Mitigating Background Shift in Class-Incremental Semantic Segmentation

Gilhan Park[®], WonJun Moon[®], SuBeen Lee[®], Tae-Young Kim[®], Jae-Pil Heo^{*}

Sungkyunkwan University, South Korea

Abstract. Class-Incremental Semantic Segmentation (CISS) aims to learn new classes without forgetting the old ones, using only the labels of the new classes. To achieve this, two popular strategies are employed: 1) pseudo-labeling and knowledge distillation to preserve prior knowledge; and 2) background weight transfer, which leverages the broad coverage of background in learning new classes by transferring background weight to the new class classifier. However, the first strategy heavily relies on the old model in detecting old classes while undetected pixels are regarded as the background, thereby leading to the background shift towards the old classes (*i.e.*, misclassification of old class as background). Additionally, in the case of the second approach, initializing the new class classifier with background knowledge triggers a similar background shift issue, but towards the new classes. To address these issues, we propose a backgroundclass separation framework for CISS. To begin with, selective pseudolabeling and adaptive feature distillation are to distill only trustworthy past knowledge. On the other hand, we encourage the separation between the background and new classes with a novel orthogonal objective along with label-guided output distillation. Our state-of-the-art results validate the effectiveness of these proposed methods. Our code is available at: https://github.com/RoadoneP/ECCV2024_MBS.

Keywords: Class-Incremental Semantic Segmentation · Continual Learning · Semantic Segmentation · Knowledge Distillation

1 Introduction

Semantic segmentation, fundamental for applications such as autonomous driving and medical imaging [4, 34, 38, 39], is a task to classify pixel-wise semantics within predefined classes on specific datasets. However, in practical scenarios, it is often necessary for models to learn additional classes after deployment. The primary goal of Class-Incremental Semantic Segmentation (CISS) is to extend the model's capability only with the supervision of newly introduced classes, without forgetting the old ones.

To achieve this, existing strategies [2,11,36,48] have adopted 1) pseudo-labeling and knowledge distillation to retain old class knowledge, *i.e.*, prevent catastrophic forgetting [14,27], as described in Fig. 1 (a). Furthermore, 2) background weight

^{*} Corresponding author



Fig. 1: (a) Pseudo-labeling is used to learn unlabeled old classes in the image based on the prediction of the old model, while knowledge distillation is to retain the intermediate knowledge of old classes by minimizing the difference between features from the old and new models. (b) However, typical models hardly recognize all pixels precisely. Therefore, the ambiguous pixels (with low prediction confidences from old model) are labeled as the background, *i.e.*, chair and sofa in Case 1 and horse in Case 2, causing the background shift towards old classes (*i.e.*, misclassification of old classes as background). In contrast, our method alleviates the background shift towards old classes by ignoring these ambiguous pixels. (c) On the other hand, background weight transfer leverages the broad category coverage of the background class by initializing the new class token with the background token parameters. (d) Despite its advantages, the baseline model faces challenges in clearly distinguishing new class tokens from the background ones (i.e., background shift towards new classes). Conversely, our method demonstrates improved separation of these classes while preserving the benefits of background weight transfer.

transfer is another popular technique [1] that duplicates background classifier weight to new class weight to exploit the broad semantic coverage of the background in learning new classes as illustrated in Fig. 1 (c). Nevertheless, we point out that these strategies remain vulnerable to the issue of background shift, which complicates the distinction between background and object classes. Specifically, the first strategy highly relies on the old model and regards ambiguous pixels as the background, triggering the background shift towards the old classes, as described in Fig. 1 (b). On the other hand, the second approach triggers the background shift towards new classes since the separation of the new class from the background is a challenging problem, as illustrated in Fig. 1 (d).

With these motivations, we propose a background-class separation framework for CISS to prevent background shifts towards both old and new classes. Initially, we propose a selective strategy during the pseudo-labeling process. To illustrate, whereas the pixels with low prediction confidence from the old model are regarded as the background class for pseudo-labeling, our strategy selectively assigns the background label only if the probability of being the actual object is low. This approach effectively reduces the background shift towards old classes by preventing the possible old class samples from being treated as the background. Similarly, adaptive feature distillation shares the same motivation that features for corresponding ambiguous patches are not trustworthy; we use the patch-wise prediction confidence as the weight to calibrate the degree of distillation for each patch. Although we mitigate the background shift towards old classes with two formerly discussed distillation methods, the background shift towards new classes remains due to the background weight transfer strategy. To tackle this, we employ label-guided output distillation which enables the decoupling of background and new classes [45] while benefiting from the background weight transfer strategy. In addition, we also introduce the novel orthogonal objective between the new classes and background tokens to further mitigate the background shift towards new classes.

To sum up, our contributions are as follows:

- We devise a selective pseudo-labeling strategy that excludes ambiguous pixels to prevent the object pixels from being misclassified as the background, mitigating the background shift towards old classes.
- To further alleviate the background shift towards old classes, we propose an adaptive feature distillation to distill only the reliable representations.
- We introduce an orthogonal objective to mitigate the background shift towards new classes while keeping the advantage of it.
- The benefits of each component and the superiority of our full method are validated with ablation studies and state-of-the-art results on all datasets.

2 Related Work

2.1 Incremental Learning

Recently, incremental learning has been extensively studied to overcome the catastrophic forgetting [14, 27] issue of deep learning models in incremental scenarios. Previous works fall into three major categories: 1) regularization-based, 2) architectural-based, and 3) replay-based. Regularization-based methods [3,9, 20, 46] aim to devise the regularizer that is to alleviate the forgetting problem by preventing the model parameters from drastic changes. On the other hand, architectural-based methods [21, 23, 25, 43] focus on proposing a new network architecture or developing adaptive designs for new tasks. Finally, replay-based methods [18,32,37,40] store the past data or utilize the generative model to replay the simplified version of the old task while learning the new tasks. Among these streams, our proposed method falls into the first category, *i.e.*, regularization-based. Yet, our goal is to highlight the significance of being discerned when distilling the prior knowledge to alleviate background shift.

