^t = f_{\ell}^t(\cdot)$ (in practice, each stage of a ResNet [11]). As a reminder, c is the channel and $w \times h$ are the spatial coordinates. All POD variants use the Euclidean distance of ℓ^2 -normalize tensors, here noted as $\|\cdot - \cdot\|$. They differ on the type of pooling applied before that distance is computed. On one extreme, one can apply no pooling at all, resulting in the most strict loss, the most rigid constrains, and the lowest plasticity:

$$\mathcal{L}_{\text{POD-pixel}}(\mathbf{h}_{\ell}^{t-1}, \mathbf{h}_{\ell}^t) = \sum_{c=1}^C \sum_{w=1}^W \sum_{h=1}^H \left\| \mathbf{h}_{\ell,c,w,h}^{t-1} - \mathbf{h}_{\ell,c,w,h}^t \right\|^2. \quad (1)$$

By pooling the channels, one preserves only the spatial coordinates, resulting in a more permissive loss, allowing the activations to reorganize across the channels, but penalizing global changes of those activations across the space,

$$\mathcal{L}_{\text{POD-channel}}(\mathbf{h}_{\ell}^{t-1}, \mathbf{h}_{\ell}^t) = \sum_{w=1}^W \sum_{h=1}^H \left\| \sum_{c=1}^C \mathbf{h}_{\ell,c,w,h}^{t-1} - \sum_{c=1}^C \mathbf{h}_{\ell,c,w,h}^t \right\|^2; \quad (2)$$

or, contrarily, by pooling the space (equivalent, up to a factor, to a Global Average Pooling), one preserves *only* the channels:

$$\mathcal{L}_{\text{POD-gap}}(\mathbf{h}_{\ell}^{t-1}, \mathbf{h}_{\ell}^t) = \sum_{c=1}^C \left\| \sum_{w=1}^W \sum_{h=1}^H \mathbf{h}_{\ell,c,w,h}^{t-1} - \sum_{w=1}^W \sum_{h=1}^H \mathbf{h}_{\ell,c,w,h}^t \right\|^2. \quad (3)$$

Note that the only difference between the variants is in the position of the summation. For example, contrast equations Eqs. 1 and 2: in the former the differences are computed between activation pixels, and then totaled; in the latter, first the channel axis is flattened, then the differences are computed, resulting in a more permissive loss.

We can trade a little plasticity for rigidity, with less aggressive pooling by aggregating statistics across just one of the spatial dimensions:

$$\mathcal{L}_{\text{POD-width}}(\mathbf{h}_{\ell}^{t-1}, \mathbf{h}_{\ell}^t) = \sum_{c=1}^C \sum_{h=1}^H \left\| \sum_{w=1}^W \mathbf{h}_{\ell,c,w,h}^{t-1} - \sum_{w=1}^W \mathbf{h}_{\ell,c,w,h}^t \right\|^2; \quad (4)$$

or, likewise, for the vertical dimension, resulting in POD-height. Each of those variants measure the distribution of activation pixels across their respective

axis. These two complementary intermediate statistics can be further combined together:

$$\mathcal{L}_{\text{POD-spatial}}(\mathbf{h}_\ell^{t-1}, \mathbf{h}_\ell^t) = \mathcal{L}_{\text{POD-width}}(\mathbf{h}_\ell^{t-1}, \mathbf{h}_\ell^t) + \mathcal{L}_{\text{POD-height}}(\mathbf{h}_\ell^{t-1}, \mathbf{h}_\ell^t). \quad (5)$$

$\mathcal{L}_{\text{POD-spatial}}$ is minimal when the average statistics over the dataset, on both width and height axes, are similar for the previous and current model. It brings the right balance between being too rigid (Eq. 1) and being too permissive (Eqs. 2 and 3).

