






Learning from the Web: Language Drives Weakly-Supervised Incremental Learning for Semantic Segmentation

Chang Liu¹, Giulia Rizzoli², Pietro Zanuttigh², Fu Li¹, and Yi Niu¹^{*}

¹ School of Artificial Intelligence, Xidian University, China

² Department of Information Engineering, University of Padova, Italy

Abstract. Current weakly-supervised incremental learning for semantic segmentation (WILSS) approaches only consider replacing pixel-level annotations with image-level labels, while the training images are still from well-designed datasets. In this work, we argue that widely available web images can also be considered for the learning of new classes. To achieve this, firstly we introduce a strategy to select web images which are similar to previously seen examples in the latent space using a Fourier-based domain discriminator. Then, an effective caption-driven rehearsal strategy is proposed to preserve previously learnt classes. To our knowledge, this is the first work to rely solely on web images for both the learning of new concepts and the preservation of the already learned ones in WILSS. Experimental results show that the proposed approach can reach state-of-the-art performances without using manually selected and annotated data in the incremental steps.

Keywords: Continual Semantic Segmentation · Weakly-Supervised Learning · Web-based learning

1 Introduction

The rapid development of deep learning techniques allowed for impressive results in Semantic segmentation (SS), i.e., pixel-level labeling of real-world scenes. Traditional supervised SS approaches are constrained to predetermined classes, leading to the problem of *catastrophic forgetting* when learning from new data. To address this, incremental learning techniques, such as Class Incremental Learning for Semantic Segmentation (CILSS), have been proposed. Several CILSS strategies aim to balance new and old knowledge by using regularization strategies or storing past examples [8, 29]. However, these methods heavily rely on expensive pixel-level annotations. Weakly-Supervised Incremental Learning for Semantic Segmentation (WILSS) has recently emerged to reduce annotation costs, utilizing image-level labels for guidance. Nevertheless, existing WILSS approaches face limitations, such as requiring multiple classes at each step and

^{*} Corresponding author




Fig. 1: General overview of the proposed method.

performance degradation with multiple steps [7, 49]. In this study, we focus on a novel paradigm, Web-based Weakly-Supervised Incremental Learning, which leverages both existing knowledge and new information from the web. This approach finds applications across various domains, particularly when pre-trained models need adaptation to new classes where data are limited or unavailable, e.g., due to privacy concerns [38]. Leveraging web resources becomes crucial to supplement the existing model with additional class information, enabling adaptation without extensive data collection efforts. However, within this framework, new challenges emerge: 1) Although external sources such as the internet provide a potential solution, web data often come with noisy image-level supervision, containing undesired classes or unlabeled classes that were previously learned. We tackle this issue by selecting images using adversarial learning in the Fourier domain and leveraging language models to provide multi-class label supervision. 2) Additionally, without access to previous training data, the model may suffer from catastrophic forgetting. To address this issue, we propose leveraging web resources also for preserving the proficiency of the model on old classes. Nevertheless, integrating this data into the existing model poses challenges, especially when the new data deviates from the original distribution. This misalignment affects the training process, requiring innovative strategies to ensure effective model adaptation to the new classes. To this extent, instead of querying through the image label (i.e., the class name), we used captions that better represent the image content. An overview of our framework is shown in Fig. 1, in summary the main contributions are:

- We present a comprehensive framework that leverages web images to tackle the challenges of learning new classes in an incremental fashion and of knowledge preservation in scenarios with limited or inaccessible training data;
- By integrating vision-language captioning, we enhance image-level supervision with textual descriptions;
- By downloading images using vision-language captioning, we can improve the searched queries and re-generate them to further verify that images have the expected content.
- Concluding, we introduce a new flexible data pipeline that allows to train continual learning schemes even without resorting to curated datasets.

2 Related Work

Class-Incremental Learning for SS (CILSS) Class-incremental continual learning has been extensively explored within the image classification domain, with recent efforts extending into the more complex task of semantic segmentation [29]. Initially, techniques such as knowledge distillation, parameters freezing, and class re-weighting were employed to introduce new classes while retaining knowledge from prior ones [8, 19, 29]. Subsequently, researchers delved into regularization and contrastive mechanisms at the feature level to enhance class-specific capabilities and maintain spatial relationships [13, 30, 46]. Further advancements include the proposal of novel objective functions to account for class similarity [32], and biased-context-insensitive consistency [51]. Other proposed strategies consider dynamic architectures to balance the network parameters [45, 50] and rehearsal data to preserve old knowledge [24].