2.2 Class-Incremental Semantic Segmentation

Semantic segmentation is a per-pixel classification task that aims to classify each pixel in the target image [4, 34, 35, 39]. Employing popular convolutional networks [5, 15, 17, 19, 24] and transformer-based architectures [7, 39, 42], we have observed the drastic improvements in semantic segmentation within the past few years. However, these methods are susceptible to incremental scenarios due to catastrophic forgetting. Tackling such vulnerability, the problem of Class-Incremental Semantic Segmentation has been proposed by ILT [28]. Since then, the need for CISS has been highlighted; MiB [1] utilized unbiased knowledge distillation to address the background shift problem, PLOP [11] used the multiscale distillation and pseudo-labeling strategy to maintain past knowledge, and Incrementer [36] proposed transformer-based framework and eased the learning process of new classes by distilling only the old-classes-relevant features. Our work resembles [1,11,36] in that we utilize prior knowledge to prevent the catastrophic forgetting of the old classes. However, our methods differ in that our goal is to effectively address background shift by separating objects (either of old and new classes) and the background.

3 Method

3.1 Problem Formulation

Class-Incremental Semantic Segmentation (CISS) aims for the model to learn new classes while retaining previously learned knowledge, using supervision only for the novel classes. Basically, the model for CISS is sequentially trained across multiple timesteps (t = 1, 2, ..., T), and its stability in preserving old classes is evaluated along with its plasticity in learning new classes at each step. At each time step t, a set of classes C^t newly appears and is distinct, ensuring that there is no overlap with the class sets from other steps formally denoted as $(C^i \cap C^j = \emptyset$ for all $i \neq j$, where $1 \leq i, j \leq T$). Then, the model learns new classes with data $D^t = \{(x_i^t \in \mathbb{R}^{H^I \times W^I \times 3}, y_i^t \in \mathbb{R}^{H^I \times W^I})\}_{i=1}^N$, where H^I and W^I denote the height and width of the image, respectively, and x_i^t, y_i^t represent the image, corresponding label map, respectively. Since only the supervision for novel classes is given at each step, the label map y^t consists of a set of novel classes c in C^t and the background class c_0 . Thus, the only available annotation in learning at t-th time step is whether each pixel belongs to the novel classes or not. In short, the pixels labeled as the background class in step t may contain the pixels that belong to the seen class sets $C^{1:t-1}$ or future class sets $C^{t+1:T}$.

3.2 Overview

The overview of our model is illustrated in Fig. 2. Employing Segmenter [39] with popular techniques for CISS, *i.e.*, Pseudo-Labeling (PL), Knowledge Distillation (KD), and the Background token Transfer (BT) [1], as the base framework,

 $\mathbf{5}$



Fig. 2: Overview of our background-class separation framework for CISS. Given a new class 'car' (grey) at step t (purple), a new class token is added to $[cls]^t$ where the token weight is initially duplicated from the background class. (Blue) Initially, the image is fed into the old model from the previous step t - 1 to generate the prediction S^{t-1} , (green) which is then used to calculate the object identifier O^t and the pseudo-label \tilde{y}^t . Subsequently, these are combined to produce the selective pseudo-label map \bar{y}^t to train the new model at step t. (yellow) Along with the selective pseudo-labeling, we further calibrate the degree of distilling the old knowledge based on patch-wise reliability, *i.e.*, prediction confidence, through adaptive feature distillation \mathcal{L}_{AFD} . Briefly, the degree for each patch to be distilled is derived from the combination of S^{t-1} and \bar{y}^t . (orange) For the semantic separation of the new class from the background at step t, background probability of old model is distilled into the new class probability of the new model. To further support the separation, an orthogonality loss between the new class and the background is implemented.

we propose multiple strategies to address the background shift in CISS. Despite the recognized their effectiveness of PL, KD, and BT, we claim that the background shift remains a significant unresolved challenge. First, the general protocol of PL is to regard the ambiguous pixels (lower than the threshold of PL) in the perspective of the old model as the background class which leads to a background shift towards old classes. Although KD is another method to distill the old model's knowledge, it forwards all features, including those that do not reflect the actual semantics, since conventional KD does not consider the old model's reliability. In continual scenarios, this leads to an error propagation problem (often triggering the background shift due to the large semantic coverage of the background). Finally, BT is to exploit the large semantic coverage in the background so that future classes are semantically included in the coverage of the background. Thus, it is shown to ease the initial training process of new classes when the background classifier weights are transferred to the ones of new classes. Yet, as the initial points of background and new classes are the same, it naturally makes the background shift towards the new classes. To address these, we propose selective pseudo-labeling (Sec. 3.3), adaptive feature distillation (Sec. 3.4), and

orthogonal objective (Sec. 3.5) to mitigate the background shift problem caused by aforementioned in CISS.

3.3 Selective Pseudo-Labeling Strategy

At every continual step in CISS, labels are provided only for the pixels belonging to the new classes. In other words, the pixels categorized as background may be part of either the old classes, including the semantic background class, or future classes. Because treating them as background could exacerbate catastrophic forgetting due to background shift towards old classes, many previous methods [1,11,36] have employed PL to address the issue. Generally, PL is conventional to assign each pixel with the specific object class as the pseudo-label if the prediction confidence of the old model for each pixel surpasses a predefined threshold τ and otherwise labeled as a background class [2,11,36,48]. Therefore, the pseudo-label map $\tilde{y} \in \mathbb{R}^{H^I \times W^I}$ [11] is described below:

$$\tilde{y}_{h,w}^{t} = \begin{cases} y_{h,w}^{t} & \text{if } (y_{h,w}^{t} \neq c_{0}) \\ \underset{c \in C^{1:t-1}}{\operatorname{argmax}} S_{h,w,c}^{t-1} & \text{if } (y_{h,w}^{t} = c_{0}) \land (\exists_{c \in C^{1:t-1}} S_{h,w,c}^{t-1} > \tau) \\ c_{0} & \text{otherwise} \end{cases}$$
(1)

where $S^{t-1} \in \mathbb{R}^{H^I \times W^I \times (|C^{1:t-1}|+1)}$ denotes the segmentation prediction of the old model after applying the softmax.

However, these methods highly rely on the old model's capability to generate pseudo-labels for all other unlabeled pixels, and if the predictions of the old model are incorrect, this can result in pixels of the old classes being mislabeled as the background in the new model's training. As a result, this aggravates the background shift towards old classes.