Constraining the Final Embedding. After the convolutional layers, the network, by design, flattens the spatial coordinates, and the formalism above needs adjustment, as a summation over w and h is no longer possible. Instead, we set a flat constraint on the final embedding $\mathbf{h}^t = f^t(\mathbf{x})$:

$$\mathcal{L}_{\text{POD-flat}}(\mathbf{h}^{t-1}, \mathbf{h}^t) = \|\mathbf{h}^{t-1} - \mathbf{h}^t\|^2. \quad (6)$$

Combining the Losses, Analysis. The final POD loss combines the two components:

$$\begin{aligned} \mathcal{L}_{\text{POD-final}}(\mathbf{x}) &= \frac{\lambda_c}{L-1} \sum_{\ell=1}^{L-1} \mathcal{L}_{\text{POD-spatial}}(f_\ell^{t-1}(\mathbf{x}), f_\ell^t(\mathbf{x})) \\ &\quad + \lambda_f \mathcal{L}_{\text{POD-flat}}(f^{t-1}(\mathbf{x}), f^t(\mathbf{x})). \end{aligned} \quad (7)$$

The hyperparameters λ_c and λ_f are necessary to balance the two terms, due to the different nature of the intermediate outputs (spatial and flat).

As mentioned, the strategy above generalizes disparate propositions existing both in the literature of incremental learning, and elsewhere. When $\lambda_c = 0$, it reduces to the cosine constraint of *Less-Forget*, proposed by Hou et al. for incremental learning, which constrains only the final embedding [13]. When $\lambda_f = 0$ and POD-spatial is replaced by POD-pixel, it suggests the Perceptual Features loss, proposed for style transfer [14]. When $\lambda_f = 0$ and POD-spatial is replaced by POD-channel, the strategy hints at the loss proposed by Komodakis et al. [17] to allow distillation across different networks, a situation in which the channel pooling responds to the very practical need to allow the comparison of architectures with different number of channels.

As we will see in our evaluations of pooling strategies (Subsect. 4.2), what proved optimal was a completely novel idea, POD-spatial, combining two poolings, each of which flattens one of the spatial coordinates. That relatively rigid strategy (channels and one of the spatial coordinates are considered in each half of the loss) makes intuitive sense in our context, which is *small-task* incremental learning, and thus where we expect a slow drift of the model across a single task.

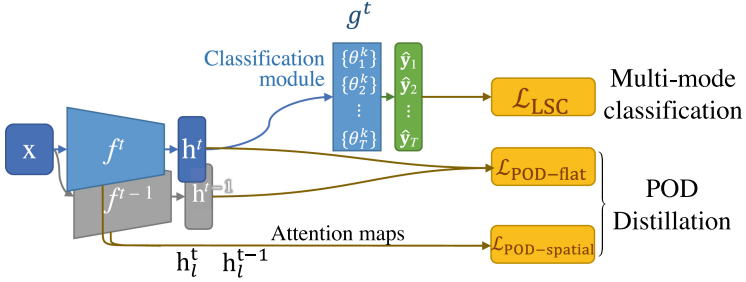


Fig. 3. Overview of PODNet: the distillation loss POD prevent excessive model drift by constraining intermediate outputs of the ConvNet f and the LSC classifier g learns a more expressive multi-modal representation.

3.2 Local Similarity Classifier

Hou et al. [13] observed that the class imbalance of incremental learning have concrete manifestations on the parameters of the final layer on classifiers, namely the weights for the over-represented (new) classes becoming much larger than those for the underrepresented (old) classes. To overcome this issue, their method (called here UCIR) ℓ^2 -normalizes both the weights and the activations, which corresponds to taking the cosine similarity instead of the dot product. For each class c , their last layer becomes

$$\hat{y}_c = \frac{\exp(\eta \langle \theta_c, \mathbf{h} \rangle)}{\sum_i \exp(\eta \langle \theta_i, \mathbf{h} \rangle)}, \quad (8)$$

where θ_c are the last-layer weights for class c , η is a learned scaling parameter, and $\langle \cdot, \cdot \rangle$ is the cosine similarity.

