

# Continual Adaptation of Visual Representations via Domain Randomization and Meta-learning

Riccardo Volpi

Diane Larlus

Grégory Rogez

NAVER LABS Europe\*

{name.lastname}@naverlabs.com

## Abstract

Most standard learning approaches lead to fragile models which are prone to drift when sequentially trained on samples of a different nature—the well-known catastrophic forgetting issue. In particular, when a model consecutively learns from different visual domains, it tends to forget the past domains in favor of the most recent ones. In this context, we show that one way to learn models that are inherently more robust against forgetting is domain randomization—for vision tasks, randomizing the current domain’s distribution with heavy image manipulations. Building on this result, we devise a meta-learning strategy where a regularizer explicitly penalizes any loss associated with transferring the model from the current domain to different “auxiliary” meta-domains, while also easing adaptation to them. Such meta-domains are also generated through randomized image manipulations. We empirically demonstrate in a variety of experiments—spanning from classification to semantic segmentation—that our approach results in models that are less prone to catastrophic forgetting when transferred to new domains.

## 1. Introduction

Modern computer vision approaches can reach super-human performance in a variety of well-defined and isolated tasks at the expense of versatility. When confronted to a plurality of new tasks or new visual domains, they have trouble adapting, or adapt at the cost of forgetting what they had been initially trained for. This phenomenon, that has been observed for decades [39], is known as *catastrophic forgetting*. Directly tackling this issue, *lifelong learning* approaches, also known as *continual learning* approaches, are designed to continuously learn from new samples without forgetting the past.

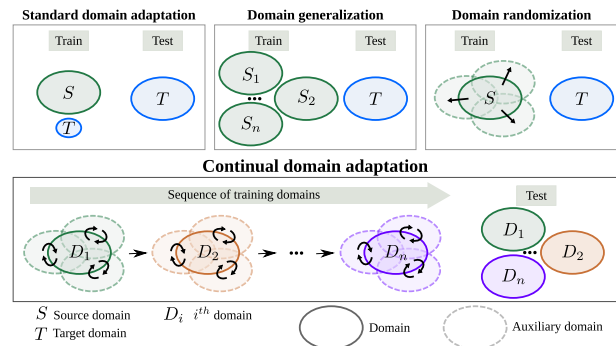


Figure 1. This paper tackles the **continual domain adaptation** task (*bottom*), compared here with standard domain adaptation, domain generalization, and domain randomization (*top*).

In this work, we are concerned with the problem of *continual* and *supervised* adaptation to new visual domains (see Figure 1). Framing this task as *continual domain adaptation*, we assume that a model has to learn to perform a given task while being exposed to conditions which constantly change throughout its lifespan. This is of particular interest when deploying applications to the real-world where a model is expected to seamlessly adapt to its environment and can encounter different domains from the one(s) observed originally at training time. A possible solution to mitigate catastrophic forgetting is storing samples from all domains encountered throughout the lifespan of the model [41]. While effective, this solution may not be appropriate when retaining data is not allowed (*e.g.*, due to privacy concerns) or when working under strong memory constraints (*e.g.*, in mobile applications). For these reasons, we are interested in developing methods for learning visual representations that are inherently more robust against catastrophic forgetting, without necessarily requiring any information storage or model expansion.

To tackle this problem, we start from a simple intuition: when adapting a model trained on a domain  $\mathcal{D}_1$  to a second domain  $\mathcal{D}_2$ , we can anticipate the severity of forgetting to depend on how demanding the adaptation process is—

\*[www.europe.naverlabs.com](http://www.europe.naverlabs.com)

that is, how close  $\mathcal{D}_2$  is to  $\mathcal{D}_1$ . The natural issue is that we cannot control whether sequential domains will be more or less similar to each other. Motivated by results on domain randomization [58, 65] and single-source domain generalization [60], we propose to heavily perturb the distribution of the current domain to increase the probability that samples from future domains will be closer to the current data distribution—hence (generally) requiring a lighter adaptation process. Focusing on computer vision tasks, we use image transformations for the randomization process, and show that models trained in this fashion are significantly more robust against catastrophic forgetting in the context of continual, supervised domain adaptation. This result represents our first contribution.

Further, we question whether we can learn representations that are *inherently* robust against transfer to new domains (that is, against gradient updates on samples from distributions different than the current one). We tackle this problem through the lens of meta-learning and devise a regularization strategy that forces the model to train for the task of interest on the current domain, while learning to be resilient to potential parameter updates on domains different from the current one. In general, meta-learning approaches require access to a number of different meta-tasks (or *meta-domains* in our case), but our setting only allows access to samples from the *current domain* at any point in time. To overcome this issue, we introduce “auxiliary” meta-domains that are produced by randomizing the original distribution—also here, using standard image transformations. Additionally, inspired by Finn et al. [18], we encourage our model to train in a way that will allow it to efficiently adapt to new domains. The devised meta-learning algorithm, based on the new concept of *auxiliary meta-domains*, constitutes our second contribution.

To extensively assess the effectiveness of our continual domain adaptation strategies, we start with an experimental protocol which tackles digit recognition. Further, we increase the difficulty of the task and focus on the PACS dataset [33] from the domain generalization literature, used here to define learning trajectories across different visual domains. Finally, we focus on semantic segmentation, exploring learning sequences across different simulated urban environments and weather conditions. In all the aforementioned experiments, we show the benefits of the proposed approaches. To conclude our analysis, we show that our methods can further be improved by combining them with a small memory of samples from previous domains [6].

## 2. Related work

Our work lies at the intersection of lifelong learning, data augmentation, meta-learning, and domain adaptation. We provide here an abridged overview of the relevant background and refer the reader to Parisi *et al.* [41],

Hospedales *et al.* [24] and Csurka [11] for detailed reviews on lifelong learning, meta-learning and domain adaptation.

**Lifelong learning.** The main goal of lifelong learning research is devising models that can learn new information throughout their lifespan, without forgetting old patterns. The survey of Parisi *et al.* [41] divides lifelong learning approaches into three categories: (i) dynamic architectures, where the model’s underlying architecture is modified as it learns new patterns [53, 67, 10, 62, 15, 64, 49]; (ii) rehearsal methods, that rely on memory replay and overcome catastrophic forgetting by storing samples from old tasks/distributions and regularly feeding them again to the model [23, 37, 5, 7, 50, 21, 46]; (iii) regularization methods, that propose ways of constraining the tasks’ objectives in order to avoid forgetting [35, 29, 66, 17]. Our work is of the third flavor: we tackle lifelong learning without necessarily replaying old data nor increasing the model capacity over time—even though we also show that our methods can be used in tandem with a small memory of old samples. In contrast with most of the prior art that focuses on task/class incremental learning, we tackle scenarios where the domain sequentially changes but the task remains the same; we refer to this problem as continual domain adaptation. See Van de Ven and Tolias for a closer look at the different problem formulations [59].