To address the above problem, we suggest being selective when distilling the knowledge of the old model via PL. This involves defining the object identifier O for each pixel, based on the old model's pixel-wise prediction confidence, to identify pixels distinct from the background class. We design the object identifier to be 1 when the sum of confidence scores over object classes is bigger than that of the background class. This is because the low confidence for the background is a strong signal that the pixel is highly likely to contain salient objects. Formally, the object identifier $O^t \in \mathbb{R}^{H^I \times W^I}$ at step t is defined as:

$$O_{h,w}^{t} = \begin{cases} 1 & \text{if } S_{h,w,c_{0}}^{t-1} < \sum_{c \in C^{1:t-1}} S_{h,w,c}^{t-1} \\ 0 & \text{otherwise} \end{cases}$$
(2)

Then, we calculate the selective pseudo-label map with the object identifier. we apply the selection strategy to detect background-misclassified pixels as below:

$$\bar{y}_{h,w}^{t} = \begin{cases} \tilde{y}_{h,w}^{t} & \text{if } (\tilde{y}_{h,w}^{t} \neq c_{0}) \\ c_{0} & \text{if } (\tilde{y}_{h,w}^{t} = c_{0}) \land (O_{h,w}^{t} = 0) \\ c_{\text{ignore}} & \text{otherwise} \end{cases}$$
(3)

where $\bar{y}^t \in \mathbb{R}^{H^I \times W^I}$ is the final pseudo-label map for step t and c_{ignore} indicates whether the pixel is excluded in training. By detecting the objects and preventing them from being labeled as the background class, we expect to mitigate the background shift. Consequently, our cross-entropy loss is expressed as:

$$\mathcal{L}_{\rm CE} = -\frac{1}{H^I W^I} \sum_{h=1}^{H^I} \sum_{w=1}^{W^I} \mathbb{1}_{\bar{y}_{h,w}^t \neq c_{\rm ignore}} \log(S_{h,w,\bar{y}_{h,w}^t}^t), \tag{4}$$

where $\mathbb{1}_{\alpha} = 1$ if α is true, and 0 otherwise.

3.4 Adaptive Feature Distillation

Along with the PL, KD is another tool to alleviate catastrophic forgetting [9,11, 12,18,36]. Despite its effectiveness in previous works, we point out that naïve KD methods without consideration of the old model's confidence can induce the background shift towards old classes.

To address the issue, we propose an adaptive feature distillation strategy, which accounts for the old model's extensive background knowledge and adaptively distills trustworthy features. Given the transformer backbone, let us denote the patch features at step t as $F^t \in \mathbb{R}^{H^F \times W^F \times D}$ where $H^F = \frac{H^I}{P}$ and $W^F = \frac{W^I}{P}$. Here, P indicates the patch size. Then, whereas the general feature distillation protocol is to teach F^t to mimic F^{t-1} , we aim to identify the unreliable patches that may trigger the error propagation. To achieve this, we utilize a patch reliability map M that measures the reliability of each patch. Specifically, we design the reliability map to retain the value of 1 in patches where pseudo-labels are assigned to one of the learned object classes with high confidence, signifying a strong learning signal. Conversely, we assign a value of 0 to patches corresponding spatially to c_{ignore} assigned in Eq. (3) or new classes; this is due to the old model's lack of knowledge about new classes and uncertainty in ambiguous regions. Additionally, for the patches pseudo-labeled as the background class, we assign a value corresponding to each patch's confidence score. Formally, the reliability map at step t, $M^t \in \mathbb{R}^{H^F \times W^F}$ is defined as:

$$M_{h,w}^{t} = \begin{cases} 1 & \text{if } (\hat{y}_{h,w}^{t} \in C^{1:t-1}) \\ 0 & \text{if } (\hat{y}_{h,w}^{t} \in C^{t}) \lor (\hat{y}_{h,w}^{t} = c_{\text{ignore}}) \\ \hat{S}_{h,w,c_{0}}^{t-1} & \text{if } (\hat{y}_{h,w}^{t} = c_{0}) \end{cases}$$
(5)

where $\hat{y} \in \mathbb{R}^{H^F \times W^F}$ and $\hat{S} \in \mathbb{R}^{H^F \times W^F}$ are the downsampled versions of the selective pseudo-label \bar{y} (defined in Eq. (3)) and segmentation prediction S to address the spatial mismatch with M at the feature level, and $M_{h,w}$ denotes the reliability score of the patch spatially located in (h, w), respectively. We follow previous works [31, 41, 44] to utilize the interpolation for downsampling.

In short, the old model's prediction for a patch located at (h, w) is considered to be reliable if the value of $M_{h,w}^t$ is close to 1. Consequently, we employ the reliability map as the weight map in distilling the old model's patch-wise

knowledge to prevent the new model from learning possibly incorrect supervision. With the reliability-based weighting scheme, our adaptive feature distillation loss is expressed as:

$$\mathcal{L}_{AFD} = \frac{1}{H^F W^F} \sum_{h=1}^{H^F} \sum_{w=1}^{W^F} M_{h,w}^t \left\| F_{h,w}^t - F_{h,w}^{t-1} \right\|^2.$$
(6)

Note that adaptive feature distillation is implemented after the decoding layers following existing works [1, 36].

3.5 Separating Background-Class after Background Weight Transfer

The background pixels predicted by the old model may include the novel classes at the current step. In this regard, MiB [1] eased the process of learning new classes by transferring the knowledge of the background to the new classes. Yet, this triggers the background shift towards the new classes, resulting in high correlations between the background and new classes as shown in Fig. 1 (d).

