

Adaptive Prototype Replay for Class Incremental Semantic Segmentation

Guilin Zhu¹, Dongyue Wu¹, Changxin Gao^{1*}, Runmin Wang², Weidong Yang¹, Nong Sang¹

¹School of Artificial Intelligence and Automation, Huazhong University of Science and Technology

²College of Information Science and Engineering, Hunan Normal University

{gzhu, dongyue_wu, cgao, Yangwd, nsang}@hust.edu.cn, runminwang@hunnu.edu.cn

Abstract

Class incremental semantic segmentation (CISS) aims to segment new classes during continual steps while preventing the forgetting of old knowledge. Existing methods alleviate catastrophic forgetting by replaying distributions of previously learned classes using stored prototypes or features. However, they overlook a critical issue: in CISS, the representation of class knowledge is updated continuously through incremental learning, whereas prototype replay methods maintain fixed prototypes. This mismatch between updated representation and fixed prototypes limits the effectiveness of the prototype replay strategy. To address this issue, we propose the **Adaptive prototype replay (Adapter)** for CISS in this paper. Adapter comprises an adaptive deviation compensation (ADC) strategy and an uncertainty-aware constraint (UAC) loss. Specifically, the ADC strategy dynamically updates the stored prototypes based on the estimated representation shift distance to match the updated representation of old class. The UAC loss reduces prediction uncertainty, aggregating discriminative features to aid in generating compact prototypes. Additionally, we introduce a compensation-based prototype similarity discriminative (CPD) loss to ensure adequate differentiation between similar prototypes, thereby enhancing the efficiency of the adaptive prototype replay strategy. Extensive experiments on Pascal VOC and ADE20K datasets demonstrate that Adapter achieves state-of-the-art results and proves effective across various CISS tasks, particularly in challenging multi-step scenarios.

Code — <https://github.com/zhu-gl-ux/Adapter>

Introduction

The class incremental semantic segmentation (CISS) models learn all classes through multiple steps. During each step, models focus on different classes with corresponding labels. When completing the all steps, the trained models are expected to grasp not only with the original classes but also with the newly introduced ones. A common challenge observed in other incremental visual tasks (Rebuffi et al. 2017; Shmelkov, Schmid, and Alahari 2017) is *catastrophic forgetting* (McCloskey and Cohen 1989). This occurs when the model learns a new task, the network weights are fine-tuned

*Corresponding author.

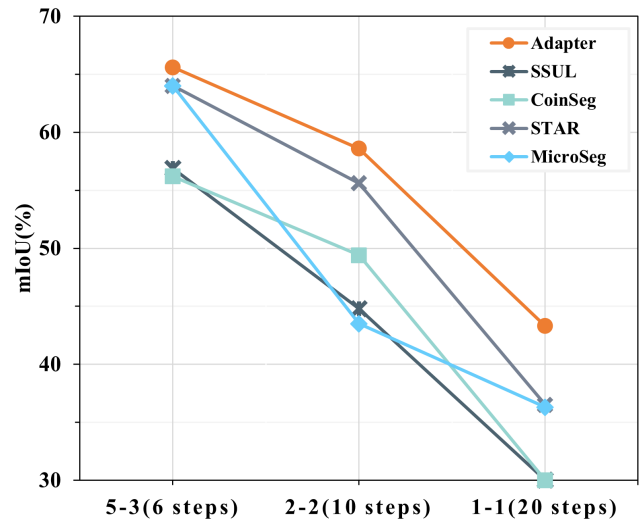


Figure 1: Performance comparison of prior works on the Pascal VOC 2012 challenging multi-step *overlapped* scenarios, where contains the small classes in the initial step and long-term incremental steps.

to accommodate the new data while the old shared weights are disrupted.

Intuitive methods (Cha et al. 2021; Maracani et al. 2021) attempted to alleviate catastrophic forgetting by replaying a few past images, but this suffer from storage burdens and privacy concerns. Previous CISS methods (Michieli and Zanuttigh 2019; Yang et al. 2022) utilized knowledge distillation (KD) to reduce catastrophic forgetting. However, these methods overlook the fact that the proportion of pixels belonging to new classes is higher than that of previously learned classes during the incremental steps, which leads to bias in the classifiers. Recent work (Chen et al. 2023) properly addressed this problem by replaying the distributions of old-class in the new-class classifiers, based on the efficient storage of representative prototypes and necessary statistics. However, the distributions replayed by static prototypes only represent the representation of previous models, neglecting the representation deviation between incremental models due to the updated shared weights.

In this paper, we propose a method called Adapter (**Adaptive prototype replay**) to tackle the issue of representation deviation in prototype replay based CISS methods. In the incremental scenario, as new classes are acquired, the models may inevitably generate representation deviations towards learned classes due to constantly updated shared weights. Therefore, Adapter avoids replaying the old-class feature distributions with fixed old-class prototypes in all the steps. Instead, Adapter exploits an adaptive deviation compensation (ADC) strategy to dynamically update representative prototypes of old classes based on the estimated representation shift distance between previous and current models in a training-free manner. Moreover, we propose an uncertainty-aware constraint (UAC) loss, which compresses the representation of each class to a compact space, thereby promoting the model to generate compact class prototypes. Additionally, to facilitate the prototype replay strategy to be carried out effectively in incremental steps, it is necessary to ensure that the representation of each class is distinguishable. Therefore, we introduce a compensation-based prototype similarity discriminative (CPD) loss. During the training process, the CPD loss repels the distance between compensatory old-class prototypes and batch-level new-class feature centers, and maintains the difference among positive new-class feature centers and negative ones.