**Data augmentation and domain randomization.** The use of data augmentation has a long history in computer vision [55, 57, 9, 8, 30, 51]. Applying geometric or photometric transformations to the images of the training set generates new training samples for free and constitutes an effective strategy to improve performance. Randomizing the input distribution has been shown to be particularly effective to improve sim-to-real performance (domain randomization [58, 65]), and also to improve out-of-domain performance in single-source domain generalization problems [60] (see Figure 1, top). In this work, we first show that randomizing the domain at hand with heavy image manipulations helps preventing catastrophic forgetting in a continual domain adaptation setting. Then, we leverage similar transformations to automatically generate the samples that compose our “auxiliary” meta-domains. Our experimental results show that the proposed meta-learning strategy improves over simply using these additional samples in a standard data augmentation fashion.

**Meta-learning.** We draw inspiration from meta-learning approaches [34, 1, 48], and in particular from the idea of learning representations that can be efficiently transferred to new tasks [18]. Meta-learning generally relies on a series of meta-train and meta-test splits, and the optimization process enforces that a few gradient descent steps on the meta-train splits lead to good generalization performance

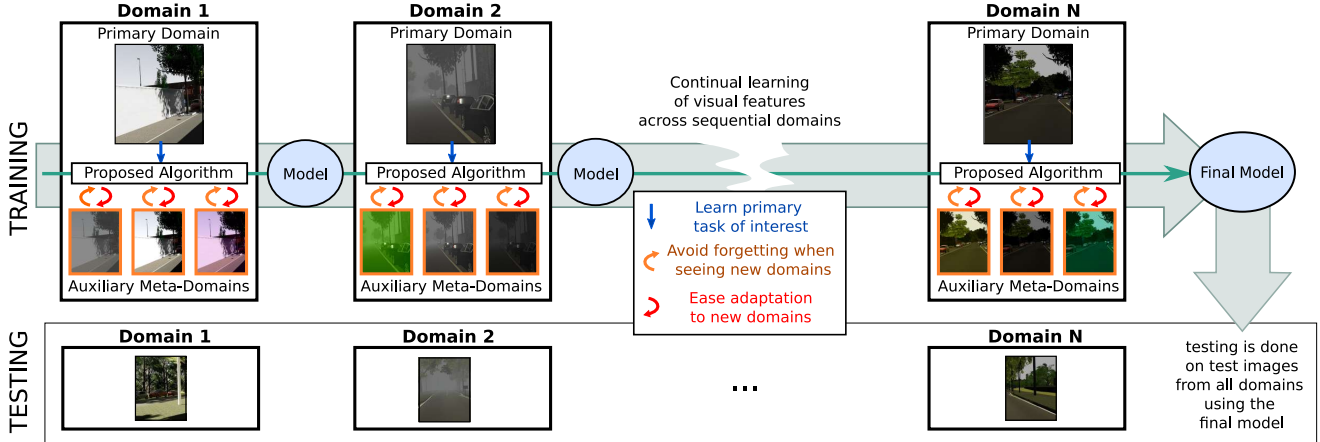


Figure 2. Life-cycle of a model when training for continual domain adaptation. At every new encountered domain, our proposed method (Algorithm 1: *Meta-DR*) is applied on the training set of that domain (Primary Domain) and on the generated “auxiliary” Meta-Domains. The final model is evaluated on test images from all the encountered domains to evaluate resilience to catastrophic forgetting.

on the meta-test. Similar to this approach, we devise a two-fold regularizer: on the one hand, it encourages models to remember previously encountered domains when exposed to new ones (by means of gradient descent updates on these tasks); on the other hand, it also encourages an efficient adaptation to such domains. In contrast with related work that relies on meta-learning to handle continual learning problems, we do not require access to a buffer of old samples [50], nor focus on learning from data streams [26].

**Domain adaptation.** Focusing on the resilience to domain-shifts rather than the more standard task-shift, our work can also be put in the context of the domain adaptation literature [56, 2, 13, 54, 19]. Our aim is indeed to perform *continual domain adaptation*, without degrading performances on past domains while facing new ones. Yet, there is a fundamental difference with the premises of our work. While, in the standard domain adaptation literature, the sole role of the source domain(s) is to compensate for the scarcity of annotated data from the target domain, in either supervised or unsupervised settings [11], our formulation does not assume such scarcity and cares about performance on the entire sequence of encountered domains.

### 3. Notations and problem formulation

First, we formally introduce the notions of catastrophic forgetting and continual domain adaptation. Then, we formalize the problem tackled in this paper and introduce the baseline we build upon.

**Notations and definitions.** Let us assume that we are interested in training a model  $\mathcal{M}_\theta$  to solve a task  $\mathcal{T}_0$ , relying on some data points that follow a distribution  $\mathcal{P}_0$ . In practice, we usually do not know this distribution but are provided with a set of samples  $\mathcal{S}_0 \sim \mathcal{P}_0$ . We fo-

cus on supervised learning and assume  $m$  training samples  $\mathcal{S}_0 = \{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^m$ , where  $\mathbf{x}_k$  and  $\mathbf{y}_k$  represent a data sample and its corresponding label, respectively. We train our model by empirical risk minimization (ERM), optimizing a loss  $\mathcal{L}_{\mathcal{T}_0}(\theta)$ . For example, for the supervised training of a multi-class classifier, this loss is typically the cross-entropy between the predictions of the model  $\hat{\mathbf{y}}$  and the ground-truth annotations  $\mathbf{y}$ :

$$\theta_{\mathcal{T}_0}^* = \min_{\theta} \left\{ \mathcal{L}_{\mathcal{T}_0}(\mathcal{S}_0; \theta) := -\frac{1}{m} \sum_{k=1}^m \mathbf{y}_k^T \log \hat{\mathbf{y}}_k \right\}. \quad (1)$$

While neural network models trained via ERM (carried out via gradient descent) have been effective in a broad range of problems, they are prone to forget about their initial task when fine-tuned on a new one, even if the two tasks appear very similar at first glance.

In practice, this means that if we use the model  $\theta_{\mathcal{T}_0}^*$  trained on the first task  $\mathcal{T}_0$  as a starting point to train for a different task  $\mathcal{T}_1$ , the newly obtained model  $\theta_{\mathcal{T}_0 \rightarrow \mathcal{T}_1}^*$  typically shows degraded performances on  $\mathcal{T}_0$ . More formally,  $\mathcal{L}_{\mathcal{T}_0}(\theta_{\mathcal{T}_0 \rightarrow \mathcal{T}_1}^*) > \mathcal{L}_{\mathcal{T}_0}(\theta_{\mathcal{T}_0}^*)$ . This undesirable property of deteriorating performance on the previously learned task is known as *catastrophic forgetting*.

**Problem formulation.** In this work, we assume that the task remains the same but that the domain varies instead. The model is sequentially exposed to a list of different domains. We wish for the model to be able to adapt to each new domain without degrading its performance on the old ones. We refer to this as *continual domain adaptation*. The goal of this paper is to mitigate catastrophic forgetting on previously seen domains.