However, this strategy optimizes a *global similarity*: its training objective increases the similarity between the extracted features and their associated weights. For each class, the normalized weight vector acts as a *single proxy* [26], towards which the learning procedure pushes all samples in the class.

We observed that such global strategy is hard to optimize in an incremental setting. To avoid forgetting, the distillation losses (Subsect. 3.1) tries to keep the final embedding \mathbf{h} consistent through time so that the class proxies stay relevant for the classifier. Unfortunately catastrophic forgetting, while alleviated by current methods, is not solved and thus the distribution of \mathbf{h} may change. The cosine classifier is very sensitive to those changes as it models a unique majority mode through its class proxies.

Local Similarity Classifier. The problem above lead us to amend the classification layer during training, in order to consider multiple proxies/modes per class. A shift in the distribution of \mathbf{h} will have less impact on the classifier as more modes are covered.

Our redesigned classification layer, which we call Local Similarity Classifier (LSC), allows for K multiple proxies/modes during training. Like before, the proxies are a way to interpret the weight vector in the cosine similarity, thus we allow for K vectors $\boldsymbol{\theta}_{c,k}$ for each class c . The similarity $s_{c,k}$ to each proxy/mode is first computed. An averaged class similarity $\hat{\mathbf{y}}_c$ is the output of the classification layer:

$$s_{c,k} = \frac{\exp \langle \boldsymbol{\theta}_{c,k}, \mathbf{h} \rangle}{\sum_i \exp \langle \boldsymbol{\theta}_{c,i}, \mathbf{h} \rangle}, \quad \hat{\mathbf{y}}_c = \sum_k s_{c,k} \langle \boldsymbol{\theta}_{c,k}, \mathbf{h} \rangle. \quad (9)$$

The multi-proxies classifier optimizes the similarity of each sample to its ground truth class representation and minimizes all others. A simple cross-entropy loss would work, but we found empirically that the NCA loss [9, 26] converged faster. We added to the original loss a hinge $[\cdot]_+$ to keep it bounded, and a small margin δ to enforce stronger class separation, resulting in the final formulation:

$$\mathcal{L}_{\text{LSC}} = \left[-\log \frac{\exp(\eta(\hat{\mathbf{y}}_y - \delta))}{\sum_{i \neq y} \exp \eta \hat{\mathbf{y}}_i} \right]_+. \quad (10)$$

Weight Initialization for New Classes. The incremental learning setting imposes detecting new classes at each new task t . New weights $\{\boldsymbol{\theta}_{c,k} \mid \forall c \in C_N^t, \forall k \in 1 \dots K\}$ must be added to predict them. We could initialize them randomly, but the class-agnostic features of the ConvNet f , extracted by the model trained so far offer a better prior. Thus, we employ a generalization of Imprinted Weights [28] procedure to multiple modes: for each new class c , we extract the features of its training samples, use a k-means algorithm to split them into K clusters, and use the centroids of those clusters as initial values for $\boldsymbol{\theta}_{c,k}$. This procedure ensures mode diversity at the beginning of a new task and resulted in a one percentage point improvement on CIFAR100 [18].

3.3 Complete Model Formulation

Our model has the classical structure of a convolutional network $f(\cdot)$ acting as a features extractor, and a classifier $g(\cdot)$ producing a score per class. We introduced two innovations to this model: (1) our main contribution is a novel distillation loss (POD) applied all over the ConvNet, from the spatial features \mathbf{h}_ℓ to the final flat embedding \mathbf{h} ; (2) as further refinement we propose that the classifier learns a multi-modal representation that explicitly keeps multiple proxy vectors per class, increasing the model expressiveness and thus making it less sensible to shift in the distribution of \mathbf{h} . The final loss for current model $g^t \circ f^t$, i.e., the model trained for task t , is simply their addition $\mathcal{L}_{\{f^t, g^t\}} = \mathcal{L}_{\text{LSC}} + \mathcal{L}_{\text{POD-final}}$.