Our main contributions are summarized as follows: 1) We propose a method called Adapter for CISS, which adopts an adaptive deviation compensation (ADC) strategy. The ADC strategy dynamically updates the stored prototypes, thereby eliminating the representation deviation between incremental models. 2) We design an uncertainty-aware constraint (UAC) loss to decrease the uncertainty predictions and aggregate class representations, facilitating the generation of compact prototypes. 3) We introduce a compensation-based prototype similarity discriminative (CPD) loss to enhance the discrimination of each class prototype and improve the efficiency of the adaptive prototype replay strategy. 4) We demonstrate the effectiveness of our Adapter through extensive experiments on two public CISS benchmarks and achieve state-of-the-art performance, particularly in more challenging and realistic multi-step scenarios (Fig. 1).

Related Work

Continual Learning

Continuous learning (CL), also referred to as incremental learning (Masana et al. 2022), seeks to acquire new knowledge from continuous data streams while minimizing the risk of forgetting previously learned concepts. In recent years, research on continuous learning has concentrated on image classification, which can be broadly categorized into regularized, architectural, and replay methods. Regularization methods utilize regularization or distillation techniques (Gou et al. 2021) to address catastrophic forgetting at both the parameter (Kirkpatrick et al. 2017; Zenke, Poole, and Ganguli 2017) and function levels (Li and Hoiem 2017; Douillard et al. 2020). Architectural methods (Jung et al. 2020; Yoon et al. 2018) involve selecting isolated parameter subspace for different tasks. Replay methods (Lopez-Paz

and Ranzato 2017; Shin et al. 2017) has focused on enhancing replay effectiveness. For instance, (Zhao et al. 2021) utilizes low-fidelity auxiliary samples of old-class knowledge transfer to enhance memory efficiency. (Yu et al. 2020) estimates feature drift caused by learning new tasks and compensates for the prototypes of previous tasks using an embedding network to enhance performance in image classification. Our method differs from prior approaches in two key aspects. First, we update the prototypes used for replay incrementally by estimating and compensating for representation deviation between incremental models without additional training. Second, while prior methods generally address task-agnostic incremental image classification, our approach specifically targets model representation deviation in class-incremental semantic segmentation with uncertain pseudo-predictions.

Class Incremental Semantic Segmentation

ILT (Michieli and Zanuttigh 2019) extends incremental learning to semantic segmentation by implementing distillation techniques to the output layer and intermediate features. MiB (Cermelli et al. 2020) addresses the unique issue in CISS called *background shift* by modeling the background semantic during learning steps. PLOP (Douillard et al. 2021) proposes a multi-scale local distillation scheme to alleviate catastrophic forgetting. SSUL (Cha et al. 2021) and MicroSeg (Zhang et al. 2022) define unknown classes within the background class to enhance plasticity, while replacing cross-entropy (CE) loss with binary cross-entropy (BCE) loss to mitigate catastrophic forgetting. DKD (Baek et al. 2022) proposes the decoupled knowledge distillation to improve the performance. IDEC (Zhao, Yuan, and Shi 2023) and CoinSeg (Zhang et al. 2023) utilize contrast learning to obtain discriminative representation. In addition, several methods (Oh, Baek, and Ham 2022; Chen et al. 2023) explored the efficiency of no-exemplar memory replay. For instance, STAR (Chen et al. 2023) stored a compact prototype of each class and the necessary statistics used to replay the feature distributions in subsequent steps, addressing classifier bias. However, failing to account for the representation deviation generated by the incremental models on old-class knowledge limits the effectiveness of prototype replay. In contrast, our method replays old-class distributions based on considering the representation deviation of incremental models, thereby providing a distinct advantage.

Method

Problem Definition

In CISS, the model training process involves a series of incremental steps, as $t = 1, \dots, T$. At each incremental step t , a sub-dataset D^t is used to learn the classes C^t (full classes as C) for the current step. The sub-dataset D^t consists of N^t pairs set $\{x_i^t, y_i^t\}_{i=1}^{N^t}$, where $x_i^t \in \mathbb{R}^{H \times W \times 3}$ denotes the image with size of $H \times W$ and $y_i^t \in \mathbb{R}^{H \times W}$ denotes the corresponding ground-truth mask. It is worth noting that the labels $\{y_i^t\}_{i=1}^{N^t}$ within sub-dataset D^t only consider the current learning classes pixels as foreground classes C^t , while other classes are treated as background.

