

This CVPR 2021 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

PLOP: Learning without Forgetting for Continual Semantic Segmentation

Arthur Douillard^{1,2}, Yifu Chen¹, Arnaud Dapogny³, Matthieu Cord^{1,4} ¹Sorbonne Université, ²Heuritech, ³Datakalab, ⁴valeo.ai

arthur.douillard@heuritech.com, {yifu.chen, matthieu.cord}@lip6.fr, ad@datakalab.com

Abstract

Deep learning approaches are nowadays ubiquitously used to tackle computer vision tasks such as semantic segmentation, requiring large datasets and substantial computational power. Continual learning for semantic segmentation (CSS) is an emerging trend that consists in updating an old model by sequentially adding new classes. However, continual learning methods are usually prone to catastrophic forgetting. This issue is further aggravated in CSS where, at each step, old classes from previous iterations are collapsed into the background. In this paper, we propose Local POD, a multi-scale pooling distillation scheme that preserves long- and short-range spatial relationships at feature level. Furthermore, we design an entropy-based pseudo-labelling of the background w.r.t. classes predicted by the old model to deal with background shift and avoid catastrophic forgetting of the old classes. Our approach, called PLOP, significantly outperforms state-of-the-art methods in existing CSS scenarios, as well as in newly proposed challenging benchmarks¹.

1. Introduction

Semantic segmentation is a fundamental problem of computer vision, that aims at assigning a label to each pixel of an image. In recent years, the introduction of Convolutional Neural Networks (CNNs) has addressed semantic segmentation in a traditional framework, where all classes are known beforehand and learned at once [61, 73, 11]. This setup, however, is quite limited for practical applications. In a more realistic scenario, the model should be able to continuously learn new classes without retraining from scratch. This setup, referred here as Continual Semantic Segmentation (CSS), has emerged very recently for medical applications [51, 52] before being proposed for general segmentation datasets [49, 7].

Deep learning approaches that deal with CSS face two main challenges. The first one, inherited from continual

¹Code is available at https://github.com/arthurdouillard/CVPR2021_PLOP



Figure 1: Our two-part strategy aims at learning a segmentation network in a continual learning framework, where old class pixels are collapsed into the background at current stage. We generate pseudo labels from old predictions (blue) to deal with the background shift, and retain shortand long-range spatial dependencies by Local POD distillation (red) to prevent catastrophic forgetting.

learning, is called *catastrophic forgetting* [55, 21, 62], and points to the fact that neural networks tend to completely and abruptly forget previously learned knowledge when learning new information [35]. Catastrophic forgetting presents a real challenge for continual learning applications based on deep learning methods, especially when storing previously seen data is not allowed for privacy reasons.

The second issue, CSS specific, is the semantic shift of the background class. In a traditional semantic segmentation setup, the background contains pixels that don't belong to any other class. However, in CSS, the background contains pixels that don't belong to any of the *current* classes. Thus, for a specific learning step, the background can contain both future classes, not yet seen by the model, as well as old classes. Thus, if nothing is done to distinguish pixels belonging to the real background class from old class pixels, this background shift phenomenon risks exacerbating the catastrophic forgetting even further [7].

In this paper, we propose a deep learning strategy to address these two challenges in CSS. Instead of reusing old images, our approach, called PLOP, standing for Pseudolabel and LOcal POD leverages the old model in two manners, as illustrated on Fig. 1. First, we propose a featurebased multi-scale distillation scheme to alleviate catastrophic forgetting. Second, we employ a confidence-based pseudo-labeling strategy to retrieve old class pixels within the background. For instance, if a current ground truth mask only distinguish pixels from class sofa and background, our approach allows to assign old classes to background pixels, e.g. classes person, dog or background (the semantic class).

We thoroughly validate PLOP on several datasets, showcasing significant performance improvements compared to the state-of-the-art methods in existing CSS scenarios. Furthermore, we propose several novel scenarios to further quantify the performances of CSS methods when it comes to long term learning, class presentation order and domain shift. Last but not least, we show that PLOP largely outperforms every CSS approach in these scenarios. To sum it up, our contributions are three-folds:

- We propose a multi-scale spatial distillation loss to better retain knowledge through the continual learning steps, by preserving long- and short-range spatial relationships, avoiding catastrophic forgetting.
- We introduce a confidence-based pseudo-labeling strategy to identify old classes for the current back-ground pixels and deal with background shift.
- We show that PLOP significantly outperforms state-ofthe-art approaches in existing scenarios and datasets for CSS, as well as in several newly proposed challenging benchmarks.

2. Related Work

CSS is a relatively new field where only a few recent papers addressed this specific problem. We thus start this section with a brief overview of the recent advances in semantic segmentation as well as continual learning and follow with a more in-depth discussion of existing approaches to CSS.

Semantic Segmentation methods based on Fully Convolutional Networks (FCN) [46, 59] have achieved impressive results on several segmentation benchmarks [18, 14, 77, 5]. These methods improve the segmentation accuracy by incorporating more spatial information or exploiting contextual information specifically. Atrous convolution [12, 48] and encoder-decoder architecture [57, 50, 2] are the most common methods for retaining spatial information. Examples of recent works exploiting contextual information include attention mechanisms [69, 76, 22, 30, 68, 61, 73], and

fixed-scale aggregation [75, 12, 11, 72]. More recently, Strip Pooling [28] consists in pooling along the width or height dimensions similarly to POD [17] as a complement to a spatial pyramid pooling [25] to capture both global and local statistics.

Continual Learning models generally face the challenge of catastrophic forgetting of the old classes [55, 62, 21]. Several solutions exist to address this problem: for instance, rehearsal learning consists in keeping a limited amount of training data from old classes either as raw images [55, 54, 6, 10], compressed features [24, 33], or generated training data [34, 60, 43]. Other works focus on adaptive architectures that can extend themselves to integrate new classes [67, 40] or dynamically re-arrange coexisting sub-networks [20] each specialized in one specific task [19, 23, 32], or to explicitly correct the classifier drift [66, 74, 3, 4] that happens with continually changing class distributions. Last but not least, distillation-based methods aim at constraining the model as it changes, either directly on the weights [37, 1, 8, 71], the gradients [47, 9], the output probabilities [42, 54, 6, 7], intermediary features [29, 16, 78, 17], or combinations thereof.