More formally, given a task  $\mathcal{T}$  that remains constant, we assume that the model is exposed to a sequence of  $N$  domains  $(D_i)_{i=1}^N$ , each characterized by a distribution  $\mathcal{P}_i$  from which specific samples  $\mathcal{S}_i$  are drawn. As in previous

work [37], we assume locally i.i.d. data distributions. With this new formulation, and slightly abusing the notations, the problem of catastrophic forgetting mentioned before can be rewritten as  $\mathcal{L}_{\mathcal{T}}(\theta_{D_i \rightarrow D_{i+1}}^*) > \mathcal{L}_{\mathcal{T}}(\theta_{D_i}^*)$ . We assume that each set of samples  $\mathcal{S}_i$  becomes unavailable when the next domain  $D_{i+1}$  with samples  $\mathcal{S}_{i+1}$  is encountered. We are interested in assessing the performance of the model at the end of the training sequence, and for every domain  $D_i$  (see Figure 2).

**Baseline.** A naive approach to tackle the problem defined above is simply to start from the model  $\mathcal{M}_{\theta}^i$  obtained after training on domain  $D_i$  and to fine-tune it using samples from  $D_{i+1}$ . Due to catastrophic forgetting, this baseline typically performs poorly on older domains  $i < N$  when it reaches the end of its training cycle. This will be our experimental lower-bound.

## 4. Method

### 4.1. Image transformation sets

A core part of our study is the definition of proper image transformation sets, used to drive the domain randomization process—and, for what concerns the proposed meta-learning solution, to define the auxiliary meta-domains. We assume access to a set  $\Psi$ , where each element is a specific transformation with a specific magnitude level (*e.g.*, “increase brightness by 10%”). Drawing from Volpi and Murino [60], we consider a larger set covering all the possible transformations obtained by combining  $N$  given basic ones (*e.g.*, with  $N = 2$ , “increase brightness by 10% and then reduce contrast by 5%”).

Given the so-defined set and a sample set  $\mathcal{S} = \{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^m \sim \mathcal{P}$ , we can craft novel data points by sampling an object from the set  $T \sim \Psi$ , and then applying it to the given data points, obtaining  $\hat{\mathcal{S}} = \{(T(\mathbf{x}_k), \mathbf{y}_k)\}_{k=1}^m$ .

We compare different sets, with different combinations of color/geometric transformations and noise injection. As we will detail later, we can also generate arbitrary auxiliary meta-domains in this way.

**Domain randomization.** As mentioned in Section 1, part of our analysis is aimed at showing that domain randomization helps mitigating catastrophic forgetting. We perform it in a simple fashion [60]: given an annotated sample  $(\mathbf{x}, \mathbf{y}) \sim \mathcal{P}$ , before feeding it to the current model  $\theta$  and optimizing the objective in Eq. (1), we transform the image with a transformation uniformly sampled from the given set  $T \sim \Psi$ , obtaining the sample  $(T(\mathbf{x}), \mathbf{y})$ .

### 4.2. A meta-learning algorithm

**Motivation.** We aim at formulating a training objective that, at the same time, allows (i) learning a task of interest  $\mathcal{T}$ ; (ii) mitigating catastrophic forgetting when the model is

transferred to a different domain; (iii) easing adaptation to a new domain.

To achieve the second and the third goals, we follow an approach which borrows from the meta-learning literature [18]; we assume that we have access to a number of meta-domains, and exploit them to run meta-gradient updates throughout the training procedure. We enforce the loss associated with both the original domain (our training set) and the meta-domains (described later) to be small in the points reached in the weight space, both reducing catastrophic forgetting and easing adaptation.

Unfortunately, when dealing with domain  $D_i$  we do not have access to the other domains  $D_k, k \neq i$ , and hence cannot use them as meta-domains. Instead, we propose to automatically produce these meta-domains using standard image transformations, as mentioned in Section 4.1. We detail this process at the end of this section. Throughout the next paragraph, when dealing with a domain  $D_i$ , we assume the availability of different meta-domains  $\hat{D}_{i,j}$ , each defined by a set of samples  $\hat{\mathcal{S}}_{i,j}$ . To ease notation, unless in potentially ambiguous cases, we will drop the index  $i$ , *i.e.*,  $(\hat{D}_j, \hat{\mathcal{S}}_j)$ .

**Optimization problem.** Training neural networks typically involves a number of gradient descent steps that minimize a loss, see for instance Eq. (1) for the classification task.

In our training procedure, prior to every gradient descent step associated with the current domain, we simulate an arbitrary number of optimization steps to minimize the losses associated with the given meta-domains. If we are given  $K$  different meta-domains, and run a single gradient descent step on each of them at iteration  $t$ , we obtain  $K$  different points in the weight space, defined as  $\{\hat{\theta}_j^t = \theta^t - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\hat{\mathcal{S}}_j; \theta^t)\}_{j=1}^K$ , where  $j$  indicates the  $j^{\text{th}}$  meta-domain.

Our core idea is to use these weight configurations to compute the loss associated with the primary domain  $D_i$  (observed through the provided training set  $\mathcal{S}_i$ ) after adaptation,  $\{\mathcal{L}_{\mathcal{T}}(\mathcal{S}_i; \hat{\theta}_j^t)\}_{j=1}^K$ . Minimizing these loss values enforces the model to be less prone to catastrophic forgetting, according to the definition we have provided in Section 3; we define their sum as  $\mathcal{L}_{recall}$ .

Furthermore, we compute the loss values associated with samples from the meta-domains  $\{\mathcal{L}_{\mathcal{T}}(\hat{\mathcal{S}}_j; \hat{\theta}_j^t)\}_{j=1}^K$ , and define their sum as  $\mathcal{L}_{adapt}$ . If one divides the samples from the meta-domains in meta-train and meta-test subsets, minimizing  $\mathcal{L}_{adapt}$  is equivalent to running the MAML algorithm [18]. Combining the pieces together, the loss that we minimize at each step is

$$\mathcal{L} := \mathcal{L}_{\mathcal{T}}(\mathcal{S}_i; \theta^t) + \underbrace{\beta \frac{1}{K} \sum_{j=1}^K \mathcal{L}_{\mathcal{T}}(\mathcal{S}_i; \hat{\theta}_j^t)}_{\mathcal{L}_{recall}} + \underbrace{\gamma \frac{1}{K} \sum_{j=1}^K \mathcal{L}_{\mathcal{T}}(\hat{\mathcal{S}}_j; \hat{\theta}_j^t)}_{\mathcal{L}_{adapt}} \quad (2)$$

Intuitively, the three terms of this objective embody the points (i), (ii) and (iii) detailed at the beginning of this section (learning one task, avoiding catastrophic forgetting, and encouraging adaptation, respectively).

In our exposition above, we have assumed that only a single meta-optimization step is performed for each auxiliary meta-domain. In this case, computing the gradients  $\nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\hat{\theta}_j^t)$  involves the computation of the gradient of a gradient, since  $\nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\hat{\theta}_j^t) = \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\theta^t - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(\theta^t))$  [18]. One could define multi-step meta-optimization procedures, but we do not cover that extension in this work.