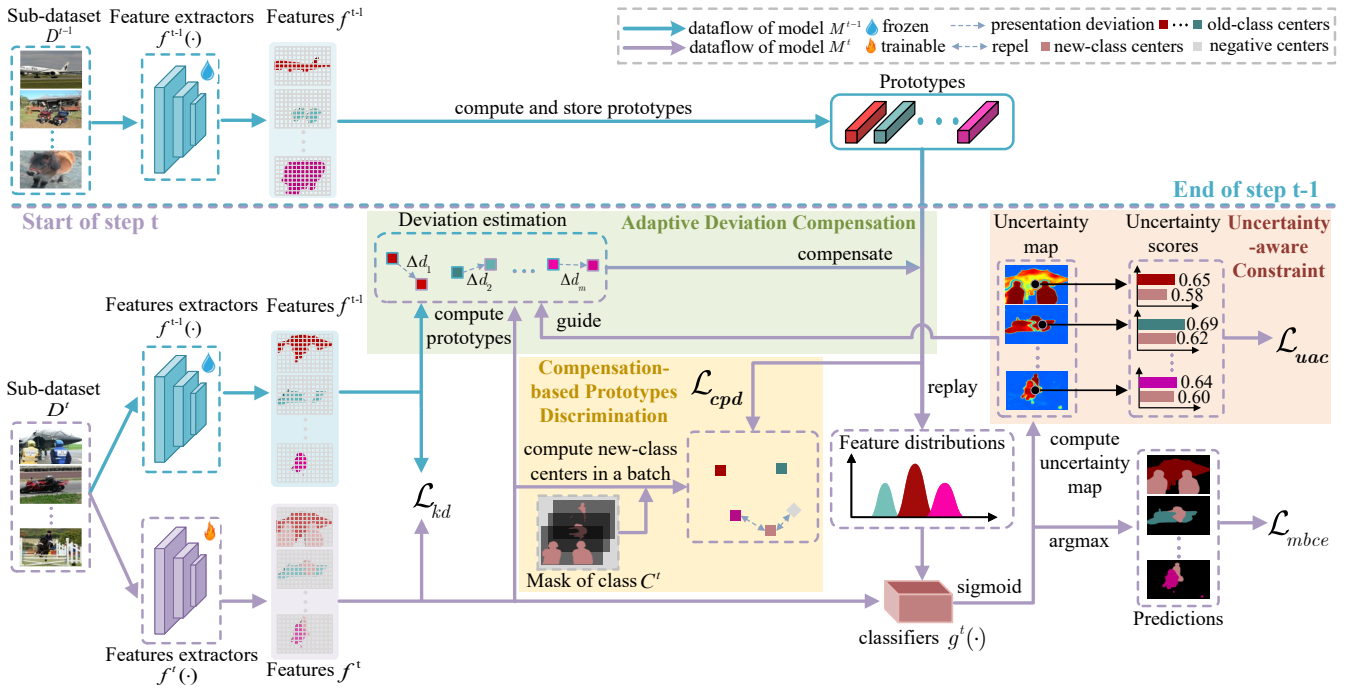


Figure 2: Overview of the proposed Adapter. After training of step $t - 1$, old-class prototypes are computed and saved. At the current step t , representation deviations towards old-class are estimated and prototypes are updated with the ADC strategy. Old-class feature distributions are replayed with updated prototypes for classifiers $g^t(\cdot)$. The UAC loss enhances the consistency of the same class by aggregating the representation of each class with uncertainty-aware loss. The CPD loss facilitates discrimination between new-class features and updated old-class prototypes, improving the efficiency of prototype replay.

We define the model in step t as M^t , which consists of feature extractors $f^t(\cdot)$ and classifiers $g^t(\cdot)$. For an input image \mathbf{x}_i^t , the output of model is $\hat{\mathbf{y}}_i^t \in \mathbb{R}^{H \times W}$ that computed by $\hat{\mathbf{y}}_i^t = M^t(\mathbf{x}_i^t) = g^t \circ f^t(\mathbf{x}_i^t)$. We denote extracted features as $f_i^t \in \mathbb{R}^{h \times w \times d}$ computed by $f_i^t = f^t(\mathbf{x}_i^t)$, where d represents the dimension of features. After completing all incremental steps, the model is desired to have considerable performance not only on old classes but also on new ones.

Overview

We show the overview of our proposed Adapter in Fig 2. After the training process of step $t - 1$, the model inference all training samples in sub-dataset D^{t-1} to obtain and store class prototypes in C^{t-1} and essential statistics. At early training epochs of steps t , the saved old-class prototypes are used for replaying feature distributions of old-class as additional background samples in classifier $g^t(\cdot)$. As parameters of model M^t are updated to adapt new-class, the representation of old-class appears deviation with high probability. We introduce the ADC strategy to estimate the representation deviation of each old-class and dynamically update saved prototypes with deviation compensation. In addition, we use UAC and CPD loss to constrain prediction uncertainty while enhancing the distinguish ability between new-class features and old-class prototypes. Consequently, the integrated objective of each incremental step is defined as:

$$\mathcal{L} = \mathcal{L}_{mbce} + \alpha \mathcal{L}_{kd} + \beta \mathcal{L}_{uac} + \gamma \mathcal{L}_{cpd}, \quad (1)$$

where \mathcal{L}_{mbce} and \mathcal{L}_{kd} denote multiple binary cross-entropy (mBCE) loss and knowledge distillation loss as in (Cha et al. 2021; Chen et al. 2023) respectively. α, β, γ are hyperparameters to balance the weights among above terms.

Adaptive Deviation Compensation

After the training process of step $t - 1$, we follow the prototype replay strategy (Chen et al. 2023) to store the old-class prototypes and necessary statistics for replaying the gaussian-distributed old-class features. These replayed features are treated as additional background samples to mitigate the bias in the new-class classifiers. However, the representation of old-class features exhibits deviation due to the gradient-updated parameters within different new training data distributions. This renders the fixed prototype replay strategy of limited effectiveness.

Intuitively, the deviation in representation should be computed to correct the replayed feature distributions of old classes. Unfortunately, due to ethical and privacy considerations, the old data is unavailable. As a result, we indirectly estimate the representation deviation in the current incremental step using a limited number of old-class pixels present in the current new-class training dataset (e.g., an image containing a *car* (old class) and a *person* (new class)).

Specifically, after several training epochs in step t , we first

obtain the prediction of each image with current model by:

$$\bar{y}_i^t = \begin{cases} \hat{y}_i^t & \text{if } y_i^t = c_b^t \wedge \phi_i^t \geq \tau \\ 0 & \text{otherwise} \end{cases}, \quad (2)$$

where c_b^t represents the background class that contains old-class and truth background in step t , and \wedge represents the co-taking of conditions. ϕ_i^t denotes the certainty scores of prediction, the details of which will be given in Eq. 9. Analogously, the prediction \bar{y}_c^{t-1} is obtained by replacing ϕ_i^t with ϕ_i^{t-1} in Eq. 2. τ is a hyperparameter to screen out the high-confidence old-class predictions.