Weakly Supervised Incremental Learning for SS (WILSS) Weakly Supervised Semantic Segmentation techniques aim to develop effective segmentation models using more affordable forms of supervision, such as bounding boxes [11, 39], scribbles [22], saliency maps [20, 43], points [28], and image-level labels [1, 3]. Following this concept, Weakly Supervised Incremental Learning for Semantic Segmentation (WILSS) involves enhancing pre-trained segmentation models with new classes using only image-level supervision. To this extent, WILSON [7] proposes a comprehensive framework, which can learn progressively from pseudo-labels generated online by a localizer attached to the model. FMWISS [49] seeks to enhance pseudo-labels by leveraging complementary foundational models. Despite the efforts, both methods are unable to perform single-class incremental learning steps, requiring negative examples to properly guide the training. Lastly, RaSP [37] relies on foundation models to build semantic connections between image-level labels.

Webly-Supervised Learning is an emerging field in research where deep learning models are trained using large amounts of data from the web [10, 31, 40]. Current research efforts are focused on understanding how to effectively query and select samples [14, 33] and how to leverage weakly supervised data [26, 47], e.g., by generating pseudo-labels. In CILSS, challenges arise due to the absence of pixel-level supervision and the presence of distribution shifts compared to the original data. The first approach to exploit web data in CILSS is RECALL [24, 27], which investigates preserving previous knowledge by leveraging pixel-level pseudo-labels obtained through either web data querying or the generation of synthetic samples. However, it is essential to acknowledge that utilizing queried images not only lacks pixel-level labeling but also introduces inconsistency at the image-level. In our work, we specifically address the scenario where images may contain multiple classes, despite being queried and labeled under a single class label. This aspect holds particular significance in WILSS - as new knowledge is solely acquired from image-level labels - and incorrect supervision significantly worsens the results. Moreover, to our knowledge, this is the first work exploring the usage of web samples for both learning new knowledge and preserving previous one.

Foundation Models in WILSS The early stage of pre-trained foundation models starts from the text field, such as BERT [12], before expanding into the domain of visual language models (VLMs) [34]. Recent years have witnessed a notable shift towards the integration of image and text, leading to active research in VLMs for a variety of tasks, including text-to-image generation [21], zero-shot segmentation [18, 52], and visual question answering [35]. CLIP [34] serves as a milestone in VLM development, where both image and text encoders are jointly trained to align image-text pairs. Motivated by this, recent studies have integrated VLMs into frameworks like WILSS [37, 49]. Yu et al. [49] exploit the zero-shot capabilities of foundational models, combining DINO [6] and MaskCLIP [53] for pseudo-supervision. Instead, Roy et al. utilize BERT [12] to derive textual descriptions and compute semantic matches between new and old classes. In contrast, our approach utilizes a captioning model, specifically Flamingo [2, 4], to extract meaningful captions solely from image inputs. Given its capability to adapt to novel tasks, we leverage it for tasks such as providing multi-label supervision, querying, and refining web data.

3 Method

In the considered setting we assume to have access to images in a fully-supervised dataset only for the first training phase on an initial set of classes, after which it is no longer available. Instead, data for newly identified classes are exclusively sourced from external platforms, notably the Internet³. Underlying this setup is the assumption that only the names of the new target classes are known, without any supplementary guidance. In Section 3.1, we outline the problem and the considered framework, as established in [7]. Subsequently, in Sections 3.2 and 3.3 we present the novel proposed scenario and our approach to tackling it by encompassing previous limitations. An outline of our method is in Fig. 2.

3.1 Problem Definition

WILSS task. Lets consider an image $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{|\mathcal{I}| \times 3}$ drawn from a distribution \mathcal{D} , where $|\mathcal{I}| = H \times W$. Its semantic description can be in the form of an image-level label $y \in \mathcal{Y} \subset \mathcal{C}$ (we represent it using one-hot vectors, i.e., $y \in \mathbb{R}^C$) or of a more detailed pixel-level segmentation map $\mathbf{y} \in \mathbb{R}^{|\mathcal{I}| \times C}$. The label set $\mathcal{C} = \{b, c_1, \dots, c_{C-1}\}$ is the same for image-level and pixel-level labeling and contains C target classes (b is the background class).