4 Experiments

We compare our technique (PODNet) with three state-of-the-art models. Those models are particularly comparable to ours since they all employ a sample memory with a fixed capacity. Both iCaRL [30] and UCIR [13] use the same inference method – *Nearest-Mean-Examplars* (NME), although UCIR also proposes a

second inference method based on the classifier probabilities (called here UCIR-CNN). We evaluate PODNet with both inference methods for a small scale dataset, and the later for larger scale datasets. BiC [35], while not focused on representation learning, is specially designed to be effective on large scale datasets, and thus provided an interesting baseline.

Datasets. We employ three images datasets—extensively used in the literature of incremental learning—for our experiments: CIFAR100 [18], ImageNet100 [4, 13, 35], and ImageNet1000 [4]. ImageNet100 is a subset of ImageNet1000 with only 100 classes, randomly sampled from the original 1000.

Protocol. We validate our model and the compared baselines using the challenging protocol introduced by Hou et al. [13]: we start by training the models on half the classes (i.e., 50 for CIFAR100 and ImageNet100, and 500 for ImageNet1000). Then the classes are added incrementally in steps. We divide the remaining classes equally among the steps, e.g., for CIFAR100 we could have 5 steps of 10 classes or 50 steps of 1 class. Note that a training of 50 steps is actually made of 51 different tasks: the initial training followed by the incremental steps. Models are evaluated after each step on *all the classes seen until then*. To facilitate comparison, the accuracies at the end of each step are averaged into a unique score called *average incremental accuracy* [30]. If not specified otherwise, the average incremental accuracy is the score reported in all our results.

Following Hou et al. [13], for all datasets, and all compared models, we limit the memory M_{per} to 20 images per old class. For results with different memory settings, refer to Subsect. 4.2.

Implementation Details. For fair comparison, all compared models employ the same ConvNet backbone: ResNet-32 for CIFAR100, and ResNet-18 for ImageNet. We remove the ReLU activation at the last block of each ResNet end-of-stage to provide a signed input to POD (Subsect. 3.1). We implemented our method (called here PODNet) in PyTorch [27]. We compare both ours and UCIR’s implementation [13] of iCaRL. Results of UCIR come from the implementation of Hou et al. [13]. We provide their reported results and also run their code ourselves. We used our implementation of BiC in order to compare with the same backbone. We sample our memory images using *herding selection* [30] and perform the inference with two different methods: the *Nearest-Mean-Exemplars* (NME) proposed for iCaRL, and also adopted on one of the variants of UCIR [13], and the “CNN” method introduced for UCIR (see Sect. 2). Please see the supplementary materials for the full implementation details.

4.1 Quantitative Results

The comparisons with all the state of the art are tabulated in Table 1 for CIFAR100 and Table 2 for ImageNet100 and ImageNet1000. All tables shows

Table 1. Average incremental accuracy for PODNet *vs.* state of the art. We run experiments three times (random class orders) on CIFAR100 and report averages \pm standard deviations. Models with an asterisk * are reported directly from Hou et al. [13]

New classes per step	CIFAR100			
	50 steps	25 steps	10 steps	5 steps
	1	2	5	10
<i>iCaRL</i> * [30]	—	—	52.57	57.17
iCaRL	44.20 \pm 0.98	50.60 \pm 1.06	53.78 \pm 1.16	58.08 \pm 0.59
BiC [35]	47.09 \pm 1.48	48.96 \pm 1.03	53.21 \pm 1.01	56.86 \pm 0.46
<i>UCIR (NME)</i> * [13]	—	—	60.12	63.12
UCIR (NME) [13]	48.57 \pm 0.37	56.82 \pm 0.19	60.83 \pm 0.70	63.63 \pm 0.87
<i>UCIR (CNN)</i> * [13]	—	—	60.18	63.42
UCIR (CNN) [13]	49.30 \pm 0.32	57.57 \pm 0.23	61.22 \pm 0.69	64.01 \pm 0.91
PODNet (NME)	61.40 \pm 0.68	62.71 \pm 1.26	64.03 \pm 1.30	64.48 \pm 1.32
PODNet (CNN)	57.98 \pm 0.46	60.72 \pm 1.36	63.19 \pm 1.16	64.83 \pm 0.98

the average incremental accuracy for each considered models with various number of steps on the incremental learning run. The “New classes per step” row shows the amount of new classes introduced per task.