Continual Semantic segmentation: Despite enormous progress in the two aforementioned areas respectively, segmentation algorithms are mostly used in an offline setting, while continual learning methods generally focus on image classification. Recent works extend existing continual learning methods [42, 29] for medical applications [51, 52] and general semantic segmentation [49]. The latter considers that the previously learned categories are properly annotated in the images of the new dataset. This is an unrealistic assumption that fails to consider the background shift: pixels labeled as background at the current step are semantically ambiguous, in that they can contain pixels from old classes (including the real semantic background class, which is generally deciphered first) as well as pixels from future classes. To the best of our knowledge, Cermelli et al. [7] are the first to address this background shift problem along with catastrophic forgetting. To do so, they apply two loss terms at the output level. First, they use a knowledge distillation loss to reduce forgetting. However, only constraining the output of the network with a distillation term is not enough to preserve the knowledge of the old classes, leading to too much plasticity and, ultimately, catastrophic forgetting. Second, they propose to modify the traditional cross-entropy loss for background pixels to propagate only the sum probability of old classes throughout the continual learning steps. We argue that this constraint is not strong enough to preserve a high discriminative power w.r.t. the old classes when learning new classes under background shift. On the contrary, in what follows, we introduce our PLOP framework and show how it enables learning without forgetting for CSS.

3. PLOP Segmentation Learning Framework

3.1. Continual semantic segmentation framework

CSS aims at learning a model in $t = 1 \dots T$ steps. For each step, we present a dataset \mathcal{D}_t that consists in a set of pairs (I^t, S^t) , where I^t denotes an input image of size $W \times H$ and S^t the corresponding ground truth segmentation mask. The latter only contains the labels of current classes C^t , and all other labels (e.g. old classes $C^{1:t-1}$ or future classes $C^{t+1:T}$) are collapsed into the background class $c_{\rm bg}$. However, the model at step t shall be able to predict all the classes seen over time $\mathcal{C}^{1:t}$. Consequently, we identify two major pitfalls in CSS: the first one, catastrophic forgetting [55, 21], suggests that the network will completely forget the old classes $C^{1:t-1}$ when learning C^t . Furthermore, catastrophic forgetting is aggravated by the second pitfall, the background shift: at step t, the pixels labeled as background are indeed ambiguous, as they may contain either old (including the real background class, predicted in \mathcal{C}^1) or future classes. Fig. 2 (top row) illustrates background shift.

Classically, a deep model at step t can be written as the composition of a feature extractor $f^t(\cdot)$ and a classifier $g^t(\cdot)$. Features can be extracted at any layer l of the former $f_l^t(\cdot), l \in \{1, ... L\}$. We denote $\hat{S}^t = g^t \circ f^t(I)$ the output predicted segmentation mask and Θ^t the set of learnable parameters for the current network at step t.

3.2. Multi-scale local distillation with Local POD

A common solution to alleviate catastrophic forgetting in continual learning consists of using a distillation loss between the predictions of the old and current models [42]. This distillation loss should constitute a suitable trade-off between too much rigidity (*i.e.* enforcing too strong constraints, resulting in not being able to learn new classes) and too much plasticity (*i.e.* enforcing loose constraints, which leads to catastrophic forgetting of the old classes).

Among existing distillation schemes based on intermediate features [17, 70, 56, 16, 78, 29], POD [17] consists in matching global statistics at different feature levels between the old and current models. Let x denote an embedding tensor of size $H \times W \times C$. Extracting a POD embedding Φ consists in concatenating the $H \times C$ width-pooled slices and the $W \times C$ height-pooled slices of x:

$$\Phi(\mathbf{x}) = \left[\frac{1}{W}\sum_{w=1}^{W} \mathbf{x}[:,w,:] \middle\| \frac{1}{H}\sum_{h=1}^{H} \mathbf{x}[h,:,:] \right] \in \mathcal{R}^{(H+W) \times C}$$
(1)

where $[\cdot \| \cdot]$ denotes concatenation over the channel axis. In our case, this embedding is computed at several layers, for both the old and current model. Then the POD loss consists in minimizing the L2 distance between the two sets of embeddings over the current network parameters Θ^t :



Figure 2: Background shift example in ground truth masks (top row). At step 2 background pixels contain old (person) and future classes (bottle). The model's target (middle row) is the union of the ground-truth and the pseudo-labels (with transparent filtered uncertain pixels) generated by the previous model. The latter helps the current model predictions (bottom row) to retain information of the old classes (table).

$$\mathcal{L}_{\text{pod}}(\Theta^{t}) = \frac{1}{L} \sum_{l=1}^{L} \left\| \Phi(f_{l}^{t}(I)) - \Phi(f_{l}^{t-1}(I)) \right\|^{2}.$$
 (2)

Due to its ability to constraint spatial statistics instead of raw pixel values, this approach yields state-of-the-art results in the context of continual learning for classification. In the frame of CSS, another interest arises: its ability to model long-range dependencies across a whole axis (horizontal or vertical). However, while spatial information is discarded by global pooling in classification, semantic segmentation requires a higher degree of spatial precision. Therefore, modeling statistics across the whole width or height leads to blurring local statistics important for smaller objects.