In the class-incremental setting, the training is assumed to happen over a sequence of steps $t = 0, \dots, T$. In contrast to pixel-level segmentation, within the framework of WILSS, at the initial step ($t = 0$) pixel-level labels \mathbf{y} are available, whereas for all subsequent future steps, it becomes solely the image-level label $y \in \mathbb{R}^C$. More in detail, at each step $t > 0$, new classes \mathcal{C}^t with image-level annotations are introduced, leading to the creation of a new label

³ Images were downloaded from the *Flickr* website adhering to its term of use.

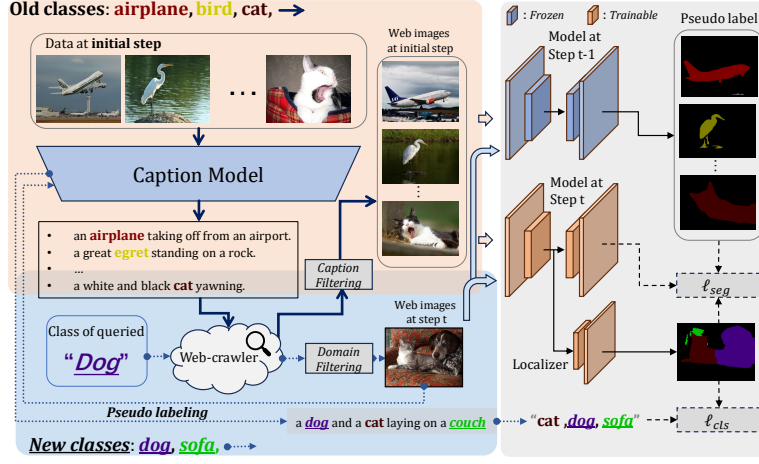


Fig. 2: The proposed method employs a web crawler to gather new knowledge and retrieve past data. The new knowledge (Sec. 3.2) is acquired by querying class names and subsequently filtered in the Fourier domain. Simultaneously, image-level labels are provided by a captioning model. On the other hand, the preservation of old knowledge (Sec. 3.3) involves querying captions from previous data and filtering them based on semantic similarity with the regenerated captions.

set $\mathcal{Y}^t = \mathcal{Y}^{t-1} \cup \mathcal{C}^t$. As in exemplar-free settings, by the time step t is reached, the images from the previous step ($t - 1$) are no longer accessible.

WILSON framework [7]. At each step t , the network is composed of a shared encoder E^t , a segmentation decoder D^t , which is incrementally extended to accommodate new classes, and a localizer head L^t , that is trained from scratch for every task. Additionally, the model from the preceding task is retained and denoted as $(E \circ D)^{t-1}$. The localizer is utilized to generate pseudo-supervision for the segmentation model, providing two guidance: one at the image level and the other at the pixel level. Let $\mathbf{y}_L^t \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{Y}^t|}$ be the pixel-level prediction of the localizer $(E \circ L)^t$. As the localizer is trained from scratch at each step, we simplify the notation by making $\mathbf{y}_L^t = \mathbf{y}_L$. Additionally, we define $y_L \in \mathbb{R}^{|\mathcal{Y}^t|}$ as the image-level counterpart, which is derived from the localizer’s scores by using normalized Global Weighted Pooling as in [3]. Specifically, the image-level supervision is computed as a multi-label soft-margin loss:

$$\mathcal{L}_{CLS}(y, y_L) = -\frac{1}{|\mathcal{C}^t|} \sum_{c \in \mathcal{C}^t} y^c \log(y_L^c) + (1 - y^c) \log(1 - y_L^c) \quad (1)$$