CIFAR100. We run our comparisons on 5, 10, 25, and 50 steps with respectively 10, 5, 2, and 1 classes per step. We created three random class orders to ran each experiment thrice, reporting averages and standard deviations. For CIFAR100 only, we evaluated our model with two different kind of inference: NME and CNN. With both methods, our model surpasses all previous state of the art models on all steps. Moreover, our model relative improvement grows as the number the steps increases, surpassing existing models by 0.82, 2.81, 5.14, and 12.1 percent points (*p.p.*) for respectively 5, 10, 25, and 50 steps. Larger numbers of steps imply stronger forgetting; those results confirm that PODNet manages to reduce drastically the said forgetting. While PODNet with NME has the largest gain, PODNet with CNN also outperforms the previous state of the art by up to 8.68*p.p.* See Fig. 4 for a plot of the incremental accuracies on this dataset. In the extreme setting of 50 increments of 1 class (Fig. 4a), our model showcases large differences, with slow degradation (“*gradual forgetting*” [7]) due to forgetting throughout the run, while the other models show a quick performance collapse (“*catastrophic forgetting*”) at the start of the run.

ImageNet100. We run our comparisons on 5, 10, 25, and 50 steps with respectively 10, 5, 2, and 1 classes per step. For both ImageNet100, and ImageNet1000 we report only PODNet with CNN, as the kNN-based NME classifier did not generalize as well to larger-scale datasets. With the more complex images of

Table 2. Average incremental accuracy, PODNet *vs.* state of the art. Models with an asterisk * are reported directly from Hou et al. [13]

New classes per step	ImageNet100				Imagenet1000	
	50 steps	25 steps	10 steps	5 steps	10 steps	5 steps
	1	2	5	10	50	100
iCaRL* [30]	—	—	59.53	65.04	46.72	51.36
iCaRL [30]	54.97	54.56	60.90	65.56	—	—
BiC [35]	46.49	59.65	65.14	68.97	44.31	45.72
UCIR (NME)* [13]	—	—	66.16	68.43	59.92	61.56
UCIR (NME) [13]	55.44	60.81	65.83	69.07	—	—
UCIR (CNN)* [13]	—	—	68.09	70.47	61.28	64.34
UCIR (CNN) [13]	57.25	62.94	67.82	71.04	—	—
PODNet (CNN)	62.48	68.31	74.33	75.54	64.13	66.95
	± 0.59	± 2.45	± 0.93	± 0.26		

ImageNet100, our model also outperforms the state of the art on all tested runs, by up to 6.51*p.p.*

ImageNet1000. This dataset is the most challenging, with much greater image complexity than CIFAR100, and ten times the number of classes as ImageNet100. We evaluate the models in 5 and 10 steps, and results confirm the consistent improvement of PODNet against existing arts by up to 2.85*p.p.*

4.2 Further Analysis and Ablation Studies

Ablation Studies. Our model has two components: the distillation loss POD and the LSC classifier. An ablation study showcasing the contribution of each component is displayed in Table 3a: each additional component improves the model performance. We evaluate every ablation on CIFAR100 with 50 steps of 1 new class each. The reported metric is the average incremental accuracy. The table shows that our novel method of constraining the whole ConvNet is beneficial. Furthermore applying only POD-spatial still beats the previous state of the art by a significant margin. Using both POD-spatial and POD-flat then further increases results with a large gain. We also compare the results with the Cosine classifier [13, 25] against the Local Similarity Classifier (LSC) with NCA loss. Finally, we add LSC-CE: our classifier with multi-mode but with a simple cross-entropy loss instead of our modified NCA loss. This version brings to mind SoftTriple [29] and Infinited Mixture Prototypes [2], used in the different context of few-shot learning. The latter only considers the closest mode of each class in its class assignment, while LSC considers all modes of a class, thus, taking into account the intra-class variance. That allows LSC to decrease class similarity

Table 3. Ablation studies performed on CIFAR100 with 50 steps. We report the average incremental accuracy.