To tackle this, we adopt a label-guided knowledge distillation (LGKD) [45], distilling the old model's refined logit \bar{S}^{t-1} in which the background logit is transferred to the ones of new classes (when the pixel belongs to the new class). The refined logits $\bar{S}_{h,w,c}^{t-1}$ and the loss for LGKD are expressed as follows:

$$\bar{S}_{h,w,c}^{t-1} = \begin{cases} 0 & \text{if } \{(c=c_0) \land (\bar{y}_{h,w}^t \in C^t)\} \lor \{(c \in C^t) \land (\bar{y}_{h,w}^t \neq c)\} \\ S_{h,w,c_0}^{t-1} & \text{if } (c \in C^t) \land (\bar{y}_{h,w}^t = c) \\ S_{h,w,c_0}^{t-1} & \text{otherwise,} \end{cases}$$
(7)

$$\mathcal{L}_{\rm LGKD} = -\frac{1}{HW} \sum_{h=1}^{H} \sum_{w=1}^{W} \bar{S}_{h,w}^{t-1} \log S_{h,w}^{t}.$$
(8)

However, we argue that applying only the distillation objective does not sufficiently separate between the background and new classes (discussed in Tab. 5). Therefore, we propose an orthogonal loss to the correlations between the classifier weights of the new class and other classes with a novel view to directly train to separate the new classes from the old ones including the background class. Formally, the process of applying orthogonality in the classifier weights is expressed as:

$$\mathcal{L}_{\text{Ortho}} = \frac{1}{|C^t||C^{1:t}|} \sum_{c_i \in C^t} \sum_{c_j \in (C^{1:t} \cup \{c_0\})} \mathbb{1}_{i \neq j} |[cls]_{c_i}^t \odot \text{sg}([cls]_{c_j}^t)|, \tag{9}$$

where $[cls]_{c_i}$ denotes the *i*-th class weight and $sg(\cdot)$ stands for stop gradient operation [6, 16]. We use the stop gradient operation on the old class weights to prevent the significant modification of the previously learned embedding space. Consequently, the background shift towards the new classes can be alleviated as

the new class logit learns to deviate from the background class weight. For the overall loss formulation, we combine \mathcal{L}_{Ortho} with \mathcal{L}_{LGKD} as they are compatible with one another in that they share the same goal of separating the new class from the old ones but with different strategies. We express it as follows:

$$\mathcal{L}_{\text{Sep}} = \lambda_{\text{LGKD}} \mathcal{L}_{\text{LGKD}} + \lambda_{\text{Ortho}} \mathcal{L}_{\text{Ortho}}, \tag{10}$$

where λ_{LGKD} and λ_{Ortho} are coefficients for each objective.

3.6 Final Objective

Our proposed methods employ cross-entropy with selective pseudo-labeling \mathcal{L}_{CE} as described in Eq. (4). Also, we use the adaptive feature distillation loss \mathcal{L}_{AFD} as described in Eq. (6) and loss objective for background-class separation Eq. (10). To sum it up, our overall objective \mathcal{L} is expressed as:

$$\mathcal{L} = \mathcal{L}_{\rm CE} + \mathcal{L}_{\rm AFD} + \mathcal{L}_{\rm Sep}.$$
 (11)

4 Experiment

4.1 Experimental Details

Datasets. Following the experimental setting of [36], we evaluate our method on two public datasets: Pascal VOC [13] and ADE20k [49]. Pascal VOC contains 10,582 fully annotated images for training and 1,449 images for testing, over 20 foreground object classes. ADE20k has 20,210 training images and 2,000 testing images in 150 classes.

Experimental Protocols. We evaluate the performance of our method under two different CISS settings following MiB [1]: Disjoint and Overlapped. In both settings, labels are provided only for the class set C^{t} that newly appears in the current step t. However, data D^t is composed of pixels of classes belonging to the old or new class sets $(C^{1:t-1} \cup C^{t})$ in the *disjoint* setting. In the overlapping setting, pixels from both sets appear, and even those from future class sets $(C^{1:t-1} \cup C^t \cup C^{t+1:T})$. Among them, the overlapped setting is usually considered more challenging and more realistic in the CISS scenario. Additionally, we follow previous [1, 36, 48] to organize the class compositions with the number of continual steps. For example, the scenario named 15-1 (6 steps) indicates that we initially train the model on 15 classes and continually learn 1 additional class at every continual step. For evaluation metrics, we use the mean Intersection over Union (mIoU). Specifically, we compute the mIoU for the initial class set C^1 , incremented sets $C^{2:T}$, and all class sets $C^{1:T}$ (overall), respectively. Briefly, each mIoU score in order can be seen as a measure to evaluate the robustness to catastrophic forgetting, plasticity to new classes, and the balance between the two. The experimental results of our method are compared with those of previous CISS methods, such as EWC [20], ILT [28], MiB [1], SDR [29], PLOP [11], RECALL [26], REMIND [30], RCIL [47], SPPA [22], RBC [48], and INC [36]. 'Joint' denotes the results when all classes are trained at once (oracle scenario).

Table 1: Performance comparison on Pascal VOC under various scenarios. CNN and Transformer column indicates the type of the backbone network. † indicates the reproduced results from [11, 36, 48]