For Eq. (2) to be practical, we need access to meta-domains  $\hat{\mathcal{D}}_j$  during training—concretely, sample sets  $\hat{\mathcal{S}}_j$  to run the meta-updates. The core idea is to generate an arbitrary number of “auxiliary” meta-domains by modifying data points from the original training set  $\mathcal{S}_i$  via data manipulations. As already mentioned, we rely on standard image transformations introduced in Section 4.1 to do so; in practice, the idea is to start from the original training set  $\mathcal{S}_i = \{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^m \sim \mathcal{P}_i$ , sample image transformations  $T_j$  from a given set  $\Psi$ ,  $T_j \sim \Psi$ , and craft auxiliary sets as  $\hat{\mathcal{S}}_j = \{(T_j(\mathbf{x}_k), \mathbf{y}_k)\}_{k=1}^m$ —that can be used for meta-learning. The learning procedure, that we named *Meta-DR*, is detailed in Algorithm 1. In our implementation, we set  $K = 1$  in Eq. (2), and approach it via gradient descent—by randomly sampling one different auxiliary transformation prior to each step ( $T$  in line 6 being the transformation employed to generate the current auxiliary meta-domain).

---

#### Algorithm 1 *Meta-DR*

---

- 1: **Input:** training set  $\mathcal{S}_i$ , auxiliary transformation set  $\Psi = \{T_q\}_{q=1}^M$ , initial weights  $\theta^{i-1}$ , hyper-parameters  $\eta$  (learning rate),  $\alpha$  (meta-learning rate),  $\beta$  and  $\gamma$
  - 2: **Output:** weights  $\theta^i$
  - 3: **Initialize:**  $\theta \leftarrow \theta^{i-1}$
  - 4: **for**  $t = 1, \dots, H$  **do**
  - 5:   Sample  $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$  uniformly from  $\mathcal{S}_i$                    ▷ Batch for meta-update
  - 6:   Sample  $T$  uniformly from  $\Psi$                        ▷ Random transformation
  - 7:    $\hat{\theta}_T^t \leftarrow \theta^t - \alpha \nabla_{\theta} \mathcal{L}_{\mathcal{T}}(T(\hat{\mathbf{x}}), \hat{\mathbf{y}}; \theta^t)$                    ▷ Meta-update
  - 8:   Sample  $(\mathbf{x}, \mathbf{y})$  uniformly from  $\mathcal{S}_i$                    ▷ Batch for update
  - 9:    $\theta^{t+1} \leftarrow \theta^t - \eta \nabla_{\theta} \left( \underbrace{\mathcal{L}_{\mathcal{T}}(\mathbf{x}, \mathbf{y}; \theta^t)}_{\text{Current task}} + \right.$
  - 10:        $\left. + \beta \underbrace{\mathcal{L}_{\mathcal{T}}(\mathbf{x}, \mathbf{y}; \hat{\theta}_T^t)}_{\text{Backward transfer}} + \gamma \underbrace{\mathcal{L}_{\mathcal{T}}(T(\mathbf{x}), \mathbf{y}; \hat{\theta}_T^t)}_{\text{Forward transfer}} \right)$
  - 11:  $\theta^i \leftarrow \theta^{H+1}$
- 

Again, while for clarity we report one single gradient descent step for the auxiliary meta-domains in the Algorithm box (line 7), the procedure is general and can be implemented with arbitrary gradient descent trajectories. Given a sequence of domains  $(D_i)_{i=1}^N$ , we run Algorithm 1 on each of them sequentially (see Figure 2), providing the respective datasets  $\mathcal{S}_i$  as input.

## 5. Experiments

In this section, we first detail our experimental protocols for evaluating lifelong learning algorithms (Section 5.1) in a continual adaptation setting. We detail the different baselines we benchmark against in Section 5.2. Finally, we report our results in Section 5.3.

### 5.1. Experimental protocols

**Digit recognition.** We consider standard digit datasets broadly adopted by the computer vision community: MNIST [31], SVHN [40], MNIST-M [19] and SYN [19]. To assess lifelong learning performance, we propose training trajectories in which first we train on samples from one dataset, then train on samples from a second one, and so on. Given these four datasets, we propose two distinct protocols, defined by the following sequences: MNIST  $\rightarrow$  MNIST-M  $\rightarrow$  SYN  $\rightarrow$  SVHN and SVHN  $\rightarrow$  SYN  $\rightarrow$  MNIST-M  $\rightarrow$  MNIST, referred to as  $P1$  and  $P2$ , respectively. These allow to assess performance on two different scenarios, respectively: starting from easy datasets and moving to harder ones, and vice-versa.

For both protocols, we use final accuracy on *every* test set as a metric (in %); each experiment is repeated  $n = 3$  times and we report averaged results and standard deviations. For compatibility, we resize all images to  $32 \times 32$  pixels, and, for each dataset, we use 10,000 training samples. We use the standard PyTorch [42] implementation of ResNet-18 [22] in both protocols. We train models on each domain for  $H = 3 \cdot 10^3$  gradient descent steps, setting the batch size to 64. We use Adam optimizer [28] with a learning rate  $\eta = 3 \cdot 10^{-4}$ , which is re-initialized with  $\eta = 3 \cdot 10^{-5}$  after the first domain. For *Meta-DR*, we set  $\beta = \gamma = 1.0$  and  $\alpha = 0.1$  (results associated with more hyper-parameters are reported in Appendix D). We consider one set with transformations that only allow for color perturbations ( $\Psi_1$ ), one that also allows for rotations ( $\Psi_2$ ), and one that also allows for noise perturbations ( $\Psi_3$ ). See Appendix B for more details.

**PACS.** We consider the PACS dataset [33], typically used by the domain generalization community. It comprises images belonging to seven classes, drawn from four distinct visual domains: Paintings, Photos, Cartoons, and Sketches. We propose to use this dataset to assess continual learning capabilities of our models: as in the Digits protocol, we train a model in one domain, then in a second one, and so on. At the end of the learning sequence, we assess performance on the test sets of all domains. We consider the sequence *Sketches*  $\rightarrow$  *Cartoons*  $\rightarrow$  *Paintings*  $\rightarrow$  *Photos* (increasing the level of realism over time), and repeat each experiment  $n = 5$  times, reporting mean and standard deviation. In this case, we start from an ImageNet [14] pre-trained model, and resize images to  $224 \times 224$  pixels. We rely on SGD optimizer with learning rate  $\eta = 0.01$ , reduced

to  $\eta = 0.001$  after the first domain. We consider a transformation set with color transformations ( $\Psi_4$ ).

**Semantic scene segmentation.** We rely on the Virtual KITTI 2 [3] dataset to generate sequences of domains. We use 30 simulated scenes, each corresponding to one of the 5 different urban city environments and one of the 6 different weather/daylight conditions (see Figure 2). Ground-truth for several tasks is given for each datapoint; here, we focus on the semantic segmentation task.