(a) Comparison of the performance of the model when disabling parts of the complete PODNet loss					(b) Comparison of distillation losses based on intermediary features. All losses evaluated with POD-flat		
Classifier	POD-flat	POD-spatial	NME	CNN	Loss	NME	CNN
Cosine			40.76	37.93	<i>None</i>	53.29	52.98
Cosine	✓		48.03	46.73	POD-pixels	49.74	52.34
Cosine		✓	54.32	57.27	POD-channels	57.21	54.64
Cosine	✓	✓	56.69	55.72	POD-gap	58.80	55.95
LSC-CE	✓	✓	59.86	57.45	POD-width	60.92	57.51
LSC			41.56	40.76	POD-height	60.64	57.50
LSC	✓		53.29	52.98	POD-spatial	61.40	57.98
LSC		✓	61.42	57.64	GradCam [5]	54.13	52.48
LSC	✓	✓	61.40	57.98	Perceptual style [14]	51.01	52.25

when intra-class variance is high (which could signal a lack of confidence in the class).

Spatial-Based Distillation. We apply our distillation loss POD differently for the flat final embedding \mathbf{h} (POD-flat) and the ConvNet’s intermediate features maps \mathbf{h}_ℓ (POD-spatial). We designed and evaluated several alternative for the latter whose results are shown in Table 3b. Refer to Sect. 3.1 and Fig. 2 for their definition. All losses are evaluated with POD-flat. “None” is using only POD-flat. Overall, we see that not using pooling results in bad performance (POD-pixels). Our final loss, POD-spatial, surpasses all others by taking advantages of the statistics aggregated from both spatial axis. For the sake of completeness we also included losses not designed by us: GradCam distillation [5] and Perceptual Style [14]. The former uses a gradient-based attention while the later – used for style transfer – computes a gram matrix for each channel.

Forgetting and Plasticity Balance. Forgetting can be drastically reduced by imposing a high factor on the distillation losses. Unfortunately, it will also degrade the capacity (its *plasticity*) to learn new classes. When POD-spatial is added on top of POD-flat, we manage to increase the oldest classes performance (+7 percentage points) while the newest classes performance were barely reduced ($-0.2p.p.$). Because our loss POD-spatial constraints only statistics, it is less stringent than a loss based on exact pixels values as POD-pixel. The latter hurts the newest classes ($-2p.p.$) for a smaller improvement of old classes ($+5p.p.$). Furthermore our experiments confirmed that LSC reduced the sensibility of the model to distribution shift, as the performance it brings was localized on the old classes.

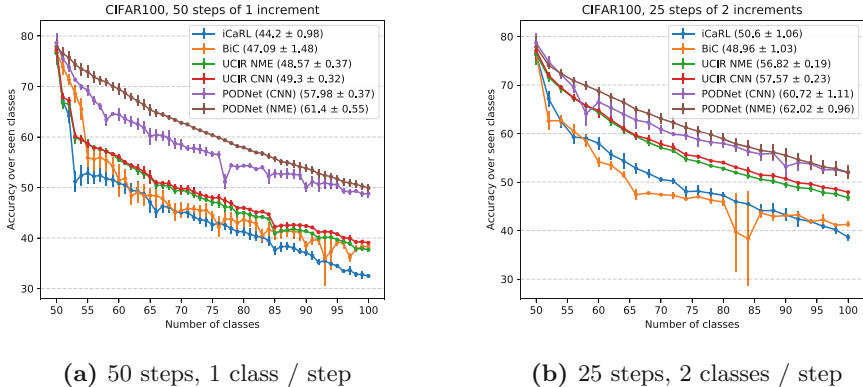


Fig. 4. Incremental Accuracy on CIFAR100 over three orders for two different step sizes. The legend reports the average incremental accuracy.