Since the predictions of the current model differ from the previous predictions, computing the sub-prototype of each old class using the respective predictions results in misaligned distributions. Therefore, we generate the unified prediction masks by:

$$\hat{y}_i^{C^{t-1}} = \begin{cases} \bar{y}_i^t & \text{if } \bar{y}_i^t = \bar{y}_i^{t-1} \\ 0 & \text{otherwise} \end{cases}. \quad (3)$$

Next, we calculate the feature centers of old-class on current data space. We utilize previous feature extractors and current ones to capture features $\{f_i^{t-1}\}_{i=1}^{N^t}$ and $\{f_i^t\}_{i=1}^{N^t}$, where $f_i^{t-1} \in \mathbb{R}^{h \times w \times d}$ is computed by $f_i^{t-1} = f^{t-1}(\mathbf{x}_i^t)$ and f_i^t is so on. We perform the average operation on normalized features of previous extractors with unified prediction masks to obtain sub-prototype $\hat{P}_c^{t-1} \in \mathbb{R}^d$ of each old-class by:

$$\hat{P}_c^{t-1} = \frac{\sum_{i=1}^{N^{t-1}} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^{t-1} \odot \delta\{\tilde{y}_{i,j}^{C^{t-1}} = c\})}{\left\| \sum_{i=1}^{N^t} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^{t-1} \odot \delta\{\tilde{y}_{i,j}^{C^{t-1}} = c\}) \right\|_2}, \quad (4)$$

where j represents the spatial location, \odot and $\delta\{\cdot\}$ indicate element-wise multiplication and indicator function respectively. $\tilde{y}_i^{C^{t-1}}$ and \hat{f}_i^{t-1} denote downsampled unified masks and features. $\|\cdot\|_2$ represents L2-norm. Analogously, the sub-prototype \hat{P}_c^t is obtained by replacing $\hat{f}_{i,j}^{t-1}$ with $\hat{f}_{i,j}^t$ in Eq. 4.

With \hat{P}_c^{t-1} and \hat{P}_c^t obtained from old and new model inference on the same prediction region, we intuitively compute the deviation towards old-class in current sub-dataset by:

$$\Delta_{c^{t-1} \rightarrow t} = \vec{\Omega}(\hat{P}_c^{t-1}, \hat{P}_c^t), \quad (5)$$

where $\vec{\Omega}(\cdot, \cdot)$ indicates the displacement vector difference operation. Hence, the shifted prototype of each old-class is estimated by:

$$\hat{P}_c^{t-1} = P_c^{t-1} + \Delta_{c^{t-1} \rightarrow t}, \quad (6)$$

where P_c^{t-1} is stored prototype of old-class c . However, there exists uncertainty in the prediction masks $\hat{y}_i^{C^{t-1}}$ and statistical limitation of sub-prototype compared to the certain prototype with ground truth mask and abundant statistic information in the previous steps. To ensure the robustness of replayed features, we combine the shifted prototypes and stored prototypes with adaptive weighted average to obtain the compensatory prototype of each old-class by:

$$\bar{P}_c^{t-1} = \rho_c^{t-1} \hat{P}_c^{t-1} + (1 - \rho_c^{t-1}) P_c^{t-1}. \quad (7)$$

Here, ρ_c is an adaptive weight to consider the contribution between two terms, which is formulated by:

$$\rho_c^{t-1} = \frac{\sum_{i=1}^{N^t} \sum_{j=1}^{h \times w} \delta\{\tilde{y}_{i,j}^{C^{t-1}} = c\}}{\eta_c^{1:t-1} + \sum_{i=1}^{N^t} \sum_{j=1}^{h \times w} \delta\{\tilde{y}_{i,j}^{C^{t-1}} = c\}}, \quad (8)$$

where $\eta_c^{1:t-1}$ denotes the sum of pixels belonging to old-class c (contain the labels and unified prediction masks in previous 1, ..., $t-1$ steps).

The compensatory prototypes of the old classes, obtained through a training-free approach, serve to generate sufficient feature distributions that are then replayed in the current new-class classifiers. These distributions not only encompass the original distributions of the old data but also account for the representation deviation between the incremental models. With these conditions in place, the adaptive deviation compensation strategy enhances the efficiency and robustness of the prototype replay.

Uncertainty-Aware Constraint

As a standard practice, pseudo-labels generated from the previous model are employed to tackle the problem of background shift. Specifically, the background values in the current labels at the current step are replaced by the predictions from previous models using a threshold filtering strategy, which may be either fixed (Douillard et al. 2021) or dynamic (Zhao, Yuan, and Shi 2023). However, these strategies focus solely on the maximum value of the logistic outputs, overlooking the interrelationships within the logistic distributions. This introduces uncertainty in the pseudo-labels, which hampers the effective updating of prototypes stored from previous steps.

To address this issue, we propose an uncertainty-aware constraint loss that takes into account the interrelationships among class logistic outputs. This loss encourages the model to produce compact representations for pixels belonging to the same old class. At step t , we first compute the certainty scores of predictions mentioned in Eq. 2.

Specifically, let $\hat{S}_i^t \in \mathbb{R}^{H, W, C^1 + \dots + C^t}$ denote as the logistic outputs of image i obtained by the current model (which include all classes need to be learning in current step t). We compute the certainty scores of predictions by:

$$\phi_i^t = \Lambda(\varphi(\hat{S}_i^t), 1) - \Lambda(\varphi(\hat{S}_i^t), 2), \quad (9)$$

where $\Lambda(\cdot)$ represents the maximum heap sorting and selection function, and $\varphi(\cdot)$ denotes the sigmoid function. Since our primary focus is on enhancing low-probability predictions, high-probability predictions need to be filtered out to avoid overconfidence. We compute the masks by utilizing the ground truth and predictions:

$$m_i^t = \begin{cases} 0 & \text{if } (y_i^t = \hat{y}_i^t) \in C^t \text{ or } \max(\varphi(\hat{S}_i^t)) \geq \tau \\ 1 & \text{otherwise} \end{cases}. \quad (10)$$

Finally, we achieve the UAC loss as:

$$\mathcal{L}_{uac} = d(u_i^t \odot m_i^t, \Omega_i^t), \quad (11)$$

where $d(\cdot)$ denotes a distance measurement function, $u_i^t = 1 - \phi_i^t$ is denoted as the uncertainty scores of predictions,