As we will show in the next section, the most severe forgetting happens when the visual conditions change drastically (as expected); for this reason, we focus on cases where we train an initial model on samples from a particular scene, and adapt it to a novel urban environment with a different condition. In concrete terms, given three urban environments  $A, B, C$  sampled from the five available, we consider the learning sequences Clean  $\rightarrow$  Foggy  $\rightarrow$  Cloudy ( $P1$ ), Clean  $\rightarrow$  Rainy  $\rightarrow$  Foggy ( $P2$ ) and Clean  $\rightarrow$  Sunset  $\rightarrow$  Morning ( $P3$ )—where by “clean” we refer to synthetic samples cloned from the original KITTI [20] scenes. For each protocol, we randomly sample  $n = 10$  different permutations of environments  $A, B, C$  and report averaged results and standard deviations (details in Appendix A).

Since Virtual KITTI 2 [3] does not provide any default train/validation/test split, for each scene/condition we use the first 70% of the sequence for training, the next 15% for validation and the final 15% for test. We use samples from both cameras and use horizontal mirroring for data augmentation in every experiment. We consider the U-Net [52] architecture with a ResNet-34 [22] backbone pre-trained on ImageNet [14]. We train for 20 epochs on the first sequence, and for 10 epochs on the following ones, setting the batch size to 8. We use Adam optimizer [28] with a learning rate  $\eta = 3 \cdot 10^{-4}$ , which is re-initialized with  $\eta = 3 \cdot 10^{-5}$  after the first domain. For *Meta-DR*, we set  $\beta = \gamma = 10.0$  and  $\alpha = 0.01$  (results associated with more hyper-parameters are reported in Appendix D). We consider a transformation set that allows for color perturbations (see Appendix B). We rely on a publicly available semantic segmentation suite [63] that is based on PyTorch [42]. We assess performance on every domain explored during the learning trajectory, using mean intersection over union (mIoU, in  $[0, 100]$ ) as a metric.

## 5.2. Training methods

We are interested in assessing the performance of models trained via domain randomization (*Naive + DR*) and with our proposed meta-learning solution (*Meta-DR*).

First, we benchmark these methods against the *Naive* baseline introduced in Section 4, which simply finetunes the model as new training domains come along. Then, we consider two *oracle* methods: If we assume access to every domain at every point in time, we can either train

on samples from the joint distribution from the beginning ( $\mathcal{P}_0 \cup \mathcal{P}_1 \dots \cup \mathcal{P}_T$ , *oracle (all)*), or grow the distribution over iterations (first train on  $\mathcal{P}_0$ , then on  $\mathcal{P}_0 \cup \mathcal{P}_1$ , etc., *oracle (cumulative)*). With access to samples from any domain, oracles are not exposed to catastrophic forgetting; yet, their performance is not necessarily an upper bound [36].

Concerning methods devised ad hoc for continual learning, we compare against L2-regularization and EWC approaches [29]. For a fair comparison, these algorithms—and the ones below—are implemented with the same domain randomization strategies used for *Meta-DR* and *Naive + DR*. We also assess performance when an episodic memory of  $M$  samples per encountered domain is stored (by default,  $M = 100$ ). Learning procedures do not change, except that at every training step the current domain’s batch of samples is stacked with a batch from the memory (Experience Replay, or ER [6]). We further benchmark our strategies against the GEM method [37], and using different values for the memory size. Note that replay-based approaches come with storage overheads, since the episodic memory grows with the number of domains as  $\mathcal{O}(n)$ .

## 5.3. Results

**Classification (Digits and PACS).** We report in Table 1 and Table 2 the comparison between models trained on Digits and on PACS, respectively.

The first result to highlight, is the significantly improved performance of *Naive + DR* with respect to the baseline (*Naive*). For instance, observe the improved performance on SVHN samples at the end of the training sequence of protocol  $P2$ , from  $\sim 54\%$  to  $\sim 72\%$  (Table 1); or the improved performance on Sketches, from  $\sim 74\%$  to  $\sim 84\%$  (Table 2). This comes without additional overhead on the learning process, except for the neglectable computation spent to transform image samples.

Next, in some cases we can appreciate the benefit of combining domain randomization with methods devised ad hoc for continual learning, namely L2 and EWC [29]; without data augmentation, we could not detect any significant improvement when using these algorithms in our setting. Note that such regularization strategies generally perform better under the availability of a task label at test time [38, 4, 16, 27, 21]; in our settings, we desire our model to be able to process samples from any arbitrary domains without further information such as domain labels.

Our meta-learning approach compares favorably with all non-oracle approaches (see *Meta-DR* row in Tables 1 and 2). A testbed in which we perform worse than a competing algorithm is SVHN in protocol  $P2$ , where L2 regularization [29] performs better. SVHN is a more complex domains than the others, so it might be less effective simulating auxiliary meta-domains from this starting point.

To better highlight the performance gains obtained by re-

Digits experiment: comparison

Method	Protocol $P_1$				Protocol $P_2$			
	MNIST (1)	MNIST-M (2)	SYN (3)	SVHN (4)	SVHN (1)	SYN (2)	MNIST-M (3)	MNIST (4)
Naive	83.7 ± 6.4	68.8 ± 3.4	92.3 ± 0.4	86.9 ± 0.1	54.0 ± 5.8	74.9 ± 3.1	71.1 ± 1.5	98.5 ± 0.0
Naive + DR	83.4 ± 3.6	72.0 ± 1.1	95.0 ± 0.3	91.4 ± 0.1	72.3 ± 0.9	80.8 ± 0.6	89.5 ± 0.6	99.0 ± 0.0
L2 [29] + DR	85.9 ± 2.8	71.8 ± 1.8	95.4 ± 0.2	91.4 ± 0.1	75.3 ± 1.4	82.0 ± 1.3	89.4 ± 0.9	98.8 ± 0.0
EWC [29] + DR	87.2 ± 1.8	70.7 ± 1.0	95.4 ± 0.3	91.8 ± 0.1	73.3 ± 0.5	80.5 ± 0.8	89.8 ± 0.6	98.8 ± 0.1
<i>Meta-DR</i>	92.0 ± 0.6	75.1 ± 0.5	95.3 ± 0.3	91.9 ± 0.2	73.8 ± 2.1	82.4 ± 1.1	90.1 ± 0.1	99.0 ± 0.1
ER [6]	93.2 ± 0.9	77.7 ± 1.2	94.1 ± 0.2	89.2 ± 0.5	74.4 ± 0.6	86.1 ± 0.1	89.9 ± 0.3	98.6 ± 0.2
ER [6] + DR	93.9 ± 0.3	79.9 ± 0.4	95.8 ± 0.2	91.5 ± 0.3	80.6 ± 0.5	90.1 ± 1.2	89.8 ± 0.7	98.8 ± 0.1
ER [6] + <i>Meta-DR</i>	93.4 ± 0.8	79.7 ± 0.4	95.8 ± 0.5	92.4 ± 0.1	82.4 ± 0.4	90.5 ± 0.2	90.4 ± 0.1	99.0 ± 0.1
Oracle (all)	99.3 ± 0.0	93.4 ± 0.4	97.1 ± 0.2	89.9 ± 0.5	89.9 ± 0.5	97.1 ± 0.2	93.4 ± 0.4	99.3 ± 0.0
Oracle (cumul.)	99.3 ± 0.1	93.3 ± 0.2	96.6 ± 0.1	88.6 ± 0.7	90.2 ± 0.2	97.0 ± 0.1	92.5 ± 0.1	98.5 ± 0.1

Table 1. Digit classification accuracy on MNIST, MNIST-M, SYN and SVHN at the end of protocols  $P_1$  (left) and  $P_2$  (right). *Meta-DR* indicates results obtained via Algorithm 1. The same transformation set  $\Psi_3$  is used for *Meta-DR* and the methods that rely on domain randomization (+ DR). Oracles are not comparable as they can access data from all domains at anytime during training. ER-based approaches [6] rely on an episodic memory, with 100 samples per domain here.