Table 4. Effect of the memory size per class M_{per} on the models performance. Results from CIFAR100 with 50 steps, we report the average incremental accuracy

M_{per}	5	10	20	50	100	200
iCaRL [30]	16.44	28.57	44.20	48.29	54.10	57.82
BiC [35]	20.84	21.97	47.09	55.01	62.23	67.47
UCIR (NME) [13]	21.81	41.92	48.57	56.09	60.31	64.24
UCIR (CNN) [13]	22.17	42.70	49.30	57.02	61.37	65.99
PODNet (NME)	48.37	57.20	61.40	62.27	63.14	63.63
PODNet (CNN)	35.59	48.54	57.98	63.69	66.48	67.62

Robustness of Our Model. While previous results showed that PODNet improved significantly over the state-of-the-arts, we wish here to demonstrate here the robustness of our model to various factors. In Table 4, we compared how PODNet behaves against the baseline when the memory size per class M_{per} changes: PODNet improvements increase as the memory size decrease, up to a gain of 26.20*p.p.* with NME (resp. 13.42*p.p.* for CNN) with $M_{per} = 5$. Notice that by default, the memory size is 20 in Subsect. 4.1. We also compared our model against baselines with a more flexible memory $M_{total} = 2000$ [30, 35], and with various initial task size (by default it is 50 on CIFAR100). In the former case, models benefit from a larger memory per class in the early tasks. In the later case, models initialization is worse because of a smaller initial task size. In these settings very different from Sect. 4.1, PODNet still outperformed significantly the compared models, proving the robustness of our model. The full results of those experiments can be found in the supplementary material.

5 Conclusion

We introduced in this paper a novel distillation loss (POD) constraining the whole convolutional network. This loss strikes a balance between reducing forgetting of old classes and learning new classes, essential for long incremental runs, by carefully chosen pooling. As a further refinement, we proposed a multi-mode similarity classifier, more robust to shift in the distribution inherent to incremental learning. Those innovations allow PODNet to outperform the previous state of the art in a challenging experimental context, with severe sample-per-class memory limitation, and long runs of many small-sized tasks, by a large margin. Extensive experiments over three datasets show the robustness of our model on different settings.

Acknowledgement. E. Valle is funded by FAPESP grant 2019/05018-1 and CNPq grants 424958/2016-3 and 311905/2017-0. This work was performed using HPC resources from GENCI-IDRIS (Grant 2019-AD011011588). We also wish to thanks Estelle Thou for the helpful discussion.

References

1. Aljundi, R., Babiloni, F., Elhoseiny, M., Rohrbach, M., Tuytelaars, T.: Memory aware synapses: learning what (not) to forget. In: Proceedings of the IEEE European Conference on Computer Vision (ECCV) (2018)
2. Allen, K., Shelhamer, E., Shin, H., Tenenbaum, J.: Infinite mixture prototypes for few-shot learning. In: International Conference on Machine Learning (ICML) (2019)
3. Castro, F.M., Marín-Jiménez, M.J., Guil, N., Schmid, C., Alahari, K.: End-to-end incremental learning. In: Proceedings of the IEEE European Conference on Computer Vision (ECCV) (2018)
4. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: a large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2009)
5. Dhar, P., Singh, R.V., Peng, K.C., Wu, Z., Chellappa, R.: Learning without memorizing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
6. Fernando, C., et al.: PathNet: evolution channels gradient descent in super neural networks. arXiv preprint library (2017)
7. French, R.: Catastrophic forgetting in connectionist networks. *Trends Cogn. Sci.* **3**(4), 128–135 (1999)
8. Gidaris, S., Komodakis, N.: Dynamic few-shot visual learning without forgetting. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
9. Goldberger, J., Hinton, G.E., Roweis, S.T., Salakhutdinov, R.R.: Neighbourhood components analysis. In: Advances in Neural Information Processing Systems (NeurIPS) (2005)
10. Golkar, S., Kagan, M., Cho, K.: Continual learning via neural pruning. In: Advances in Neural Information Processing Systems (NeurIPS), Neuro AI Workshop (2019)