Hence, a suitable distillation scheme for CSS shall retain both long-range and short-range spatial relationships. Thus, inspired from the multi-scale literature [38, 25], we propose a novel Local POD feature distillation scheme, that consists in computing width and height-pooled slices on multiple regions extracted at different scales $\{1/2^s\}_{s=0...S}$, as shown on Fig. 3. For an embedding tensor **x** of size $H \times W \times C$, and at scale $1/2^s$, the Local POD embedding $\Psi^s(\mathbf{x})$ at scale *s* is computed as the concatenation of s^2 POD embeddings:

$$\Psi^{s}(\mathbf{x}) = \left[\Phi(\mathbf{x}_{0,0}^{s}) \| \dots \|\Phi(\mathbf{x}_{s-1,s-1}^{s})\right] \in \mathcal{R}^{(H+W) \times C},$$
(3)

where $\forall i = 0 \dots s - 1$, $\forall j = 0 \dots s - 1$, $\mathbf{x}_{i,j}^s = \mathbf{x}[iH/s:$ (i+1)H/s, jW/s: (j+1)W/s, :] is a sub-region of the embedding tensor \mathbf{x} of size $W/s \times H/s$. We then concatenate (along channel axis) the Local POD embeddings $\Psi^s(\mathbf{x})$ of each scale s to form the final embedding:



Figure 3: Illustration of local POD. An embedding of size $W \times H \times C$ is pooled at *S* scales with POD with a spatialpyramid scheme. Here applying local POD with S = 2and scales 1 and 1/2 respectively produces 1, and 4 POD embeddings making $S \times C \times (H + W)$ dimensions total.

$$\Psi(\mathbf{x}) = \left[\Psi^{1}(\mathbf{x})\|\dots\|\Psi^{S}(\mathbf{x})\right] \in \mathcal{R}^{(H+W) \times C \times S}.$$
 (4)

We provide in the supplementary materials the complete algorithm of Local POD embedding extraction. We compute Local POD embeddings for several layers of both old and current models. The final Local POD loss is:

$$\mathcal{L}_{\text{LocalPod}}(\Theta^{t}) = \frac{1}{L} \sum_{l=1}^{L} \left\| \Psi(f_{l}^{t}(I)) - \Psi(f_{l}^{t-1}(I)) \right\|^{2} .$$
(5)

Note that while the first scale of Local POD $(1/2^0)$ is equivalent to POD and models long-range dependencies, which are important for segmentation [64, 31, 53, 28], the subsequent scales ($s = 1/2^1, 1/2^2...$) enforce short-range dependencies. This constrains the old and current models to have similar statistics over more local regions. Thus, Local POD allows retaining both long-range and short-range spatial relationships, thus alleviating catastrophic forgetting.

3.3. Solving background shift with pseudo-labeling

As described above, the pixels labelled as background at step t can belong to either old (including the semantic background class) or future classes. Thus, treating them as background would result in aggravating catastrophic forgetting. Rather, we address background shift with a pseudolabeling strategy for background pixels. Pseudo-labeling [39] is commonly used in domain adaptation for semantic segmentation [63, 41, 79, 58], where a model is trained on the union of real labels of a source dataset and pseudo labels assigned to an unlabeled target dataset. In our case, we use predictions of the old model for background pixels as clues regarding their real class, most notably if they belong to any of the old classes, as illustrated on Fig. 2 (middle row). Formally, let $C^t = card(C^t) - 1$ the cardinality of the current classes excluding the background class. Let $\hat{S}^t \in \mathcal{R}^{W,H,1+C^1+\dots+C^t}$ denote the predictions of the current model (which include the real background class, all the old classes as well as the current ones). We define $\tilde{S}^t \in \mathcal{R}^{W,H,1+C^1+\dots+C^t}$ the target as step t, computed using the one-hot ground-truth segmentation map $S^t \in \mathcal{R}^{W,H,1+C^t}$ at step t as well as pseudo-labels extracted using the old model predictions $\hat{S}^{t-1} \in \mathcal{R}^{W,H,1+C^1+\dots+C^{t-1}}$ as follows:

$$\tilde{S}^{t}(w,h,c) = \begin{cases} 1 \text{ if } S^{t}(w,h,c_{bg}) = 0 \text{ and } c = \underset{c' \in \mathcal{C}^{t}}{\operatorname{argmax}} S^{t}(w,h,c') \\ 1 \text{ if } S^{t}(w,h,c_{bg}) = 1 \text{ and } c = \underset{c' \in \mathcal{C}^{1:t-1}}{\operatorname{argmax}} \hat{S}^{t-1}(w,h,c') \\ 0 \text{ otherwise} \end{cases}$$
(6)

In other words, in the case of non-background pixels we copy the ground truth label. Otherwise, we use the class predicted by the old model $g^{t-1}(f^{t-1}(\cdot))$. This pseudo-label strategy allows to assign each pixel labelled as background his real semantic label if this pixel belongs to any of the old classes. However pseudo-labeling all background pixels can be unproductive, e.g. on uncertain pixels where the old model is likely to fail. Therefore we only retain pseudo-labels where the old model is "confident" enough. Eq. 6 can be modified to take into account this uncertainty:

$$\tilde{S}^{t}(w,h,c) = \begin{cases} 1 \text{ if } S^{t}(w,h,c_{bg}) = 0 \text{ and } c = \underset{c' \in \mathcal{C}^{t}}{\operatorname{argmax}} S^{t}(w,h,c') \\ 1 \text{ if } S^{t}\!(w,h,c_{bg}) = 1 \text{ and } c = \underset{c' \in \mathcal{C}^{1:t-1}}{\operatorname{argmax}} \hat{S}^{t-1}\!(w,h,c') \text{ and } u < \tau \\ 0 \text{ otherwise,} \end{cases}$$
(7)

where u represents the uncertainty of pixel (w, h) and τ_c is a class-specific threshold. Thus, we discard all the pixels for which the old model is uncertain $(u \ge \tau_c)$ in Eq. 7 and decrement the normalization factor WH by one. We use entropy as the uncertainty measurement u. Specifically, before learning task t, we compute the median entropy for the old model over all pixels of \mathcal{D}^t predicted as c for all the previous classes $c \in C^{1:t-1}$, which provides in thresholds $\tau_c \in C^{1:t-1}$, as proposed in [58]. The cross-entropy loss with pseudo-labeling of the old classes can be written as:

$$\mathcal{L}_{\text{pseudo}}(\Theta^{t}) = -\frac{\nu}{WH} \sum_{w,h}^{W,H} \sum_{c \in \mathcal{C}^{t}} \tilde{S}(w,h,c) \log \hat{S}^{t}(w,h,c) ,$$
(8)

where ν is the ratio of accepted old classes pixels over the total number of such pixels. This ponderation allows to adaptively weight the importance of the pseudo-labeling within the total loss. We call PLOP (standing for Pseudolabeling and LOcal Pod) the proposed approach, that uses both Local POD to avoid catastrophic forgetting, and our uncertainty-based pseudo-labeling to address background shift. To sum it up, the total loss in PLOP is:

$$\mathcal{L}(\Theta^{t}) = \underbrace{\mathcal{L}_{\text{pseudo}}(\Theta^{t})}_{\text{classification}} + \lambda \underbrace{\mathcal{L}_{\text{localPod}}(\Theta^{t})}_{\text{distillation}}, \quad (9)$$

with λ an hyperparameter.

| | 19-1 (2 tasks) | | | | 15-5 (2 tasks) | | | | 15-1 (6 tasks) | | | |
|--|-----------------------|-------|-------|-------|-----------------------|-------|-------|-------|-----------------------|-------|-------|-------|
| Method | 0-19 | 20 | all | avg | 0-15 | 16-20 | all | avg | 0-15 | 16-20 | all | avg |
| EWC [†] [37] | 26.90 | 14.00 | 26.30 | | 24.30 | 35.50 | 27.10 | | 0.30 | 4.30 | 1.30 | |
| LwF-MC [†] [54] | 64.40 | 13.30 | 61.90 | | 58.10 | 35.00 | 52.30 | | 6.40 | 8.40 | 6.90 | |
| ILT [†] [49] | 67.10 | 12.30 | 64.40 | | 66.30 | 40.60 | 59.90 | | 4.90 | 7.80 | 5.70 | |
| ILT [<mark>49</mark>] | 67.75 | 10.88 | 65.05 | 71.23 | 67.08 | 39.23 | 60.45 | 70.37 | 8.75 | 7.99 | 8.56 | 40.16 |
| MiB [†] [7] | 70.20 | 22.10 | 67.80 | | 75.50 | 49.40 | 69.00 | | 35.10 | 13.50 | 29.70 | |
| MiB [7] | 71.43 | 23.59 | 69.15 | 73.28 | 76.37 | 49.97 | 70.08 | 75.12 | 34.22 | 13.50 | 29.29 | 54.19 |
| PLOP | 75.35 | 37.35 | 73.54 | 75.47 | 75.73 | 51.71 | 70.09 | 75.19 | 65.12 | 21.11 | 54.64 | 67.21 |
| | | | | | 11 | | | | | | | |
| Table 2: Continual Semantic Segmentation results on ADE20k in Mean IoI (%) | | | | | | | | | | | | |

Table 1: Continual Semantic Segmentation results on Pascal-VOC 2012 in Mean IoU (%). †: results excerpted from [7]. Other results comes from re-implementation.

| Table 2: Continual Semantic Segmentation results on ADE20k in Mean IoU (%). | | | | | | | | | | | | |
|---|-------|---------|-------|-------|------------------------|--------|-------|-------|-------------------------|---------|-------|-------|
| 100-50 (2 tasks) | | | | | 50-50 (3 tasks) | | | | 100-10 (6 tasks) | | | |
| Method | 0-100 | 101-150 | all | avg | 0-50 | 51-150 | all | avg | 0-100 | 101-150 | all | avg |
| ILT [<mark>49</mark>] | 18.29 | 14.40 | 17.00 | 29.42 | 3.53 | 12.85 | 9.70 | 30.12 | 0.11 | 3.06 | 1.09 | 12.56 |
| MiB [7] | 40.52 | 17.17 | 32.79 | 37.31 | 45.57 | 21.01 | 29.31 | 38.98 | 38.21 | 11.12 | 29.24 | 35.12 |
| PLOP | 41.87 | 14.89 | 32.94 | 37.39 | 48.83 | 20.99 | 30.40 | 39.42 | 40.48 | 13.61 | 31.59 | 36.64 |

4. Experiments

4.1. Datasets, Protocols, and Baselines

To ensure fair comparisons with state-of-the-art approaches, we follow the experimental setup of [7] for datasets, protocol, metrics, and baseline implementations.

Datasets: we evaluate PLOP on 3 segmentation datasets: Pascal-VOC 2012 [18] (20 classes), ADE20k [77] (150 classes) and CityScapes [14] (19 classes from 21 different cities). Full details are in the supplementary materials.

CSS protocols: [7] describes two different CSS settings: Disjoint and Overlapped. In both, only the current classes are labeled vs. a background class C^t . However, in the former, images of task t only contain pixels $C^{1:t-1} \cup C^t$ (old and current), while, in the latter, pixels can belong to any classes $C^{1:t-1} \cup C^t \cup C^{t+1:T}$ (old, current, and future). Thus, the Overlapped setting is the most challenging and realistic, as in a real setting there isn't any oracle method to exclude future classes from the background. Therefore, in our experiments, we focus on Overlapped CSS but more results for Disjoint CSS can be found in the supplementary materials. While the training images are only labeled for the current classes, the testing images are labeled for all seen classes. We evaluate several CSS protocols for each dataset, e.g. on VOC 19-1, 15-5, and 15-1 respectively consists in learning 19 then 1 class (T = 2 steps), 15 then 5 classes (2 steps), and 15 classes followed by five times 1 class (6 steps). The last setting is the most challenging due to its higher number of steps. Similarly, on ADE 100-50 means 100 followed by 50 classes (2 steps), 100-10 means 100 followed by 5 times 10 classes (6 steps), and so on.