Method	15-1 (6 steps)			5-3 (6 steps)			10-1 (11 steps)			2-2 (10 steps)			1-1 (20 steps)		
	old	new	all	old	new	all	old	new	all	old	new	all	old	new	all
MiB(CVPR20)	35.1	13.5	29.7	57.1	42.6	46.7	12.3	13.1	12.7	41.7	26.0	28.2	38.5	8.1	11.0
PLOP(CVPR21)	65.1	21.1	54.6	41.1	23.4	25.9	44.0	15.5	30.5	24.1	11.9	13.7	12.4	11.9	4.7
SSUL(NeurIPS21)	77.3	36.6	67.6	72.4	50.7	56.9	71.3	46.0	59.3	61.4	42.1	44.8	52.6	27.5	29.9
DKD(NeurIPS22)	78.1	42.7	69.7	69.6	53.5	58.1	73.1	46.5	60.4	60.5	45.8	47.9	<u>56.1</u>	24.6	27.6
RCIL(CVPR22)	70.6	23.7	59.4	65.3	41.5	50.3	55.4	15.1	34.3	28.3	19.0	19.4	-	-	-
MicroSeg(NeurIPS22)	<u>80.1</u>	36.8	69.8	77.6	59.0	64.3	72.6	48.7	61.2	61.4	40.6	43.5	55.9	34.2	36.3
IDEC(PAMI23)	77.0	36.5	67.3	67.1	49.0	54.1	70.7	46.3	59.1	-	-	-	-	-	-
CoinSeg(CVPR23)	80.6	36.2	70.1	66.8	51.9	56.2	73.7	45.0	60.1	68.3	46.2	49.4	53.7	27.4	29.9
STAR(NeurIPS23)	79.4	<u>50.3</u>	<u>72.5</u>	71.9	<u>61.5</u>	<u>64.4</u>	<u>73.1</u>	55.4	<u>64.7</u>	59.2	<u>55.0</u>	<u>55.6</u>	43.6	<u>35.7</u>	<u>36.5</u>
Ours	79.9	51.9	73.2	<u>73.8</u>	61.9	65.3	74.9	<u>54.3</u>	65.1	<u>62.8</u>	57.9	58.6	63.4	40.5	42.7

Table 1: Quantitative results on the validation set of PASCAL VOC (Everingham et al. 2010) for *overlapped* settings. The best and second best performances are in bold and underline, respectively.

and Ω_i^t is the ideal targets of uncertainty. With \mathcal{L}_{uac} , the uncertainty in predictions is mitigated, features of the same class are aggregated to enhance model rigidity.

Compensation-based Prototype Discrimination

The previously proposed ADC strategy and UAC loss successfully address representation deviation and prediction uncertainty during incremental steps. However, feature extractors may generate indistinguishable features in two scenarios: when background pixels resemble the foreground and when foreground pixels resemble the background. These indistinguishable features lead to confusion in current classifiers for new classes. To mitigate this issue, we propose a compensation-based prototype similarity discriminative loss to distinguish between similar representations of new and old classes.

Specifically, we consider the similarity class issue from two perspectives. First, for pixels of new classes at the current step, we can leverage ground truth labels to obtain the centers of new-class features within a batch:

$$\zeta_c^t = \frac{\sum_{i \in \mathcal{B}} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^t \odot \delta\{\tilde{y}_{i,j}^t = c\})}{\left\| \sum_{i \in \mathcal{B}} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^t \odot \delta\{\tilde{y}_{i,j}^t = c\}) \right\|_2}, \quad (12)$$

where \mathcal{B} denotes the size of batch. Intuitively, we expect to discriminate them with the compensatory old-class prototypes:

$$\mathcal{L}_{n \leftrightarrow o} = \frac{1}{|C^t|} \sum_{c \in C^t} \frac{1}{\min_{c \in C^{1:t-1}} \|\zeta^t - \bar{P}^{t-1}\|_2 + \epsilon}, \quad (13)$$

where ϵ is a constant to avoid zero denominator. On the other hand, for the pixels belonging to the background classes (containing old or feature classes) that are misclassified as foreground classes in the current step, we believe that these pixels are similar to the new classes, thus confusing the new-class classifiers. We compute the feature centers of these misclassified pixels:

$$\tilde{\zeta}_c^t = \frac{\sum_{i \in \mathcal{B}} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^t \odot \delta\{\tilde{y}_{i,j}^t \neq c \wedge \hat{y}_{i,j}^t = c\})}{\left\| \sum_{i \in \mathcal{B}} \sum_{j=1}^{h \times w} (\hat{f}_{i,j}^t \odot \delta\{\tilde{y}_{i,j}^t \neq c \wedge \hat{y}_{i,j}^t = c\}) \right\|_2}. \quad (14)$$

Accordingly, We also add a loss term to separate the features of these background pixels from the features of foreground pixels:

$$\mathcal{L}_{n_{pos} \leftrightarrow n_{neg}} = \frac{1}{|C^t|} \sum_{c \in C^t} \frac{1}{\|\zeta_c^t - \tilde{\zeta}_c^t\|_2 + \epsilon}. \quad (15)$$

Combining the above two loss terms, we can obtain the final CPD loss:

$$\mathcal{L}_{cpd} = \mathcal{L}_{n \leftrightarrow o} + \mathcal{L}_{n_{pos} \leftrightarrow n_{neg}}. \quad (16)$$

Using \mathcal{L}_{cpd} , the features of new classes are not only separated from the compensatory prototypes of old classes but also from the features of old classes within the same batch. Additionally, the CPD loss motivates the model to learn discriminative representations for new classes, which plays a crucial role in storing representative prototypes for new classes at the current step.

Experiments

Experimental Setup

Datasets and Evaluation Metrics. We evaluate our proposed Adapter on Pascal VOC 2012 (Everingham et al. 2010) and ADE20K (Zhou et al. 2017). Pascal VOC 2012 consists of 20 foreground classes and one background class while ADE20K is a large-scale semantic segmentation dataset containing 150 classes. For evaluation, we follow the work(Douillard et al. 2021) to use mean Intersection over Union (mIoU) as the evaluation metric.