Given that pixel-level predictions from a classifier may contain noise, \mathbf{y}_L is initially smoothed [25] ($\tilde{\mathbf{y}}_L$) - please refer to [7] - and then combined with the old

model predictions:

$$\hat{\mathbf{y}} = \begin{cases} \min(\mathbf{y}_D^{c(t-1)}, \tilde{\mathbf{y}}_L^c) & \text{if } c = b \\ \tilde{\mathbf{y}}_L^c & \text{if } c \in \mathcal{C}^t \\ \mathbf{y}_D^{c(t-1)} & \text{otherwise} \end{cases} \quad (2)$$

where $y_D^{t-1} \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{Y}^{t-1}|}$ is the prediction of the segmentation model $(E \circ D)^{t-1}$. Then, the pixel-level loss is obtained as:

$$\mathcal{L}_{SEG}(\hat{\mathbf{y}}, \mathbf{y}_D^t) = -\frac{1}{|\mathcal{I}||\mathcal{Y}^t|} \sum_{i \in \mathcal{I}} \sum_{c \in \mathcal{Y}^t} \hat{y}_{(i)}^c \log(y_{(i)D}^c) + (1 - \hat{y}_{(i)}^c) \log(1 - y_{(i)D}^c) \quad (3)$$

where we denoted with $\hat{y}_{(i)}$ the predictions for the i -th pixel. Furthermore, to prevent the forgetting of old classes, two additional knowledge distillation losses are proposed in [7]: 1) \mathcal{L}_{KDE} aligns the features from the current encoder E^t with the previous one E^{t-1} through mean-squared error; 2) \mathcal{L}_{KDL} improves consistency between the old model's $(E \circ D)^{t-1}$ and the localizer's $(E \circ L)^t$. Ultimately the objective function is computed as follows: $\mathcal{L} = \mathcal{L}_{SEG} + \mathcal{L}_{CLS} + \mathcal{L}_{KDE} + \mathcal{L}_{KDL}$ (see the supplementary material for details).

3.2 Learning new Knowledge from Web

In CILSS, the web has been used for rehearsal of past classes [24,27]. The samples are gathered from the web by querying each class using the corresponding class name, obtaining the corresponding image set:

$$\mathcal{X}^{web} = \{\mathbf{x} = \mathcal{D}^{web}(q) \mid q = c\} \quad (4)$$

where q is the searched query. As the query q corresponds to the class-name of c , the naive way to provide image-level supervision for $\mathbf{x} \in \mathcal{X}^{web}$ is:

$$y^c = \begin{cases} 1 & \text{if } q = c \\ 0 & \text{otherwise} \end{cases} \quad (5)$$

We introduce a scenario wherein both the images from the preceding step $t-1$ and those from the current step t are unavailable (i.e., access to the original dataset distribution \mathcal{D} is not allowed after the initial step). At each incremental step t , the newly introduced image $\mathbf{x} \in \mathcal{X}^{web} \subset \mathbb{R}^{|\mathcal{I}| \times 3}$ belonging to the new class set \mathcal{C}^t is sampled from a distribution \mathcal{D}^{web} distinct from \mathcal{D} , thus representing an external data source. Notably, in our framework, the test dataset belongs to the original distribution \mathcal{D} . Real dataset images are typically manually selected and verified, resulting in superior content quality and greater diversity. Furthermore, in most cases, their labeling is the product of costly human labour. To explore learning effectively from web images, firstly a Fourier transformation-based domain discriminator is designed for selecting web images which are considered to have more similarities with previous samples in latent space. Then, a pre-trained image caption model is adopted for generating pseudo image-level annotations.

Fourier Domain-based Discriminator Although web images are widely available, by downloading a small sample from the internet there is no guarantee that its distribution resembles the one from the dataset with which we pretrained the model at step 0. In order to preserve the same statistics from the original dataset, we propose to train a discriminator network to distinguish between the two distributions, selecting web images that closely resemble those from the original dataset. However, since no dataset images are available for the current classes in the incremental steps, the discriminator must be trained solely in the initial step and continue to function effectively in subsequent steps.

Based on the above concern, we introduce a Fourier-based domain discriminator. The rationale behind this approach is that amplitude and phase representations in the Fourier domain tend to exhibit more consistency in class statistics changes compared to the pixel domain. Recent domain adaptation techniques have leveraged this consistency, applying amplitude as a stylization to preserve the style of the target distribution [36, 48]. Given an input \mathbf{x} , the prediction from discriminator is:

$$(p_{ds}, p_{web}) = (p(\mathbf{x} \in \mathcal{X}), p(\mathbf{x} \in \mathcal{X}^{web})) = M_D(|\mathcal{F}(\mathbf{x})|) \quad (6)$$

where the two outputs correspond to the scores for the data to belong to the source dataset (e.g., PASCAL-VOC in our experiments) or to the web, respectively. \mathcal{F} represents the Fourier transform operator, the modulus operation extracts the amplitude value of the transformed sample, and M_D is the proposed discriminator. In this context, the corresponding amplitude value of the transformed sample is fed into the discriminator. An image is considered valid if it can deceive the discriminator, meaning $p_{ds}/p_{web} > 1$.