Metrics: we compare the different models using traditional mean Intersection over Union (mIoU). Specifically, we compute mIoU after the last step T for the initial classes C^1 , for the incremented classes $C^{2:T}$, and for all classes $C^{1:T}$ (all). These metrics respectively reflect the robustness to catastrophic forgetting (the model rigidity), the capacity to learn new classes (plasticity), as well as its overall performance (trade-of between both). We also introduce a novel *avg* metric (short for *average*), which measures the average of mIoU scores measured step after step, integrating performance over the whole continual learning process.

Baselines: We benchmark our model against the latest stateof-the-arts CSS methods ILT [49] and MiB [7]. We also evaluate general continual models based on weight constraints (EWC [37]) and knowledge distillation (LwF-MC [54]). More baselines are available in the supplementary materials. All models, ours included, don't use rehearsal learning [55, 54, 10] where a limited quantity of previous tasks data can be rehearsed. Finally, we also compare with a reference model learned in a traditional semantic segmentation setting ("Joint model" without continual learning), which may constitute an upper bound for CSS methods.

Implementation Details: As in [7], we use a Deeplab-V3 [13] architecture with a ResNet-101 [26] backbone pretrained on ImageNet [15] for all experiments. Full details are provided in the supplementary materials.

4.2. Quantitative Evaluation

First, we compare PLOP with state-of-the-art methods. **Pascal VOC 2012:** Table 1 shows quantitative experiments on VOC 19-1, 15-5, and 15-1. PLOP outperforms its clos-

Table 3: Mean IoU on Pascal-VOC 2012 10-1.

| | VOC 10-1 (11 tasks) | | | | | | | | |
|----------|----------------------------|-------|-------|-------|--|--|--|--|--|
| Method | 0-10 | 11-20 | all | avg | | | | | |
| ILT [49] | 7.15 | 3.67 | 5.50 | 25.71 | | | | | |
| MiB [7] | 12.25 | 13.09 | 12.65 | 42.67 | | | | | |
| PLOP | 44.03 | 15.51 | 30.45 | 52.32 | | | | | |

Table 4: Mean IoU on ADE20k 100-5.

| | ADE 100-5 (11 tasks) | | | | | |
|----------|----------------------|---------|-------|-------|--|--|
| Method | 0-100 | 101-150 | all | avg | | |
| ILT [49] | 0.08 | 1.31 | 0.49 | 7.83 | | |
| MiB [7] | 36.01 | 5.66 | 25.96 | 32.69 | | |
| PLOP | 39.11 | 7.81 | 28.75 | 35.25 | | |

est contender, MiB [7] on all evaluated settings by a significant margin. On 19-1, the forgetting of old classes (1-19) is reduced by 4.39 percentage points (p.p) while performance on new classes is greatly improved (+13.76 p.p). On 15-5, our model is on par with our re-implementation of MiB, and surpasses the original paper scores [7] by 1 p.p. On the most challenging 15-1 setting, general continual models (EWC and LwF-MC) and ILT all have very low mIoU. While MiB shows significant improvements, PLOP still outperforms it by a wide margin: +86% on all classes, +90% on old classes, and +56% on new classes. Also, the joint model mIoU is 77.40%, thus PLOP narrows the gap compared to state-of-the-art approaches on every CSS scenario. The average mIoU is also improved by +24% compared to MiB, indicating that each CSS step benefits from the improvements related to our method. This is echoed by Fig. 4, which shows that while mIoU for both ILT and MiB deteriorates after only a handful of steps, PLOP's mIoU remains very high throughout, indicating improved resilience to catastrophic forgetting and background shift.

ADE20k: Table 2 shows experiments on ADE 100-50, 100-10, and 50-50. This dataset is notoriously hard, as the joint model baseline mIoU is only 38.90%. ILT has poor performance in all three scenarios. PLOP shows comparable performance with MiB on the short setting 100-50 (only 2 tasks), improves by $1.09 \ p.p$ on the medium setting 50-50 (3 tasks), and significantly outperforms MiB with a wider margin of $2.35 \ p.p$ on the long setting 100-10 (6 tasks). In addition to being better on all settings, PLOP showcased an increased performance gain on longer CSS (e.g. 100-10) scenarios, due to increased robustness to catastrophic forgetting and background shift. To further validate this robustness, we propose harder novel CSS scenarios.



Figure 4: mIoU evolution on Pascal-VOC 2012 15-1. While MiB's mIoU quickly deteriorates, PLOP's mIoU remains high, due to improved resilience to catastrophic forgetting.



Figure 5: Boxplots of the mIoU of initial classes (1-15), new (16-20), all, and average for 20 random class orderings. PLOP is significantly better and more stable than MiB.

4.3. New Protocols and Evaluation

Longer Continual Learning: We argue that CSS experiments should push towards more steps [65, 45, 17, 6] to quantify the robustness of approaches w.r.t. catastrophic forgetting and background shift. We introduce two novel and much more challenging settings with 11 tasks, almost twice as many as the previous longest setting. We report results for VOC 10-1 in Table 3 (10 classes followed by 10 times 1 class) and ADE 100-5 in Table 4 (100 classes followed by 10 times 5 classes). The second previous State-of-the-Art method, ILT, has a very low mIoU (< 6 on VOC 10-1 and practically null on ADE 100-5). Furthermore, the gap between PLOP and MiB is even wider compared with previous benchmarks (e.g. \times 3.6 mIoU on VOC for mIoU of base classes 1-10), which confirms the superiority of PLOP when dealing with long continual processes.