| | 19-1 (2 steps) | | | | | | 15-5 (2 steps) | | | | | 15-1 (6 steps) | | | | | | |
|------------------|-------------------|--------|------|--------|-------|------|-----------------|---------|------|------|------------|----------------|----------|-------|------|--------|-------|------|
| Method | D | isjoii | nt | Ove | erlap | ped | Γ | Disjoir | ıt | Ov | Overlapped | | Disjoint | | Ov | erlapp | ped | |
| | 1-19 | 20 | All | 1 - 19 | 20 | All | 1 - 15 | 16-20 | All | 1-15 | 16-20 | All | 1 - 15 | 16-20 | All | 1-15 | 16-20 | All |
| | CNN-based Methods | | | | | | | | | | | | | | | | | |
| EWC [†] | 23.2 | 16.0 | 22.9 | 26.9 | 14.0 | 26.3 | 26.7 | 37.7 | 29.4 | 24.3 | 35.5 | 27.1 | 0.3 | 4.3 | 1.3 | 0.3 | 4.3 | 1.3 |
| ILT^{\dagger} | 69.1 | 16.4 | 66.4 | 67.8 | 10.9 | 65.1 | 63.2 | 39.5 | 57.3 | 67.1 | 39.2 | 60.5 | 3.7 | 5.7 | 4.2 | 8.8 | 8.0 | 8.6 |
| MiB^{\dagger} | 69.6 | 25.6 | 67.4 | 71.4 | 23.6 | 69.2 | 71.8 | 43.3 | 64.7 | 76.4 | 50.0 | 70.1 | 46.2 | 12.9 | 37.9 | 34.2 | 13.5 | 29.3 |
| SDR^{\dagger} | 69.9 | 37.3 | 68.4 | 69.1 | 32.6 | 67.4 | 73.5 | 47.3 | 67.2 | 75.4 | 52.6 | 69.9 | 59.2 | 12.9 | 48.1 | 44.7 | 21.8 | 39.2 |
| $PLOP^{\dagger}$ | 75.4 | 38.9 | 73.6 | 75.4 | 37.4 | 73.5 | 71.0 | 42.8 | 64.3 | 75.7 | 51.7 | 70.1 | 57.9 | 13.7 | 46.5 | 65.1 | 21.1 | 54.6 |
| RECALL | 65.2 | 50.1 | 65.8 | 67.9 | 53.5 | 68.4 | 66.3 | 49.8 | 63.5 | 66.6 | 50.9 | 64.0 | 66.0 | 44.9 | 62.1 | 65.7 | 47.8 | 62.7 |
| REMIND | - | - | - | 76.5 | 32.3 | 74.4 | - | - | - | 76.1 | 50.7 | 70.1 | - | - | - | 68.3 | 27.2 | 58.5 |
| RCIL | - | - | - | - | - | - | 75.0 | 42.8 | 67.3 | 78.8 | 52.0 | 72.4 | 66.1 | 18.2 | 54.7 | 70.6 | 23.7 | 59.4 |
| SPPA | 75.5 | 38.0 | 73.7 | 76.5 | 36.2 | 74.6 | 75.3 | 48.7 | 69.0 | 78.1 | 52.9 | 72.1 | 59.6 | 15.6 | 49.1 | 66.2 | 23.3 | 56.0 |
| RBC^{\dagger} | 76.4 | 45.8 | 75.0 | 77.3 | 55.6 | 76.2 | 75.1 | 49.7 | 69.9 | 76.6 | 52.8 | 70.9 | 61.7 | 19.5 | 51.6 | 69.5 | 38.4 | 62.1 |
| Joint | 77.4 | 78.0 | 77.4 | 77.4 | 78.0 | 77.4 | 79.1 | 72.6 | 77.4 | 79.1 | 72.6 | 77.4 | 79.1 | 72.6 | 77.4 | 79.1 | 72.6 | 77.4 |
| | | | | | | Tra | nsfor | mer- | base | d M | ethod | ls | | | | | | |
| MiB† | 80.6 | 45.2 | 79.6 | 79.9 | 47.7 | 79.1 | 75.0 | 59.9 | 72.3 | 78.6 | 63.1 | 75.6 | 66.7 | 26.3 | 58.3 | 72.6 | 23.1 | 61.7 |
| RBC^{\dagger} | 80.9 | 42.1 | 79.7 | 80.2 | 38.8 | 79.0 | 77.7 | 59.1 | 74.0 | 78.9 | 62.0 | 75.5 | 69.0 | 28.4 | 60.5 | 75.9 | 40.2 | 68.2 |
| INC^{\dagger} | 82.4 | 64.2 | 82.2 | 82.5 | 61.0 | 82.1 | 81.6 | 62.2 | 77.6 | 82.5 | 69.2 | 79.9 | 81.4 | 57.1 | 76.2 | 79.6 | 59.6 | 75.6 |
| Ours | 82.8 | 69.3 | 82.8 | 83.3 | 72.0 | 83.3 | 81.9 | 67.2 | 79.0 | 84.1 | 76.0 | 82.6 | 81.5 | 64.7 | 78.1 | 82.6 | 72.2 | 80.6 |
| Joint | 83.9 | 81.6 | 83.8 | 83.9 | 81.6 | 83.8 | 84.8 | 80.7 | 83.8 | 84.8 | 80.7 | 83.8 | 84.8 | 80.7 | 83.8 | 84.8 | 80.7 | 83.8 |

Implementation Details. We utilize the ImageNet-pretrained [8] vision transformer VIT-B/16 [10] as our encoder and two transformer layers as the decoder [39] which processes cropped input image of resolution 512×512 . Our initial learning rate is set to 1×10^{-3} with SGD [33] optimizer. Momentum and the weight decay parameters for SGD are set to 0.9 and 1×10^{-5} , respectively, and the pseudolabeling threshold τ is set to 0.7 [2] for all experiments. For Pascal VOC 2012, the training proceeds with the batch size of 16 for 32 epochs, with the learning rate of 1×10^{-4} in incremental steps and the weight for output distillation loss λ_{LGKD} is set to 25. On ADE20k, we train the model with the batch size of 8 for 64 epochs, adjusting the learning rate to 5×10^{-4} and with λ_{LGKD} set to 50 in continual steps. Finally, the λ_{Ortho} is adaptively set to $\frac{|C^t|}{|C^{1:t}|}$ for all experiments.

4.2 Comparison with the State-of-the-arts

Pascal VOC. Experimental results on Pascal VOC with different incremental learning settings are shown in Tab. 1. As observed, our method clearly outperforms previous state-of-the-art (SOTA) methods in all incremental scenarios. Particularly, we highlight the remarkable performance improvement in overlapped settings and multi-step scenarios where more objects appear as the background and the error in previous distillation methods may be propagated through multiple steps. Specifically in Tab. 1, our proposed method exceeds the previous SOTA model in a single continual step (15-5) setting by 1.4%p (disjoint) and 2.7%p (overlapped). Furthermore, our results achieve the SOTA with a large margin, up to 1.9%p (disjoint) and 5.0%p (overlapped), in the 15-1 multi-continual