Protocol. Following (Baek et al. 2022; Chen et al. 2023), we mainly consider the *overlapped* settings in CISS, which means the background pixels in step t might contain old and future classes. We denote by $(N_1 - N_2)$ the incremental scenario, where N_1 and N_2 are class numbers of the initial step and each incremental step, respectively. For each benchmark dataset, we follow (Baek et al. 2022) to evaluate our method under multiple incremental scenarios (i.e., 19-1, 15-5, 15-1, 5-3, and 10-1 in Pascal VOC, and 100-50, 50-50, 100-10 in ADE20K). In addition, we also evaluate our method on more challenging incremental scenarios (i.e., 2-2 and 1-1 in Pascal VOC, and 100-5 in ADE20K).

Implementation Details. We follow the common practice (Cermelli et al. 2020) to use DeepLabv3 (Chen et al.

Method	100-50 (2 steps)			50-50 (3 steps)			100-10 (6 steps)			100-5 (11 steps)		
	old	new	all	old	new	all	old	new	all	old	new	all
MiB(CVPR20)	40.5	17.2	32.8	45.6	21.0	29.3	38.2	11.1	29.2	36.0	5.7	26.0
PLOP(CVPR21)	41.9	14.9	32.9	48.8	21.0	30.4	40.5	13.6	31.6	39.1	7.8	28.8
SSUL(NeurIPS21)	41.3	18.0	33.6	48.4	20.2	29.6	40.2	18.8	33.1	39.9	17.4	32.5
RCIL(CVPR22)	42.3	18.8	34.5	48.3	24.6	32.5	39.3	17.7	32.1	38.5	11.5	29.6
IDEC(PAMI23)	42.0	18.2	34.1	47.4	26.0	33.1	40.3	17.6	32.7	39.2	14.6	31.0
STAR(NeurIPS23)	42.4	24.2	36.4	48.7	<u>27.2</u>	34.4	42.0	20.6	34.9	41.7	17.5	33.7
Ours	43.1	<u>23.6</u>	36.7	49.3	27.3	34.7	42.9	<u>19.9</u>	35.3	42.6	18.0	34.5

Table 2: Quantitative results on the validation set of ADE20K (Zhou et al. 2017) for *overlapped* settings. The best and second best performances are in bold and underline, respectively.

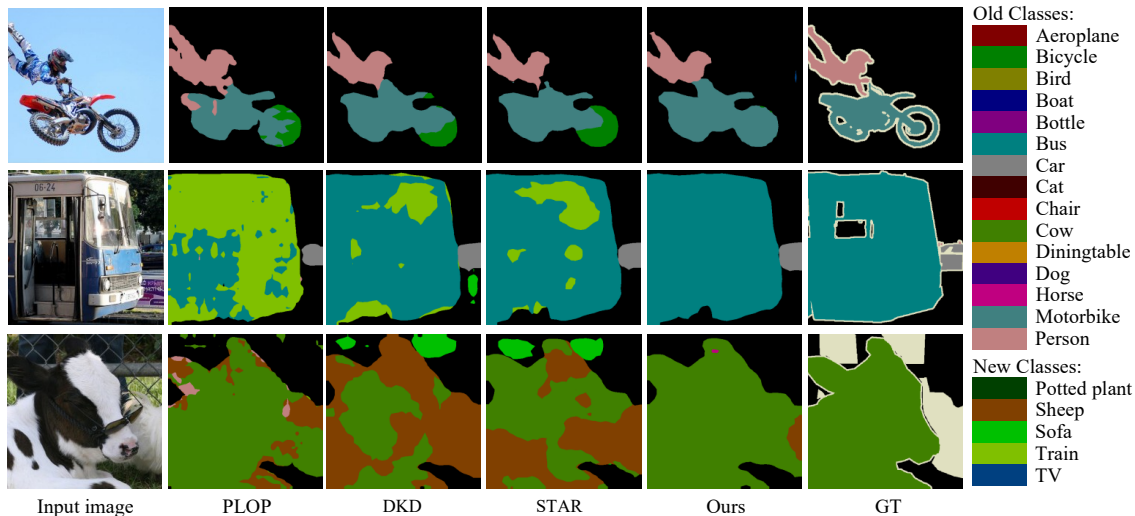


Figure 3: Qualitative comparison on Pascal VOC 2012 between Adapter and previous methods.

2017) with ResNet-101 (He et al. 2016) pre-trained on ImageNet as the segmentation network. In line with (Douillard et al. 2021; Oh, Baek, and Ham 2022), we utilize different training strategies for two datasets. For Pascal VOC 2012, we set the initial learning rate to 10^{-3} and 10^{-4} for training with 60 epochs in the initial step and incremental steps, respectively. For ADE20K, we train for 100 epochs with an initial learning rate of 2.5×10^{-4} for the initial step and 2.5×10^{-5} for incremental steps. The model is optimized by SGD with a momentum of 0.9, and the batch size is set to 24. Hyper-parameters α , β , γ , and τ are set to 5, 0.1, 0.05, and 0.7, respectively. We run our experiments on NVIDIA RTX 4090 GPUs using PyTorch.

Experimental Results

Comparison on Pascal VOC. We evaluate our Adapter on the Pascal VOC 2012 dataset across various incremental scenarios, including standard settings in other methods (e.g., 15-1, 5-3, and 10-1) and two challenging long-term scenarios, namely 2-2 and 1-1, which involve continuously introducing new classes. The comparative results of our method against classical and prior CISS methods are presented in Tab. 1. We denote "old" as the mIoU of the classes containing the background class in the initial step, and "new" as the mIoU of the classes for all incremental steps. The

results demonstrate that our method consistently delivers superior performance across various incremental scenarios. Especially in the challenging multi-step incremental settings, such as the *overlapped* 2-2 (10 steps) and 1-1 (steps) tasks, our method surpasses the state-of-the-art by 3.0% and 6.2% in terms of mIoU, respectively. Fig. 3 shows qualitative results of our approach compared with other competitors on an *overlapped* 15-1 setting, highlighting the superiority of our method. For example, our method accurately classifies the tires of the *motorbike* in row 1, unlike several competitors, who misidentify these pixels as part of bicycles. Additionally, our method distinctly differentiates the *bus* (in row 2) and the *cow* (in row 3) from the similar-looking *train* and *sheep* classes. In contrast, the *train* and *sheep* classes, added during incremental steps, lead to catastrophic forgetting of similar classes in competing methods.