Stability w.r.t. class ordering: We already showed that existing continual learning methods may be prone to instability. It has already been shown in related contexts [36] that class ordering can have a large impact on performance. However, in real-world settings, the optimal class order can never be known beforehand: thus, the performance of an ideal CSS method should be as class order-invariant as possible. In all experiments done so far, this class order has

Table 5: Final mIoU for Continual-domain Cityscapes.

| Method | 11-5 (3 tasks) | 11-1 (11 tasks) | 1-1 (21 tasks) |
|----------|----------------|------------------------|-----------------------|
| ILT [49] | 59.14 | 57.75 | 30.11 |
| MiB [7] | 61.51 | 60.02 | 42.15 |
| PLOP | 63.51 | 62.05 | 45.24 |

been kept constant, as defined in [7]. We report results in Fig. 5 under the form of boxplots obtained by applying 20 random permutations of the class order on VOC 15-1. We report in Fig. 5 (from left to right) the mIoU for the old, new classes, all classes, and average over CSS steps. In all cases, PLOP surpasses MiB in term of avg mIoU. Furthermore, the standard deviation (e.g. 10% vs 5% on *all*) is always significantly lower, showing the excellent stability of PLOP compared with existing approaches.

Domain Shift: The previous experimental setups mainly assess the capacity of CSS methods to integrate new classes, i.e. to deal with catastrophic forgetting and background shift at a semantic level. However, a domain shift can also happen in CSS scenarios. Thus, we propose a novel benchmark on Cityscapes to quantify robustness to domain shift, in which all 19 classes will be known from the start and, instead of adding new classes, each step brings a novel domain (e.g. a new city), similarly to the NI setting of [44] for image classification. Table 5 compares the performance of ILT, MiB, and PLOP on CityScapes 11-5, 11-1, and 1-1, making 3, 11 and 21 steps of 11 + 2 times 5 cities, 11 + 10 times 1 city, and 1 + 20 times 1 city respectively. PLOP performs better by a significant margin in every such scenario compared with ILT and MiB which, in this setting, is equivalent to a simple cross-entropy plus basic knowledge distillation [27]. Our Local POD, however, retains better domain-related information by modeling long and shortrange dependencies at different representation levels.

4.4. Model Introspection

We compare several distillation and classification losses on VOC 15-1 to stress the importance of the components of PLOP and report results in Table 6. All comparisons are evaluated on a val set made with 20% of the train set, therefore results are slightly different from the main experiments. **Distillation comparisons:** Table 6a compares different distillation losses when combined with our pseudo-labeling loss. As such, UNKD introduced in [7] performs better than the Knowledge Distillation (KD) of [27], but not at every step (as indicated by the *avg.* value), which indicates instability during the training process. POD, proposed in [17], improves the results on the old classes, but not on the new classes (16-20). In fact, due to too much plasticity, POD model likely overfits and predicts nothing but the new classes, hence a lower mIoU. Finally, Local POD leads to

Table 6: Comparison studies on Pascal-VOC 2012 15-1 on a validation subset of 20% of the training set.

| (a) Pseudo loss (Eq. 8) with different distillation losses. | | | | | | |
|---|----------|--------------|-------------|-----------|--|--|
| Distillation loss | 0-15 | 16-20 | all | avg | | |
| Knowledge Distillation | 29.72 | 4.42 | 23.69 | 49.18 | | |
| UNKD | 34.85 | 5.26 | 27.80 | 46.39 | | |
| POD | 43.94 | 4.82 | 34.62 | 53.35 | | |
| Local POD (Eq. 5) | 63.06 | 17.92 | 52.31 | 65.71 | | |
| (b) Local POD loss (Eq. 5 |) with d | ifferent cla | assificatio | n losses. | | |
| Classification loss | 0-15 | 16-20 | all | avg | | |
| CE only on new 1 | 2.95 | 2.54 | 10.47 | 47.02 | | |
| CE 3 | 3.80 | 4.67 | 26.87 | 50.79 | | |
| UNCE 4 | 8.46 | 4.82 | 38.62 | 53.19 | | |
| Pseudo (Eq. 8) 6 | 53.06 | 17.92 | 52.31 | 65.71 | | |
| Pseudo-Oracle 6 | 53.69 | 23.35 | 54.09 | 66.05 | | |

superior performance (+20 *p.p*) w.r.t. all metrics, due to its integration of both long and short-range dependencies. This final row represents our full PLOP strategy.

Classification comparisons: Table 6b compares different classification losses when combined with our Local POD distillation loss. Cross-Entropy (CE) variants perform poorly, especially on new classes. UNCE, introduced in [7], improves by merging the background with old classes, however, it still struggles to correctly model the new classes, whereas our pseudo-labeling propagates more finely information of the old classes, while learning to predict the new ones, dramatically enhancing the performance in both cases. This penultimate row represents our full PLOP strategy. Also notice that the performance for pseudo-labeling is very close to Pseudo-Oracle (where the incorrect pseudo-labels are removed), which may constitute a performance ceiling of our uncertainty measure. A comparison between these two results illustrates the relevance of our entropy-based uncertainty estimate.

Vizualisation: Fig. 6 shows the predictions for both MiB and PLOP on VOC 15-1 across time. At first, both models output equivalent predictions. However, MiB quickly forgets the previous classes and becomes biased towards new classes. On the other hand, PLOP predictions are much more stable on old classes while learning new classes, thanks to Local POD alleviating catastrophic forgetting by spatially constraining representations, and pseudo-labeling dealing with background shift. Fig. 7 more closely highlights this phenomenon: at first, the ground-truth only contains the class person. At step 5, the class train is introduced. As a result, MiB overfits on train and forgets person and predicts decent segmentation for both classes.



Figure 6: Visualization of MiB and PLOP predictions across time in VOC 15-1 for two test images. MiB quickly forgets the initial 15 classes (row 1: person and table, row 3: bird) in favor of new classes (plant, sheep, sofa, train) and is biased towards new classes. PLOP, however, barely suffers from catastrophic forgetting (rows 2+4).



Figure 7: Visualization of MiB and PLOP predictions across time in VOC 15-1 on a test set image. At steps 1-4 only class person has been seen. At step 5, the class train is introduced, causing dramatic background shift. While MiB overfits on the new class and forget the old class, PLOP is able to predict both classes correctly.