| Mathod | 100-50 (2 steps) | | | 50- | 50-50 (3 steps) | | | 100-10 (6 steps) | | | 100-5 (11 steps) | | |
|-------------------|---------------------------|-----------|------|------|-----------------|-------------|-------|------------------|-------------|-------|------------------|------|--|
| Method | 1-100 | 101 - 150 | All | 1-50 | 51 - 150 | All | 1-100 | 101 - 150 | All | 1-100 | 101 - 150 | All | |
| CNN-based Methods | | | | | | | | | | | | | |
| MiB | 40.5 | 17.2 | 32.8 | 45.5 | 21.0 | 29.3 | 38.2 | 11.1 | 29.2 | 36.0 | 5.7 | 26.0 | |
| SDR | 37.4 | 24.8 | 33.2 | 40.9 | 23.8 | 29.5 | 28.9 | 7.4 | 21.7 | - | - | - | |
| PLOP | 41.7 | 15.4 | 33.0 | 47.8 | 21.6 | 30.4 | 39.4 | 13.6 | 30.9 | 39.1 | 7.8 | 28.8 | |
| REMINDER | 41.6 | 19.2 | 34.1 | 47.1 | 20.4 | 29.4 | 39.0 | 21.3 | 33.1 | - | - | - | |
| RCIL | 42.3 | 18.8 | 34.5 | 48.3 | 25.0 | 32.5 | 39.3 | 17.6 | 32.0 | 38.5 | 11.5 | 29.6 | |
| SPPA | 42.9 | 19.9 | 35.2 | 49.8 | 23.9 | 32.5 | 41.0 | 12.5 | 31.5 | - | - | - | |
| RBC | 42.9 | 21.5 | 35.8 | 49.6 | 26.3 | 34.2 | 39.0 | 21.7 | 33.3 | - | - | - | |
| Joint | 43.9 | 27.2 | 38.3 | 50.9 | 32.1 | 38.3 | 43.9 | 27.2 | 38.3 | 43.9 | 27.2 | 38.3 | |
| | Transformer-based Methods | | | | | | | | | | | | |
| MiB | 46.4 | 35.0 | 42.6 | 52.2 | 35.6 | 41.1 | 43.0 | 30.8 | 38.9 | 40.2 | 26.6 | 35.7 | |
| INC | 49.4 | 35.6 | 44.8 | 56.2 | 37.8 | 43.9 | 48.5 | 34.6 | 43.9 | 46.9 | 31.3 | 41.7 | |
| Ours | 50.0 | 37.1 | 45.7 | 57.0 | 39.7 | 45.4 | 49.0 | 35.4 | 44.5 | 48.0 | 32.3 | 42.8 | |
| Joint | 50.1 | 41.4 | 47.2 | 57.6 | 42.0 | 47.2 | 50.1 | 41.4 | 47.2 | 50.1 | 41.4 | 47.2 | |

Table 2: Performance comparison on ADE20k across various scenarios in *overlapped* setting. CNN and Transformer column indicates the type of the backbone network.

benchmarking scenario. For additional experimental results in the 10-1 scenario and 5-3 scenario, we refer to the Appendix A.1. Such large improvements in multicontinual scenarios indicate the significance of error propagation in incremental learning caused by the background shift and how well our reliable distillation techniques, i.e., selective pseudo-labeling and adaptive feature distillation, address the background shift towards old classes. Furthermore, relatively large improvements in overlapped scenarios compared to disjoint scenarios are due to the inclusion of future class knowledge in the background in overlapped settings. Therefore, we believe that leveraging the background knowledge to initialize the new class token and decoupling the shared knowledge benefits the most in such settings.

ADE20k. Our experiments in the overlapped setting on ADE20k dataset are reported in Tab. 2. As demonstrated, we achieve an average improvement of 1.0%p over the state-of-the-art methods in the 100-50, 50-50, 100-10, and 100-5 scenarios. In particular, as the scenarios become more challenging, *i.e.*, more time steps, our proposed method showcases its benefit in retaining its knowledge (stability) beyond the previous SOTA. Specifically, the margins between the performances for the classes in the initial step are 0.6%p in 100-50 (2 steps), 0.8%p in 50-50 (3 steps) and 1.1%p in 100-5 (11 steps) while also exceeding the performances for measuring the new class learning. Overall, these results verify the effectiveness of ours in both the stability and the plasticity perspective which demonstrates the strength in CISS.

4.3 Ablation Studies

Component Analysis. We investigate the effectiveness of each component in the overlapped setting of Pascal VOC 15-1 and ADE20k 100-50, as observed in

| | | | | Р | ascal VO | С | ADE20k | | | |
|-----|-----------------------|--------------|------------------------------|----------------|----------|------|------------------|-----------|------|--|
| | SPL | AFD | $\mathcal{L}_{\mathrm{Sep}}$ | 15-1 (6 steps) | | | 100-50 (2 steps) | | | |
| | | | | 1 - 15 | 16 - 20 | All | 1-100 | 101 - 150 | All | |
| (a) | - | - | - | 81.1 | 49.8 | 74.2 | 47.3 | 35.0 | 43.2 | |
| (b) | ✓ | - | - | 83.5 | 63.0 | 79.1 | 50.0 | 35.5 | 45.2 | |
| (c) | - | \checkmark | - | 82.5 | 57.3 | 77.0 | 47.4 | 36.4 | 43.7 | |
| (d) | - | - | \checkmark | 80.8 | 69.6 | 78.7 | 47.1 | 37.3 | 43.8 | |
| (e) | ✓ | \checkmark | - | 83.8 | 65.3 | 79.9 | 50.0 | 35.7 | 45.2 | |
| (f) | ✓ | - | \checkmark | 82.2 | 70.2 | 79.9 | 49.7 | 36.9 | 45.5 | |
| (g) | - | \checkmark | \checkmark | 81.2 | 70.8 | 79.3 | 47.4 | 37.6 | 44.1 | |
| (h) | \checkmark | \checkmark | \checkmark | 82.6 | 72.2 | 80.6 | 50.0 | 37.1 | 45.7 | |

Table 3: Ablation study for each component on Pascal VOC 15-1 and ADE20k 100-50 in overlapped setting. SPL, AFD, \mathcal{L}_{Sep} denote selective pseudo-labeling, adaptive feature distillation, and separating objective, respectively.