Comparison on ADE20K. Table 2 presents the comparison results on the ADE 100-50, 50-50, 100-10, and 100-5 tasks. Despite the larger class number in ADE20K, our method once again outperforms the previous state-of-the-art. Specifically, for the ADE20K *overlapped* 100-5 (11 steps) setting, we obtain a 34.5 mIoU, which is a 0.8 gain over the second-best method. This further demonstrates the capability of our method to address more challenging scenarios.

Baseline	ADC	UAC	CPD	15-1 (6 steps)		
				old	new	all
✓				78.7	47.4	71.3
✓	✓			79.2	48.5	71.9
✓	✓	✓		79.5	49.8	72.4
✓	✓	✓	✓	79.9	51.9	73.2

Table 3: Ablation study of our method components on PASCAL VOC *overlapped* 15-1 (6 steps).

Copro	$\mathcal{L}_{n \leftrightarrow o}$	$\mathcal{L}_{n_{pos} \leftrightarrow n_{neg}}$	15-1 (6 steps)		
			old	new	all
			79.5	49.8	72.4
	✓		79.6	50.6	72.7
		✓	79.6	50.3	72.6
	✓	✓	79.8	51.0	72.9
✓	✓		79.8	51.4	73.0
✓	✓	✓	79.9	51.9	73.2

Table 4: Ablation study results for the component of CPD loss on PASCAL VOC *overlapped* 15-1 (6 steps).

Ablation Study

Method Components. we present an evaluation of each component of the proposed Adapter in Tab. 3. The "baseline" denotes the model using fixed prototype replay strategy and knowledge distillation in (Chen et al. 2023) with mBCE loss, which already achieves competitive performance. From the comparison between the first and second row in Tab. 3, we can see that the ADC strategy improves the both performance of old and new classes (0.5 and 1.1 mIoU respectively), which arises from ADC compensating the representation deviation toward old-class and providing the appropriate features in new-class classifiers. Finally, our ADC strategy boosts the mIoU of all classes by 0.6, further demonstrating the feasibility of it. Comparing the second row with the third row, the effectiveness of UAC loss can be recognized. Specifically, UAC obtained gains of 1.3 in terms of new-class mIoU and 0.5 for all classes. This benefit by the constraint on uncertainty predictions, thereby compressing each class representation to a compact prototype and reserve latent space for learning new classes. Removing the CPD leads to the performance declining by 2.1 and 0.8 in terms of mIoU (row 3 vs. row 4) for new and all classes, respectively. Without the CPD loss, it is challenging to differentiate between representations of similar classes. This suggests that our CPD loss is more effective for generating distinguished features of each class and facilitating the effectiveness of the adaptive prototype replay strategy.

Components of the CPD. We present an ablation study of the CPD loss components on the PASCAL VOC *overlapped* 15-1 scenario, as shown in Table 4. The term "Copro" indicates that the old-class prototypes used in $\mathcal{L}_{n \leftrightarrow o}$ are compensated by our ADC strategy. The results in the table show that, with fixed old-class prototypes, the use of the two constraint terms results in a 0.5% increase for all classes. When the old-class prototypes are updated using the ADC strategy, the mIoU improves by an additional 0.3%. Each of the two loss terms also contributes to the enhanced performance.

Method	15-1 (6 steps)		
	old	new	all
MiB	35.1	13.5	29.7
MiB+UAC	40.5	17.4	35.0
PLOP	65.1	21.1	54.6
PLOP+UAC	66.1	30.0	59.2
DKD	78.1	42.7	69.7
DKD+UAC	78.5	46.4	70.9
CoinSeg	80.6	36.2	70.1
CoinSeg+UAC	80.8	40.9	71.3

Table 5: UAC in other representative CISS methods. UAC consistently enhances segmentation performance.

β	15-1 (6 steps)	γ	15-1 (6 steps)	τ	15-1 (6 steps)
0.01	72.6	0.01	72.5	0.5	72.3
0.05	73.0	0.05	73.2	0.6	72.8
0.1	73.2	0.1	72.9	0.7	73.2
0.5	72.9	0.5	72.1	0.8	73.0
1	72.7	1	71.8	0.9	72.5

Table 6: Search of hyper-parameters: the weights of loss β and γ , and threshold τ . mIoU for all classes are reported

Generality of the UAC. We applied our UAC loss to two representative CISS methods, namely MiB (Cermelli et al. 2020) and PLOP (Douillard et al. 2021). Compared to MiB, our UAC improves performance by 5.3%, with increases of 5.4% and 3.9% in mIoU for old and new classes on PASCAL VOC *overlapped* 15-1 scenario, respectively. In addition, implementation of PLOP with our UAC also have 4.6% performance improvement. These results provide convincing evidence that our insight of reducing prediction uncertainty is effective and generalizes well across different CISS methods.

Ablation of Hyper-parameters. Tab. 6 illustrates the influence of hyper-parameters: the weight of proposed loss terms β and γ , and threshold τ in method. The results show that, in most cases, our approach is not highly sensitive to the choice of hyperparameters.

Conclusion

In this paper, we propose a novel method, called Adapter, designed to address representation deviation in prototype replay based CISS methods. First, we introduce an adaptive deviation compensation strategy to estimate representation deviation between incremental models and update the stored prototype used for replay. Second, we employ an uncertainty-aware constraint loss and a compensation-based prototype discriminative loss to aggregate the representation of each class and enhance the discrimination of prototypes. Finally, experiments demonstrate the effectiveness of our proposed Adapter, which achieves remarkable performance, especially on challenging multi-step incremental settings, outperforming the previous state-of-the-art.