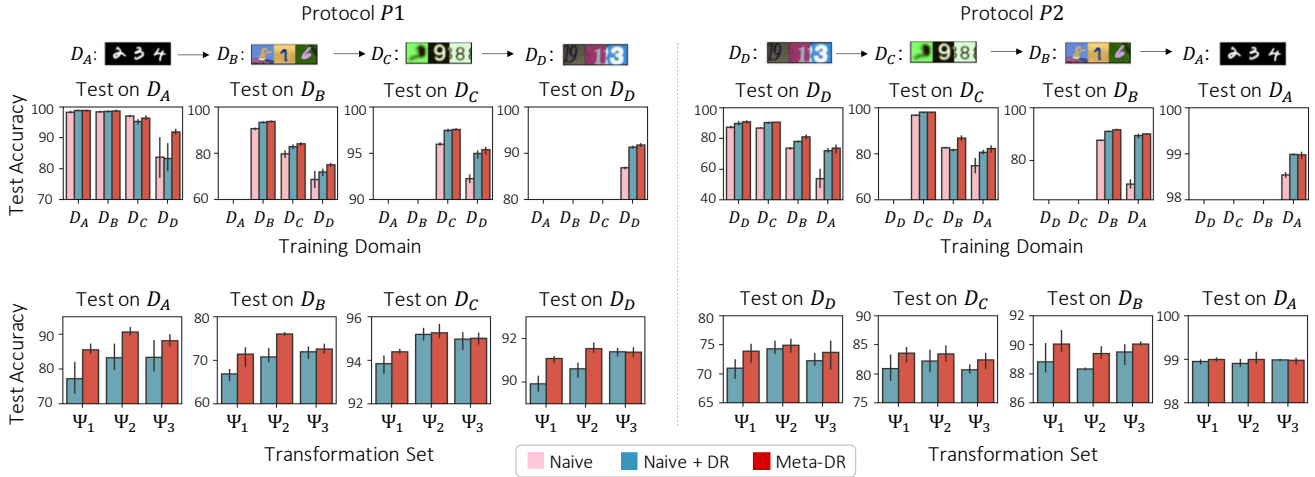


Figure 3. Digit classification accuracy for protocols  $P_1$  (left, *i.e.* MNIST  $\rightarrow$  MNIST-M  $\rightarrow$  SYN  $\rightarrow$  SVHN) and  $P_2$  (right, *i.e.* SVHN  $\rightarrow$  SYN  $\rightarrow$  MNIST-M  $\rightarrow$  MNIST) of digit experiments. (Top) Per-domain performance on the test set throughout the training sequence (after having trained on each of the four domains). (Bottom) Performance at the end of the training sequence for different transformation sets  $\Psi_i$ . The methods *Naive*, *Naive + DR* and *Meta-DR*, are associated with pink, blue and red bars, respectively.

lying on meta-learning with respect to the simpler randomization approach and to the naive approach, we report in Figure 3 (top) the accuracy evolution as the model is transferred and adapted to each of the different domains, for the two protocols ( $P_1$  and  $P_2$ , in left and right panel, respectively). *Meta-DR* (red bars) is compared with *Naive* and *Naive + DR* (pink and blue bars, respectively).

To disambiguate the contribution of the transformation sets from the contribution of the meta-learning solution, we report in Figure 3 (bottom) the performance achieved using different transformation sets  $\Psi$  to generate auxiliary meta-domains. Results are benchmarked against *Naive + DR* trained with the same sets. These results show that the meta-learning strategy consistently outperforms such baseline across several choices for the auxiliary set  $\Psi$ .

**Episodic memory.** We report the results associated with models trained with an additional memory of old samples in Tables 1 and 2 for Digits and PACS experiments, respectively (middle rows). As previous work has shown [6], the availability of a memory significantly mitigates forgetting. First, we show that if a memory is available, the methods *Naive + DR* and *Meta-DR* still positively contribute to the model performance. Next, we show that such improvement is consistent as we increase the memory size (Table 2).<sup>1</sup> More results can be found in the Supplementary Material.

**Ablation study.** We report an ablation study in Table 3, analyzing performance of models trained via *Meta-DR* (re-

<sup>1</sup>Note that, in our experiments, the ER algorithm [6] is the best performing baseline to compare against—for what concerns replay-based methods.

PACS experiment					
Methods	M.size	Sketches (1)	Cartoons (2)	Paintings (3)	Photos (4)
Naive	—	73.0 ± 2.9	71.4 ± 3.0	78.3 ± 1.6	98.8 ± 0.2
Naive + DR	—	84.3 ± 2.4	75.0 ± 2.2	80.7 ± 0.7	98.7 ± 0.1
L2 [29] + DR	—	83.9 ± 2.3	73.5 ± 2.8	81.5 ± 1.0	99.0 ± 0.2
EWG [29] + DR	—	84.8 ± 1.8	74.0 ± 2.1	80.9 ± 1.9	98.6 ± 0.1
<i>Meta-DR</i>	—	85.7 ± 1.8	75.4 ± 0.7	82.0 ± 1.9	98.5 ± 0.3
<hr/>					
GEM [37]	100	81.6 ± 2.9	85.4 ± 0.6	82.4 ± 0.9	97.4 ± 0.2
	200	84.3 ± 0.8	86.3 ± 0.8	81.4 ± 0.6	97.2 ± 0.3
	300	82.3 ± 3.0	85.6 ± 0.9	81.6 ± 0.3	96.9 ± 0.7
<hr/>					
GEM [37] + DR	100	88.3 ± 2.3	84.3 ± 0.2	84.9 ± 0.7	97.9 ± 0.2
	200	90.0 ± 0.9	85.3 ± 0.3	83.7 ± 0.2	97.6 ± 0.3
	300	89.9 ± 1.2	84.1 ± 0.5	85.0 ± 0.9	97.9 ± 0.2
<hr/>					
ER [6]	100	88.1 ± 1.1	85.0 ± 1.7	84.1 ± 1.4	98.5 ± 0.3
	200	89.5 ± 0.7	85.7 ± 0.7	84.7 ± 1.5	98.6 ± 0.4
	300	90.1 ± 0.6	86.5 ± 0.9	85.0 ± 0.9	98.5 ± 0.3
<hr/>					
ER [6] + DR	100	90.9 ± 0.5	85.1 ± 1.2	87.2 ± 1.0	98.7 ± 0.3
	200	91.3 ± 1.0	86.3 ± 0.9	87.3 ± 1.1	98.7 ± 0.2
	300	91.4 ± 0.6	87.1 ± 0.4	87.4 ± 1.2	98.5 ± 0.3
<hr/>					
ER [6] + <i>Meta-DR</i>	100	91.5 ± 0.8	87.2 ± 1.0	87.6 ± 1.4	98.4 ± 0.2
	200	92.5 ± 0.5	87.7 ± 0.8	87.8 ± 0.8	98.7 ± 0.2
	300	92.4 ± 0.6	88.1 ± 0.5	88.5 ± 0.9	98.7 ± 0.3
<hr/>					
Oracle (all)	—	93.6 ± 0.3	90.7 ± 0.7	88.4 ± 0.6	98.8 ± 0.2
Oracle (cumul)	—	92.2 ± 0.7	89.1 ± 1.0	86.3 ± 0.8	97.8 ± 0.4