Table 4: Ablation study for selective pseudo-labeling filter effectiveness on Pascal VOC 15-1 in overlapped setting.

| | Step 1 | Step 2 | Step 3 | Step 4 | Step 5 | Avg. |
|---------------------|--------|----------|--------|--------|----------|-------|
| Ignored Label Ratio | 38.21 | 31.29 | 34.59 | 34.32 | 47.52 | 37.18 |

Tab. 3. Compared to the baseline (a), rows (b) to (c) show the benefits of SPL and AFD discussed in Sec. 3.3 and Sec. 3.4. Specifically, the performances for the classes in the initial step (1-15) are enhanced by 2.4%p and 1.4%p in the 15-1 scenario. We owe this improvement to their capability of preventing the old class pixels from being labeled as the background. Surprisingly, we also find that SPL and AFD also benefit in classes that appear in later steps, SPL particularly showing a 13.2%p increase. This is because selective and adaptive strategies can also suppress the learning of future class objects with the background label if they appear in earlier steps. On the other hand, our objective \mathcal{L}_{Sep} , explained in Sec. 3.5, enhances the model's plasticity (result for new classes in row (d) is 4.5% p higher than the baseline (a) in Pascal VOC 15-1 setting) by directly separating the embeddings of the new classes from the background. Consequently, the results in (h) with all components combined demonstrate not only the effectiveness of each component but also the compatibility to one another.

Effectiveness of Selective Strategy

within Pseudo-labeling. We inspect Table 5: Ablation study evaluating the imare detected with our selective strategy. ADE20k 100-50 overlapped setting. In Tab. 4, we present the percentage of the old class pixels identified by selective strategy among mispredicted pixels by naïve pseudo-labeling for each step. Briefly, our selective method filters out

how much the incorrectly labeled pixels pact of the Orthogonal Objective in the

| | C | C | | Similar | ity |
|-----|-----------------------|------------------------------|--------------|--------------|--------------------|
| | LGKD | $\mathcal{L}_{\text{Ortho}}$ | (C^t, C^t) | (C^t, c_0) | $(C^t, C^{1:t-1})$ |
| (a) | - | - | 91.24 | 91.46 | 3.03 |
| (b) | ✓ | - | 90.15 | 90.03 | 2.93 |
| (c) | \checkmark | \checkmark | 52.68 | 53.86 | 1.58 |

about 37% mispredicted old class pixels as the background on average. In other

MBS 13



Fig. 3: Comparison of qualitative results on the 15-1 protocol of the Pascal VOC between the baseline, MiB, and ours.

words, we claim that being selective within pseudo-labeling can alleviate the background shift towards old classes up to 37% which demonstrates the strengths of our method.

Class Token Similarity. To further validate the effect of separating objective, we present the class token similarities, that more directly represent the degree of separation between classes, in Tab. 5. As shown in row (a), when only the background transfer is used, the similarities for new class pairs (C^t, C^t) and new-to-background (C^t, c_0) are exceptionally high at 91.24 and 91.46, respectively. Even if the label-guided knowledge distillation L_{LGKD} [45] is adopted, it shows a marginal separation as like 90.15 and 90.03, individually. On the other hand, when our separating objective L_{Ortho} is utilized together, the similarity scores extremely decrease to 52.68 and 53.86, respectively. Moreover, the similarity between new-to-old $(C^t, C^{1:t-1})$ decreased to 1.58. This proves that the orthogonal objective possessing a direct separating purpose is effective in isolating the new classes from the background class and mitigates the background shift towards new classes.



Fig. 4: Comparisons between the Pseudo-Labeling map (PL map) and the Selective Pseudo-Labeling map (SPL map). Both of them are yielded from our method.

4.4 Qualitative Analysis

We visualize the qualitative results in the Pascal VOC 15-1 *overlapped* setting in Fig. 3. Across examples, our proposed methods show its benefits in terms of robustness to forgetting and plasticity to learning new classes compared to both the baseline as mentioned in Sec. 4.3 and MiB. For example, whereas the baseline and MiB struggle to retain the old knowledge (a-step 5, and b-step 6), our method successively captures the old classes. Furthermore, (b-step 2) display difficulty in learning new classes in the baseline and MiB, respectively, while our methods are shown to learn the new class 'plant' precisely. These results show that our method possesses robustness to background shifts towards old and new classes. The additional qualitative results can be found in the Appendix.

To show the role of the object identifier O, we visualize the naïve pseudolabeling and selective pseudo-labeling maps in Fig. 4. The naïve pseudo-labeling map shows its vulnerability to the imprecision of the old model, labeling the ambiguous objects as the background. In contrast, our selective pseudo-labeling alleviates the background shift by detecting and filtering out those objects.

5 Conclusion

In this paper, we propose a background-class separation framework to address the background shift issue towards the old and new classes in CISS. To alleviate the prediction bias towards the background in old classes, we first introduced selective pseudo labeling and adaptive feature distillation that are to scale down the degree of distillation in parts where the false knowledge might be distilled. Subsequently, to facilitate the decoupling of background and new classes while leveraging the benefits of the background weight transfer strategy, we encourage further separation between the background and new classes. This is achieved through a novel orthogonal objective, complemented by label-guided output distillation. With all our components combined, our proposed method achieved state-of-the-art performances in terms of both robustness to forgetting and the plasticity of new class learning. Acknowledgements. This work was supported in part by MSIT&KNPA/KIPoT (Police Lab 2.0, No. 210121M06), MSIT/IITP (No. 2022-0-00680, 2019-0-00421, 2020-0-01821, RS-2024-00437102), and SEMES-SKKU collaboration funded by SEMES. Gilhan Park acknowledges support from the Hyundai Motor Chung Mong-Koo Foundation

References

- Cermelli, F., Mancini, M., Bulo, S.R., Ricci, E., Caputo, B.: Modeling the background for incremental learning in semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9233–9242 (2020)
- Cha, S., Yoo, Y., Moon, T., et al.: Ssul: Semantic segmentation with unknown label for exemplar-based class-incremental learning. Advances in neural information processing systems 34, 10919–10930 (2021)
- Chaudhry, A., Dokania, P.K., Ajanthan, T., Torr, P.H.: Riemannian walk for incremental learning: Understanding forgetting and intransigence. In: Proceedings of the European conference on computer vision (ECCV). pp. 532–547 (2018)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. arXiv preprint arXiv:1412.7062 (2014)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence 40(4), 834–848 (2017)
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 15750–15758 (2021)
- Cheng, B., Schwing, A., Kirillov, A.: Per-pixel classification is not all you need for semantic segmentation. Advances in Neural Information Processing Systems 34, 17864–17875 (2021)
- Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 248–255. Ieee (2009)
- Dhar, P., Singh, R.V., Peng, K.C., Wu, Z., Chellappa, R.: Learning without memorizing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5138–5146 (2019)
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: International Conference on Learning Representations (2021), https://openreview. net/forum?id=YicbFdNTTy
- Douillard, A., Chen, Y., Dapogny, A., Cord, M.: Plop: Learning without forgetting for continual semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4040–4050 (2021)
- Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E.: Podnet: Pooled outputs distillation for small-tasks incremental learning. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX 16. pp. 86–102. Springer (2020)