5. Conclusion

In this paper, we paved the way for future research on Continual Semantic Segmentation, which is an emerging domain in computer vision. We highlighted two main challenges in Continual Semantic Segmentation (CSS), namely catastrophic forgetting and background shift. To deal with the former, we proposed Local POD, a multi-scale pooling distillation scheme that allows preserving long and shortrange spatial relationships between pixels, leading to a suitable trade-off between rigidity and plasticity for CSS and, ultimately, alleviating catastrophic forgetting. The proposed method is general enough to be used in other related distillation settings, where preserving spatial information is a concern. In addition, we introduced a new strategy to address the background shift based on an efficient pseudolabeling method. We validate our PLOP framework, on several existing CSS scenarios involving multiple datasets. In addition, we propose novel experimental scenarios to assess the performance of future CSS approaches in terms of long term learning capacity and stability. We showed that PLOP performs significantly better than all existing baselines in every such CSS benchmark.

Acknowledgments: This work was granted access to the HPC resources of IDRIS under the allocation AD011011706 made by GENCI.

References

- Rahaf Aljundi, Francesca Babiloni, Mohamed Elhoseiny, Marcus Rohrbach, and Tinne Tuytelaars. Memory aware synapses: Learning what (not) to forget. In *Proceedings of the IEEE European Conference on Computer Vision* (ECCV), 2018.
- [2] V. Badrinarayanan, A. Kendall, and R. Cipolla. Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2017.
- [3] Eden Belouadah and Adrian Popescu. Il2m: Class incremental learning with dual memory. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [4] Eden Belouadah and Adrian Popescu. Scail: Classifier weights scaling for class incremental learning. In Proceedings of the IEEE Winter Conference on Application of Computer Vision (WACV), 2020.
- [5] Holger Caesar, Jasper R. R. Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *Proceed-ings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [6] Francisco M. Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2018.
- [7] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulò, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [8] Arslan Chaudhry, Puneet Dokania, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Riemannian walk for incremental learning: Understanding forgetting and intransigence. *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2018.
- [9] 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.
- [10] 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. In *International Conference on Machine Learning (ICML) Workshop*, 2019.
- [11] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2018.
- [12] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. In *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2018.
- [13] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. In *arXiv preprint library*, 2017.

- [14] M. Cordts, M. Omran, S. Ramos, T. Reheld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision* and Pattern Recognition (CVPR), 2016.
- [15] 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.
- [16] Prithviraj Dhar, Rajat Vikram Singh, Kuan-Chuan Peng, Ziyan Wu, and Rama Chellappa. Learning without memorizing. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
- [17] Arthur Douillard, Matthieu Cord, Charles Ollion, Thomas Robert, and Eduardo Valle. Podnet: Pooled outputs distillation for small-tasks incremental learning. In Proceedings of the IEEE European Conference on Computer Vision (ECCV), 2020.
- [18] Mark Everingham, S. M. Ali Eslami, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes challenge: A retrospective. In *International Journal of Computer Vision (IJCV)*, 2015.
- [19] Chrisantha Fernando, Dylan Banarse, Charles Blundell, Yori Zwols, David Ha, Andrei A. Rusu, Alexander Pritzel, and Daan Wierstra. PathNet: Evolution Channels Gradient Descent in Super Neural Networks. arXiv preprint library, 2017.
- [20] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In Proceedings of the International Conference on Learning Representations (ICLR), 2019.
- [21] Robert French. Catastrophic forgetting in connectionist networks. *Trends in cognitive sciences*, 1999.
- [22] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual attention network for scene segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [23] Siavash Golkar, Michael Kagan, and Kyunghyun Cho. Continual learning via neural pruning. Advances in Neural Information Processing Systems (NeurIPS) Workshop, 2019.
- [24] Tyler L. Hayes, Kushal Kafle, Robik Shrestha, Manoj Acharya, and Christopher Kanan. Remind your neural network to prevent catastrophic forgetting. In *Proceed*ings of the IEEE European Conference on Computer Vision (ECCV), 2020.
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2014.
- [26] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [27] Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. Distilling the knowledge in a neural network. In Advances in Neural Information Processing Systems (NeurIPS) Workshop, 2015.
- [28] Qibin Hou, Li Zhang, Ming-Ming Cheng, and Jiashi Feng. Strip pooling: Rethinking spatial pooling for scene parsing.

In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2020.

- [29] Saihui Hou, Xinyu Pan, Chen Change Loy, Zilei Wang, and Dahua Lin. Learning a unified classifier incrementally via rebalancing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [30] Z. Huang, X. Wang, L. Huang, C. Huang, Y. Wei, and W. Liu. Ccnet: Criss-cross attention for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [31] Zilong Huang, Xinggang Wang, Yunchao Wei, Lichao Huang, Humphrey Shi, Wenyu Liu, and Thomas S. Huang. Ccnet: Criss-cross attention for semantic segmentation. 2020.
- [32] Steven C.Y. Hung, Cheng-Hao Tu, Cheng-En Wu, Chien-Hung Chen, Yi-Ming Chan, and Chu-Song Chen. Compacting, picking and growing for unforgetting continual learning. In Advances in Neural Information Processing Systems (NeurIPS), 2019.
- [33] Ahmet Iscen, Jeffrey Zhang, Svetlana Lazebnik, and Cordelia Schmid. Memory-efficient incremental learning through feature adaptation. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2020.
- [34] Ronald Kemker and Christopher Kanan. Fearnet: Braininspired model for incremental learning. In Proceedings of the International Conference on Learning Representations (ICLR), 2018.
- [35] 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.
- [36] Dahyun Kim, Jihwan Bae, Yeonsik Jo, and Jonghyun Choi. Incremental learning with maximum entropy regularization: Rethinking forgetting and intransigence. arXiv preprint library, 2019.
- [37] 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. *Proceedings of the National Academy of Sciences*, 2017.
- [38] Svetlana Lazebnik, Cordelia Schmid, and Jean Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. *Object Categorization: Computer and Human Vision Perspectives, Cambridge Uni*versity Press, 2006.
- [39] Dong-Hyun Lee. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *International Conference on Machine Learning (ICML) Workshop*, 2013.
- [40] Xilai Li, Yingbo Zhou, Tianfu Wu, Richard Socher, and Caiming Xiong. Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019.
- [41] Yunsheng Li, Lu Yuan, and Nuno Vasconcelos. Bidirectional learning for domain adaptation of semantic segmentation. In

Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.