Table 2. Results related to PACS [33], at the end of the training sequence *Sketches* → *Cartoons* → *Paintings* → *Photos*. GEM [37] and ER-based approaches [6] rely on an episodic memory, with a number of samples per domain indicated in the 2<sup>nd</sup> column.

Digits experiment: ablation study					
Losses		Protocol: P1			
$\mathcal{L}_{recall}$	$\mathcal{L}_{adapt}$	MNIST (1)	MNIST-M (2)	SYN (3)	SVHN (4)
		83.7 ± 6.4	68.8 ± 3.4	92.3 ± 0.4	86.9 ± 0.1
	✓	94.3 ± 0.7	76.5 ± 0.6	94.4 ± 0.0	89.5 ± 0.1
	✓	89.7 ± 0.5	74.6 ± 0.1	95.4 ± 0.1	91.9 ± 0.0
✓	✓	92.0 ± 0.6	75.1 ± 0.5	95.4 ± 0.3	91.9 ± 0.2

Table 3. Ablation study of the loss terms in Eq. (2). Performance evaluated on all domains at the end of the training sequence *P1*.

lying on the set  $\Psi_3$ ) as we include the different terms in our proposed loss (in Eq. 2). We report performance of models trained on Digits (protocol *P1*). Accuracy values were computed after having trained on the four datasets. These results show that, in this setting, the first regularizer helps retaining performance on older tasks (cf. MNIST performance with and without  $\mathcal{L}_{recall}$ ). Without the second regularizer though, performance on late tasks suffers (cf. performance on SYN and SVHN with and without  $\mathcal{L}_{adapt}$ ). The two terms together (last row) allow for good performance on early tasks as well as good adaptation to new ones.

**Semantic scene segmentation.** We report in Figure 4 (left) results related with protocols *P1*, *P2* and *P3* (top, middle and bottom, respectively). We compare *Meta-DR* with *Naive* and *Naive + DR*. Also in these settings, heavy data augmentation proved to be effective to better remember the

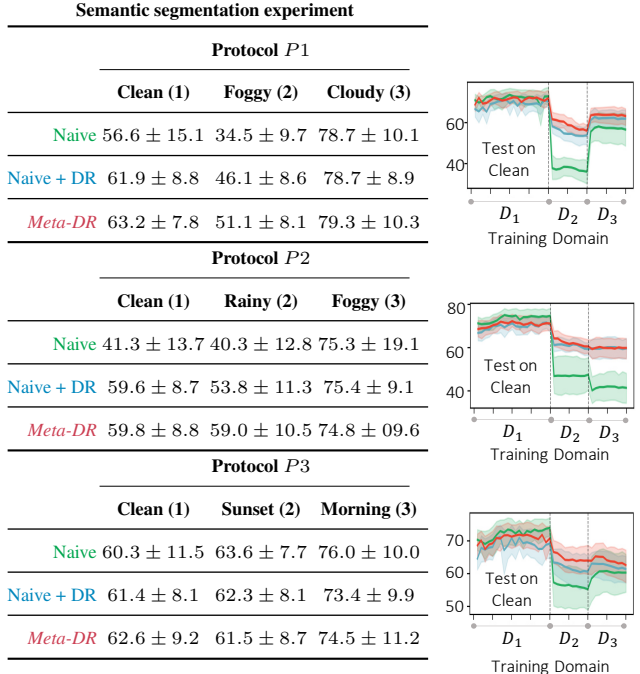


Figure 4. (Left) mIoUs on domains related to protocols *P1*, *P2*, *P3* at the end of the training sequences (top, middle and bottom, respectively). (Right) Performance on the first domain “Clean” throughout the learning sequences. Legend in Table (Left).

previous domains; in general, using *Meta-DR* allows for comparable or better performance than *Naive + DR*, using the same transformation set. We can observe smaller improvements with respect to the *Naive* baseline when the domain shift is less pronounced (*P3*): in this case, neither domain randomization nor *Meta-DR* carry the same benefit that we observed in the other protocols, or in the experiments with different benchmarks (Tables 1 and 2).

## 6. Conclusions and future directions

We propose and experimentally validate the use of domain randomization in a continual adaptation setting—more specifically, when a visual representation needs to be adapted to different domains sequentially. Through our experimental analysis, we show that heavy data augmentation plays a key role in mitigating catastrophic forgetting. Apart from the classical way of performing augmentation, we propose a more effective meta-learning approach based on the concept of “auxiliary” meta-domains, that we believe the broader meta-learning field can draw inspiration from.

For future work, we believe that designing more effective auxiliary sets might significantly improve adaptation to more diverse domains. For example, previous works focused on augmentation strategies to improve in-domain [12, 25] and out-of-domain [61, 60, 47] performance might prove helpful in our settings.