- 16 G. Park et al.
- Everingham, M., Eslami, S.A., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes challenge: A retrospective. International journal of computer vision 111, 98–136 (2015)
- French, R.M.: Catastrophic forgetting in connectionist networks. Trends in cognitive sciences 3(4), 128–135 (1999)
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3146–3154 (2019)
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems 33, 21271–21284 (2020)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
- Hou, S., Pan, X., Loy, C.C., Wang, Z., Lin, D.: Learning a unified classifier incrementally via rebalancing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 831–839 (2019)
- Huang, Z., Wang, X., Huang, L., Huang, C., Wei, Y., Liu, W.: Ccnet: Criss-cross attention for semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 603–612 (2019)
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al.: Overcoming catastrophic forgetting in neural networks. Proceedings of the national academy of sciences 114(13), 3521–3526 (2017)
- Li, X., Zhou, Y., Wu, T., Socher, R., Xiong, C.: Learn to grow: A continual structure learning framework for overcoming catastrophic forgetting. In: International Conference on Machine Learning. pp. 3925–3934. PMLR (2019)
- Lin, Z., Wang, Z., Zhang, Y.: Continual semantic segmentation via structure preserving and projected feature alignment. In: European Conference on Computer Vision. pp. 345–361. Springer (2022)
- Liu, Y., Schiele, B., Sun, Q.: Adaptive aggregation networks for class-incremental learning. In: Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition. pp. 2544–2553 (2021)
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
- Mallya, A., Lazebnik, S.: Packnet: Adding multiple tasks to a single network by iterative pruning. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 7765–7773 (2018)
- Maracani, A., Michieli, U., Toldo, M., Zanuttigh, P.: Recall: Replay-based continual learning in semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7026–7035 (2021)
- McCloskey, M., Cohen, N.J.: Catastrophic interference in connectionist networks: The sequential learning problem. In: Psychology of learning and motivation, vol. 24, pp. 109–165. Elsevier (1989)
- Michieli, U., Zanuttigh, P.: Incremental learning techniques for semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision workshops. pp. 0–0 (2019)
- Michieli, U., Zanuttigh, P.: Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1114–1124 (2021)

- Phan, M.H., Phung, S.L., Tran-Thanh, L., Bouzerdoum, A., et al.: Class similarity weighted knowledge distillation for continual semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16866–16875 (2022)
- Qiu, Y., Shen, Y., Sun, Z., Zheng, Y., Chang, X., Zheng, W., Wang, R.: Sats: Self-attention transfer for continual semantic segmentation. Pattern Recognition 138, 109383 (2023)
- Rebuffi, S.A., Kolesnikov, A., Sperl, G., Lampert, C.H.: icarl: Incremental classifier and representation learning. In: Proceedings of the IEEE conference on Computer Vision and Pattern Recognition. pp. 2001–2010 (2017)
- Robbins, H., Monro, S.: A stochastic approximation method. The annals of mathematical statistics pp. 400–407 (1951)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)
- Seong, H.S., Moon, W., Lee, S., Heo, J.P.: Leveraging hidden positives for unsupervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19540–19549 (2023)
- Shang, C., Li, H., Meng, F., Wu, Q., Qiu, H., Wang, L.: Incrementer: Transformer for class-incremental semantic segmentation with knowledge distillation focusing on old class. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7214–7224 (2023)
- Shin, H., Lee, J.K., Kim, J., Kim, J.: Continual learning with deep generative replay. Advances in neural information processing systems **30** (2017)
- Siam, M., Elkerdawy, S., Jagersand, M., Yogamani, S.: Deep semantic segmentation for automated driving: Taxonomy, roadmap and challenges. In: 2017 IEEE 20th international conference on intelligent transportation systems (ITSC). pp. 1–8. IEEE (2017)
- Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 7262–7272 (2021)
- Tiwari, R., Killamsetty, K., Iyer, R., Shenoy, P.: Gcr: Gradient coreset based replay buffer selection for continual learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 99–108 (2022)
- 41. Wang, W., Zhou, T., Yu, F., Dai, J., Konukoglu, E., Van Gool, L.: Exploring crossimage pixel contrast for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7303–7313 (2021)
- 42. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems 34, 12077–12090 (2021)
- 43. Yan, S., Xie, J., He, X.: Der: Dynamically expandable representation for class incremental learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3014–3023 (2021)
- Yang, G., Fini, E., Xu, D., Rota, P., Ding, M., Nabi, M., Alameda-Pineda, X., Ricci, E.: Uncertainty-aware contrastive distillation for incremental semantic segmentation. IEEE Transactions on Pattern Analysis and Machine Intelligence 45(2), 2567–2581 (2022)
- 45. Yang, Z., Li, R., Ling, E., Zhang, C., Wang, Y., Huang, D., Ma, K.T., Hur, M., Lin, G.: Label-guided knowledge distillation for continual semantic segmentation

on 2d images and 3d point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 18601–18612 (2023)

- Zenke, F., Poole, B., Ganguli, S.: Continual learning through synaptic intelligence. In: International conference on machine learning. pp. 3987–3995. PMLR (2017)
- Zhang, C.B., Xiao, J.W., Liu, X., Chen, Y.C., Cheng, M.M.: Representation compensation networks for continual semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7053–7064 (2022)
- Zhao, H., Yang, F., Fu, X., Li, X.: Rbc: Rectifying the biased context in continual semantic segmentation. In: European Conference on Computer Vision. pp. 55–72. Springer (2022)
- 49. Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralba, A.: Scene parsing through ade20k dataset. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 633–641 (2017)