11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
12. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: Advances in Neural Information Processing Systems (NeurIPS), Deep Learning and Representation Learning Workshop (2015)
13. Hou, S., Pan, X., Change Loy, C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
14. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: Proceedings of the IEEE European Conference on Computer Vision (ECCV) (2016)
15. Kemker, R., Kanan, C.: FearNet: Brain-inspired model for incremental learning. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018)
16. Kirkpatrick, J., et al.: Overcoming catastrophic forgetting in neural networks. In: Proceedings of the National Academy of Sciences (2017)
17. Komodakis, N., Zagoruyko, S.: Paying more attention to attention: improving the performance of convolutional neural networks via attention transfer. In: Proceedings of the International Conference on Learning Representations (ICLR) (2017)
18. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Technical report (2009)
19. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: spatial pyramid matching for recognizing natural scene categories. *Object Categorization: Computer and Human Vision Perspectives*. Cambridge University Press (2006)
20. Li, X., Zhou, Y., Wu, T., Socher, R., Xiong, C.: Learn to grow: a continual structure learning framework for overcoming catastrophic forgetting (2019)
21. Li, Z., Hoiem, D.: Learning without forgetting. In: Proceedings of the IEEE European Conference on Computer Vision (ECCV) (2016)
22. Lomonaco, V., Maltoni, D.: CORE50: a new dataset and benchmark for continuous object recognition. In: Annual Conference on Robot Learning (2017)
23. Lopez-Paz, D., Ranzato, M.: Gradient episodic memory for continual learning. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)* (2017)
24. Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (1999)
25. Luo, C., Zhan, J., Xue, X., Wang, L., Ren, R., Yang, Q.: Cosine normalization: Using cosine similarity instead of dot product in neural networks. In: International Conference on Artificial Neural Networks (2018)
26. Movshovitz-Attias, Y., Toshev, A., Leung, T.K., Ioffe, S., Singh, S.: No fuss distance metric learning using proxies. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017)
27. Paszke, A., et al.: Automatic differentiation in PyTorch. In: *Advances in Neural Information Processing Systems (NeurIPS), Autodiff Workshop* (2017)
28. Qi, H., Brown, M., Lowe, D.G.: Low-shot learning with imprinted weights. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
29. Qian, Q., Shang, L., Sun, B., Hu, J., Li, H., Jin, R.: SoftTriple loss: deep metric learning without triplet sampling. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2019)

30. Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: iCaRL: incremental classifier and representation learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
31. Robins, A.: Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Sci.* **7**, 123–146 (1995)
32. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-CAM: visual explanations from deep networks via gradient-based localization. In: Proceedings of the IEEE International Conference on Computer Vision (ICCV) (2017)
33. Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
34. Thrun, S.: Lifelong learning algorithms. In: Thrun, S., Pratt, L. (eds.) *Learning to Learn*, pp. 181–209. Springer, Boston, MA (1998). https://doi.org/10.1007/978-1-4615-5529-2_8
35. Wu, Y., et al.: Large scale incremental learning. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
36. Yoon, J., Yang, E., Lee, J., Hwang, S.J.: Lifelong learning with dynamically expandable networks. In: Proceedings of the International Conference on Learning Representations (ICLR) (2018)
37. Zhou, P., Mai, L., Zhang, J., Xu, N., Wu, Z., Davis, L.S.: M2KD: multi-model and multi-level knowledge distillation for incremental learning. *arXiv preprint library* (2019)