Caption Labeling At each step t , the localizer has to identify objects within the image. Crucially, if the localizer predicts a new class for a given pixel, it is assigned as such, thus ensuring adaptability to novel classes. One of the key challenges we address is the reliance on image-level labels for supervising the localizer. When we introduce unsupervised images for the web instead of the data from the original distribution, as images are downloaded querying the class name and labelled with it (as in Eq. (5)), these labels are prone to inaccuracies, leading to potential misclassifications. Inspired by the remarkable performance of pre-trained vision language models (VLMs), we introduce a strategy that leverages a VLM to enhance image-level supervision. Specifically, we utilize a captioning model M_{CAP} , which takes as input the image $\mathbf{x} \in \mathcal{X} \subset \mathbb{R}^{|\mathcal{I}| \times 3}$ drawn from a distribution \mathcal{D}^{web} , where $\mathcal{D}^{web} \cap \mathcal{D} = \emptyset$. First, we consider a set of words \mathcal{W}^c for each class $c \in \mathcal{Y}^t$, including synonyms, plurals and linguistic alternatives (e.g. “plane” belongs to the “aeroplane” class, “woman” belongs to the “person” class, refer to the supplementary material for the full list). Then, we generate the caption corresponding to each image: $w = M_{CAP}(\mathbf{x})$. The image-level label $y \in \mathbb{R}^C$ is derived from w , $\forall c \in \mathcal{Y}^t$, as:

$$y^c = \begin{cases} 1 & \text{if } \exists w_i \in w : w_i \in \mathcal{W}^c \\ 0 & \text{otherwise} \end{cases} \quad (7)$$

providing multi-label guidance to the web images. For example, a caption mentioning “a person standing on a boat” would yield labels for both “person” and “boat”. This approach not only enhances the granularity of image-level supervision but also addresses inconsistencies in labeling. Moreover, it allows us to identify and discard images that lack the relevant class altogether, thus improving the overall quality of the dataset.

3.3 Rehearsal strategies

Several existing WILSS methodologies [7, 49] face significant limitations: 1) they require the availability of multiple classes at each incremental step. 2) their performance drops largely when multiple steps are performed. To mitigate these issues, we propose a rehearsal strategy that avoids storing original images, which may raise privacy concerns. Specifically, we explore the feasibility of saving solely the image captions and subsequently querying relevant samples from the web. Furthermore, we utilize the caption model for additional supervision to verify whether the downloaded images contain the primary class for which they were obtained. This approach not only addresses potential privacy concerns associated with storing original images but also provides a practical solution for overcoming the constraints observed in conventional WILSS methods. Notably, to adapt WILSS to training on old class data, we modify Eq. (1) and Eq. (3). Specifically, we define $\mathbf{x} \in \mathcal{X}_r$ where \mathcal{X}_r is the set of web rehearsal data. As the image \mathbf{x} should contain previous classes, the segmentation loss is simply $\mathcal{L}_{SEG}(\mathbf{y}_D^{t-1}, \mathbf{y}_D^t)$, while the classification loss of the localizer is $\mathcal{L}_{CLS}(y_D^{t-1}, y_L)$, where the image level label y_D^{t-1} comes from the predicted classes in the pixel-level label \mathbf{y}_D^{t-1} (i.e., $y_D^{c(t-1)} = 1$ if class c is present in \mathbf{y}_D^{t-1}).

Automatically obtaining web images suitable for training the network is challenging due to the lack of control on the retrieval process. Having previously utilized the image-level capabilities of a VLM to enhance image-level supervision, we decided to exploit its ability to generate representative text descriptions for previously trained samples. This brings us three main advantages: 1) we can store the text-level caption of the image sample instead of the image itself, thus strongly reducing storage and privacy concerns; 2) web images queried with accurate captions contain more classes and share more similar semantic content with respect to dataset samples if compared with image queried with just the class name; 3) by re-generating descriptions for queried images, one can easily discard images that do not contain content related to the considered class. The strategy consists of 2 steps: caption-based querying and caption-based filtering.