**Acknowledgments.** This work is part of MIAI@Grenoble Alpes (ANR-19-P3IA-0003).



## References

- [1] Marcin Andrychowicz, Misha Denil, Sergio Gómez, Matthew W Hoffman, David Pfau, Tom Schaul, Brendan Shillingford, and Nando de Freitas. Learning to learn by gradient descent by gradient descent. In D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, editors, *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [2] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2007.
- [3] Johann Cabon, Naila Murray, and Martin Humenberger. Virtual KITTI 2. *arXiv:2001.10773 [cs.CV]*, 2020.
- [4] Arslan Chaudhry, Puneet K. Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian Walk for Incremental Learning: Understanding Forgetting and Intransigence. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [5] Arslan Chaudhry, Marc’Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient Lifelong Learning with A-GEM. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [6] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet K. Dokania, Philip H. S. Torr, and Marc’Aurelio Ranzato. On Tiny Episodic Memories in Continual Learning. *arXiv:1902.10486 [cs.LG]*, 2019.
- [7] Arslan Chaudhry, Marcus Rohrbach, Mohamed Elhoseiny, Thalaiyasingam Ajanthan, Puneet Kumar Dokania, Philip H. S. Torr, and Marc’Aurelio Ranzato. Continual learning with tiny episodic memories. *arxiv1902.10486 [stat.ML]*, 2019.
- [8] Dan Ciresan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [9] Dan C. Ciresan, Ueli Meier, Jonathan Masci, Luca Maria Gambardella, and Jürgen Schmidhuber. High-performance neural networks for visual object classification. *arxiv:1102.0183 [cs.NE]*, 2011.
- [10] Corinna Cortes, Xavier Gonzalvo, Vitaly Kuznetsov, Mehryar Mohri, and Scott Yang. AdaNet: Adaptive Structural Learning of Artificial Neural Networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2017.
- [11] Gabriela Csurka, editor. *Domain Adaptation in Computer Vision Applications*. Advances in Computer Vision and Pattern Recognition. Springer, 2017.
- [12] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [13] Hal Daume III and Daniel Marcu. Domain adaptation for statistical classifiers. *Journal of artificial Intelligence research*, 26:101–126, 2006.
- [14] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. ImageNet: A Large-Scale Hierarchical Image Database. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [15] Timothy J. Draelos, Nadine E. Miner, Christopher C. Lamb, Jonathan A. Cox, Craig M. Vineyard, Kristofor D. Carlson, William M. Severa, Conrad D. James, and James B. Aimone. Neurogenesis Deep Learning. In *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 2017.
- [16] Sebastian Farquhar and Yarin Gal. Towards robust evaluations of continual learning. *arXiv:1805.09733 [stat.ML]*, 2018.
- [17] Enrico Fini, Stéphane Lathuilière, Enver Sangineto, Moin Nabi, and Elisa Ricci. Online continual learning under extreme memory constraints. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [18] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-Agnostic Meta-Learning for Fast Adaptation of Deep Networks. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2017.
- [19] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2015.
- [20] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the KITTI vision benchmark suite. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.
- [21] Tyler L. Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [23] Geoffrey E Hinton and David C Plaut. Using Fast Weights to Deblur Old Memories. In *Annual Conference of the Cognitive Science Society*, 1987.
- [24] Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. Meta-learning in neural networks: A survey. *arXiv:2004.05439 [cs.LG]*, 2020.
- [25] Zhiting Hu, Bowen Tan, Ruslan Salakhutdinov, Tom M. Mitchell, and Eric P. Xing. Learning data manipulation for augmentation and weighting. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [26] Khurram Javed and Martha White. Meta-learning representations for continual learning. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2019.

- [27] Ronald Kemker, Marc McClure, Angelina Abitino, Tyler L. Hayes, and Christopher Kanan. Measuring catastrophic forgetting in neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2018.
- [28] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2014.
- [29] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *PNAS*, 2017.
- [30] Alex Krizhevsky, Ilya Sutskever, and Geoffrey Hinton. ImageNet classification with deep convolutional neural networks. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [31] Yann Lecun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998.
- [32] Yann LeCun, Corinna Cortes, and CJ Burges. Mnist handwritten digit database. *ATT Labs [Online]*. Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.
- [33] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2017.
- [34] Ke Li and Jitendra Malik. Learning to Optimize. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [35] Zhizhong Li and Derek Hoiem. Learning without Forgetting. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2016.
- [36] Vincenzo Lomonaco and Davide Maltoni. CORE50: a New Dataset and Benchmark for Continuous Object Recognition. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2017.
- [37] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient Episodic Memory for Continual Learning. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [38] Davide Maltoni and Vincenzo Lomonaco. Continuous Learning in Single-Incremental-Task Scenarios. *Neural Networks*, 2019.
- [39] Michael McCloskey and Neil J. Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. *The Psychology of Learning and Motivation*, 24:104–169, 1989.
- [40] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y. Ng. Reading digits in natural images with unsupervised feature learning. In *NIPS Workshop on Deep Learning and Unsupervised Feature Learning*, 2011.
- [41] German I. Parisi, Ronald Kemker, Jose L. Part, Christopher Kanan, and Stefan Wermter. Continual Lifelong Learning with Neural Networks: A Review. *Neural Networks*, 113:54–71, 2019.
- [42] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2019.
- [43] Pillow imageenhance module. <https://pillow.readthedocs.io/en/3.0.x/reference/ImageEnhance.html>.
- [44] Pillow imageops module. <https://pillow.readthedocs.io/en/3.0.x/reference/ImageOps.html>.
- [45] Python imaging library. <https://github.com/python-pillow/Pillow>.
- [46] Ameya Prabhu, Puneet Kumar Dokania, and Philip H.S. Torr. Gdumb: A simple approach that questions our progress in continual learning. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2020.
- [47] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [48] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2017.
- [49] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H. Lampert. iCaRL: Incremental Classifier and Representation Learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [50] Matthew Riemer, Ignacio Cases, Robert Ajemian, Miao Liu, Irina Rish, Yuhai Tu, and Gerald Tesauro. Learning to Learn Without Forgetting by Maximizing Transfer and Maximizing Interference. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [51] Grégory Rogez and Cordelia Schmid. Mocap-guided data augmentation for 3d pose estimation in the wild. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, pages 3108–3116, 2016.
- [52] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Proceedings of the Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, volume 9351 of *LNCS*, pages 234–241. Springer, 2015.
- [53] Andrei A. Rusu, Neil C. Rabinowitz, Guillaume Desjardins, Hubert Soyer, James Kirkpatrick, Koray Kavukcuoglu, Razvan Pascanu, and Raia Hadsell. Progressive Neural Networks. *arXiv:1606.04671 [cs.LG]*, September 2016.

- [54] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2010.
- [55] Bernhard Schölkopf, Chris Burges, and Vladimir Vapnik. Incorporating invariances in support vector learning machines. In *International Conference on Artificial Neural Networks*, 1996.
- [56] Hidetoshi Shimodaira. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of Statistical Planning and Inference*, 90(2):227–244, 2000.
- [57] P. Y. Simard, D. Steinkraus, and J. C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Proceedings of the Seventh International Conference on Document Analysis and Recognition.*, 2003.
- [58] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [59] Gido M van de Ven and Andreas S Tolias. Three scenarios for continual learning. In *NeurIPS Continual Learning Workshop*, 2018.
- [60] Riccardo Volpi and Vittorio Murino. Model vulnerability to distributional shifts over image transformation sets. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- [61] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2018.
- [62] Tianjun Xiao, Jiaying Zhang, Kuiyuan Yang, Yuxin Peng, and Zheng Zhang. Error-Driven Incremental Learning in Deep Convolutional Neural Network for Large-Scale Image Classification. In *Proceedings of the ACM international conference on Multimedia*. ACM Press, 2014.
- [63] Pavel Yakubovskiy. Segmentation models pytorch. [https://github.com/qubvel/segmentation\\_models.pytorch](https://github.com/qubvel/segmentation_models.pytorch), 2020.
- [64] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong Learning with Dynamically Expandable Networks. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2018.
- [65] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
- [66] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual Learning Through Synaptic Intelligence. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, 2017.
- [67] Guanyu Zhou, Kihyuk Sohn, and Honglak Lee. Online Incremental Feature Learning with Denoising Autoencoders. In *Proceedings of the International Conference on Artificial Intelligence and Statistics, (AISTATS)*, 2012.