- [42] Z. Li and D. Hoiem. Learning without forgetting. Proceedings of the IEEE European Conference on Computer Vision (ECCV), 2016.
- [43] Yaoyao Liu, Yuting Su, An-An Liu, Bernt Schiele, and Qianru Sun. Mnemonics training: Multi-class incremental learning without forgetting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (CVPR), 2020.
- [44] Vincenzo Lomonaco and Davide Maltoni. Core50: a new dataset and benchmark for continuous object recognition. In *Annual Conference on Robot Learning*, 2017.
- [45] Vincenzo Lomonaco, Davide Maltoni, and Lorenzo Pellegrini. Rehearsal-free continual learning over small non-i.i.d. batches. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop, 2020.
- [46] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [47] David Lopez-Paz and Marc'Aurelio Ranzato. Gradient episodic memory for continual learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.
- [48] Sachin Mehta, Mohammad Rastegari, Anat Caspi, Linda G. Shapiro, and Hannaneh Hajishirzi. Efficient spatial pyramid of dilated convolutions for semantic segmentation. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2018.
- [49] Umberto Michieli and Pietro Zanuttigh. Incremental learning techniques for semantic segmentation. In Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshop, 2019.
- [50] H. Noh, S. Hong, and B. Han. Learning deconvolution network for semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015.
- [51] Firat Ozdemir, Philipp Fuernstahl, and Orcun Goksel. Learn the new, keep the old: Extending pretrained models with new anatomy and images. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2018.
- [52] Firat Ozdemir and Orcun Goksel. Extending pretrained segmentation networks with additional anatomical structures. In *International journal of computer assisted radiology and surgery*, 2019.
- [53] Sangyong Park and Yong Seok Heo. Knowledge distillation for semantic segmentation using channel and spatial correlations and adaptive cross entropy. In *Sensors*, 2020.
- [54] 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.
- [55] Anthony Robins. Catastrophic forgetting, rehearsal and pseudorehearsal. *Connection Science*, 1995.

- [56] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. arXiv preprint library, 2014.
- [57] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI), 2015.
- [58] Antoine Saporta, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Esl: Entropy-guided self-supervised learning for domain adaptation in semantic segmentation. In *Proceedings* of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop, 2020.
- [59] Pierre Sermanet, David Eigen, Xiang Zhang, Michaë Mathieu, Rob Fergus, and Yann LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. In *Proceedings of the International Conference* on Learning Representations (ICLR), 2014.
- [60] Hanul Shin, Jung Kwon Lee, Jaehong Kim, and Jiwon Kim. Continual learning with deep generative replay. In Advances in Neural Information Processing Systems (NeurIPS), 2017.
- [61] Andrew Tao, Karan Sapra, and Bryan Catanzaro. Hierarchical multi-scale attention for semantic segmentation. In *arXiv preprint library*, 2020.
- [62] Sebastian Thrun. Lifelong learning algorithms. In *Springer Learning to Learn*, 1998.
- [63] Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2019.
- [64] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Standalone axial-attention for panoptic segmentation. In Proceedings of the IEEE European Conference on Computer Vision (ECCV), 2020.
- [65] Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition for continual learning. In Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [66] Yue Wu, Yinpeng Chen, Lijuan Wang, Yuancheng Ye, Zicheng Liu, Yandong Guo, and Yun Fu. Large scale incremental learning. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition (CVPR), 2019.
- [67] 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.
- [68] Yuhui Yuan, Xilin Chen, and Jingdong Wang. Objectcontextual representations for semantic segmentation. In *Proceedings of the IEEE European Conference on Computer Vision (ECCV)*, 2020.
- [69] Yuhui Yuan and Jingdong Wang. Ocnet: Object context network for scene parsing. In *arXiv preprint library*, 2018.
- [70] Sergey Zagoruyko and Nikos Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. *Proceedings*

of the International Conference on Learning Representations (ICLR), 2016.

- [71] Friedemann Zenke, Ben Poole, and Surya Ganguli. Continual learning through synaptic intelligence. In *International Conference on Machine Learning (ICML)*, 2017.
- [72] Hang Zhang, Kristin J. Dana, Jianping Shi, Zhongyue Zhang, Xiaogang Wang, Ambrish Tyagi, and Amit Agrawal. Context encoding for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [73] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Muller, R. Manmatha, Mu Li, and Alexander Smola. Resnest: Splitattention networks. In *arXiv preprint library*, 2020.
- [74] Bowen Zhao, Xi Xiao, Guojun Gan, Bin Zhang, and Shutao Xia. Maintaining discrimination and fairness in class incremental learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [75] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia. Pyramid scene parsing network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [76] Hengshuang Zhao, Yi Zhang, Shu Liu, Jianping Shi, Chen Change Loy, Dahua Lin, and Jiaya Jia. Psanet: Pointwise spatial attention network for scene parsing. In Proceedings of the IEEE European Conference on Computer Vision (ECCV), 2018.
- [77] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [78] Peng Zhou, Long Mai, Jianming Zhang, Ning Xu, Zuxuan Wu, and Larry S. Davis. M2kd: Multi-model and multilevel knowledge distillation for incremental learning. *arXiv* preprint library, 2019.
- [79] Yang Zou, Zhiding Yu, BVK Vijaya Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In Proceedings of the IEEE European Conference on Computer Vision (ECCV), 2018.