Acknowledgments

This work was supported by the Hubei Provincial Natural Science Foundation of China No.2022CFA055, the National Natural Science Foundation of China No.62176097.

References

- Baek, D.; Oh, Y.; Lee, S.; Lee, J.; and Ham, B. 2022. Decomposed knowledge distillation for class-incremental semantic segmentation. *Advances in Neural Information Processing Systems*, 35.
- Cermelli, F.; Mancini, M.; Bulò, S. R.; Ricci, E.; and Caputo, B. 2020. Modeling the background for incremental learning in semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9233–9242.
- Cha, S.; Yoo, Y.; Moon, T.; et al. 2021. Ssul: Semantic segmentation with unknown label for exemplar-based class-incremental learning. *Advances in neural information processing systems*, 34.
- Chen, J.; Cong, R.; Luo, Y.; Ip, H.; and Kwong, S. 2023. Saving 100x Storage: Prototype Replay for Reconstructing Training Sample Distribution in Class-Incremental Semantic Segmentation. *Advances in Neural Information Processing Systems*, 36.
- Chen, L.-C.; Papandreou, G.; Kokkinos, I.; Murphy, K.; and Yuille, A. L. 2017. 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.
- Douillard, A.; Chen, Y.; Dapogny, A.; and Cord, M. 2021. Plop: Learning without forgetting for continual semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4040–4050.
- Douillard, A.; Cord, M.; Ollion, C.; Robert, T.; and Valle, E. 2020. Podnet: Pooled outputs distillation for small-tasks incremental learning. In *Proceedings of the European conference on computer vision (ECCV)*, 86–102. Springer.
- Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88: 303–338.
- Gou, J.; Yu, B.; Maybank, S. J.; and Tao, D. 2021. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6): 1789–1819.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Jung, S.; Ahn, H.; Cha, S.; and Moon, T. 2020. Continual learning with node-importance based adaptive group sparse regularization. *Advances in neural information processing systems*, 33.
- Kirkpatrick, J.; Pascanu, R.; Rabinowitz, N.; Veness, J.; Desjardins, G.; Rusu, A. A.; Milan, K.; Quan, J.; Ramalho, T.; Grabska-Barwinska, A.; et al. 2017. Overcoming catastrophic forgetting in neural networks. volume 114, 3521–3526. National Acad Sciences.
- Li, Z.; and Hoiem, D. 2017. Learning without forgetting. *IEEE transactions on pattern analysis and machine intelligence*, 40(12): 2935–2947.
- Lopez-Paz, D.; and Ranzato, M. 2017. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30.
- Maracani, A.; Michieli, U.; Toldo, M.; and Zanuttigh, P. 2021. Recall: Replay-based continual learning in semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 7026–7035.
- Masana, M.; Liu, X.; Twardowski, B.; Menta, M.; Bagdanov, A. D.; and Van De Weijer, J. 2022. Class-incremental learning: survey and performance evaluation on image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(5): 5513–5533.
- McCloskey, M.; and Cohen, N. J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, 109–165. Elsevier.
- Michieli, U.; and Zanuttigh, P. 2019. Incremental learning techniques for semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision workshops*.
- Oh, Y.; Baek, D.; and Ham, B. 2022. Alife: Adaptive logit regularizer and feature replay for incremental semantic segmentation. *Advances in Neural Information Processing Systems*, 35.
- Rebuffi, S.-A.; Kolesnikov, A.; Sperl, G.; and Lampert, C. H. 2017. icarl: Incremental classifier and representation learning. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, 2001–2010.
- Shin, H.; Lee, J. K.; Kim, J.; and Kim, J. 2017. Continual learning with deep generative replay. *Advances in neural information processing systems*, 30.
- Shmelkov, K.; Schmid, C.; and Alahari, K. 2017. Incremental learning of object detectors without catastrophic forgetting. In *Proceedings of the IEEE international conference on computer vision*, 3400–3409.
- Yang, G.; Fini, E.; Xu, D.; Rota, P.; Ding, M.; Nabi, M.; Alameda-Pineda, X.; and Ricci, E. 2022. Uncertainty-aware contrastive distillation for incremental semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(2): 2567–2581.
- Yoon, J.; Yang, E.; Lee, J.; and Hwang, S. J. 2018. Life-long Learning with Dynamically Expandable Networks. In *International Conference on Learning Representations*.
- Yu, L.; Twardowski, B.; Liu, X.; Herranz, L.; Wang, K.; Cheng, Y.; Jui, S.; and Weijer, J. v. d. 2020. Semantic drift compensation for class-incremental learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6982–6991.
- Zenke, F.; Poole, B.; and Ganguli, S. 2017. Continual learning through synaptic intelligence. In *International conference on machine learning*, 3987–3995. PMLR.
- Zhang, Z.; Gao, G.; Fang, Z.; Jiao, J.; and Wei, Y. 2022. Mining unseen classes via regional objectness: A simple

baseline for incremental segmentation. *Advances in neural information processing systems*, 35: 24340–24353.

Zhang, Z.; Gao, G.; Jiao, J.; Liu, C. H.; and Wei, Y. 2023. Coinseg: Contrast inter-and intra-class representations for incremental segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 843–853.

Zhao, D.; Yuan, B.; and Shi, Z. 2023. Inherit With Distillation and Evolve With Contrast: Exploring Class Incremental Semantic Segmentation Without Exemplar Memory. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(10): 11932–11947.

Zhao, H.; Wang, H.; Fu, Y.; Wu, F.; and Li, X. 2021. Memory-efficient class-incremental learning for image classification. *IEEE Transactions on Neural Networks and Learning Systems*, 33(10): 5966–5977.

Zhou, B.; Zhao, H.; Puig, X.; Fidler, S.; Barriuso, A.; and Torralba, A. 2017. Scene parsing through ade20k dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 633–641.