Caption-based Querying We employed the same model discussed in Sec. 3.2 for captioning, denoted as M_{CAP} . In the initial step, we obtained captions for all inputs \mathbf{x} belonging to the set \mathcal{X} such that y is in \mathcal{Y}^0 . Then, the new downloaded images are obtained as:

$$\mathcal{X}_r^{web} = \{\mathbf{x} = \mathcal{D}^{web}(q') \mid q' = M_{CAP}(\mathbf{x}) : \mathbf{x} \in \mathcal{X}\} \quad (8)$$

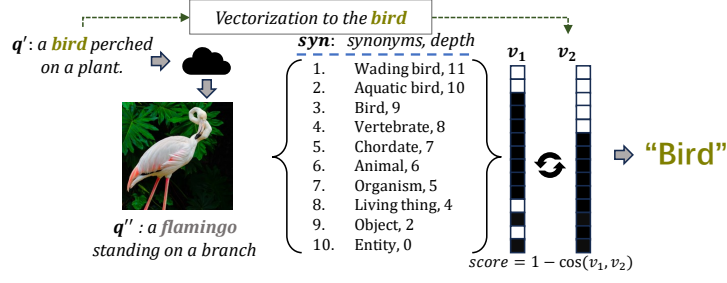


Fig. 3: Illustration of caption-based filtering approach.

replacing the naive downloading process described in Eq. (4). Additionally, leveraging replay data, at step t , we have images from classes \mathcal{Y}^{t-1} . Therefore, concerning Eq. (1), we train the localizer not only on \mathcal{C}^t but also on \mathcal{Y}^t . Specifically, we denote the memory capacity as M , i.e., the number of rehearsal images considered for each class.

Caption-based Filtering Despite querying $\mathbf{x} \in \mathcal{X}_r^{web}$ based on a more representative description rather than the class name, there is no guarantee that samples include the desired class. Moreover, there is also no guarantee that the captions contain semantic language corresponding to the original class names, therefore we needed to rely on a different strategy with respect to Sec. 3.2. Specifically, inspired by classic NLP methods, we used Penn TreeBank [42] and the lexical database WordNet [16]. The approach is structured as follows:

1. we re-generate a second caption q'' by feeding $\mathbf{x} \in \mathcal{X}_r^{web}$ into the caption model as $q'' = M_{CAP}(\mathbf{x})$;
2. we tokenize the two queries and identified two pairs of words (n'_1, n'_2) and (n''_1, n''_2) , which represent the first two nouns from sentences q' and q'' , respectively, based on their syntax tags as in [42];
3. for each noun n we extract a set of hypernyms, i.e. its cognitive synonyms (see an example in Fig. 3), from the WordNet Tree [16];
4. for each noun n we create a corresponding vectorized descriptor v (e.g., v'_1 corresponds to n'_1 , and so on) by considering the depth d of each of its hypernyms in the tree and updating $v[d] += 1$;
5. we compute the cosine distance between each vector (e.g., we compare v'_1 with both v''_1 and v''_2);
6. if at least one couple of descriptors is similar (e.g., v'_1 with v''_1 or v''_2), then the image \mathbf{x} is kept, meaning that it contains at least one concept that was contained in the original image. We assume two vectors v_i and v_j to be similar if $1 - \cos(v_i, v_j) > T$.

Essentially, we assessed whether the words in the captions belong to the same semantic family. While this approach does not guarantee that the words correspond to the set of the old classes, this strategy enforces the same image content from the original distribution, which is not limited to the learnt categories.

4 Experimental Results

WILSS Dataset and Settings Aligning with WILSON [7], we experimented on PASCAL-VOC (VOC) [15] and COCO [23] benchmarks. We considered two setups: 1) Single-step multi-class, with one single incremental step with multiple classes (e.g., **15-5**, learn 15 classes at step 0, then 5 classes together); 2) multi-step single-class, with multiple incremental steps while learning only one class (e.g., **15-1**, one class for each incremental step). In the COCO-to-VOC setting, the initial step is computed on COCO, while the incremental steps on VOC, results for this setting please refer to the Sec. 4.1. We adopted two standard settings [29]: *disjoint* and *overlapped*. In both, only the image-level labels belonging to the current task are provided (i.e., $y \in \mathcal{C}^t$). In disjoint, for each step t , images of the current task contain only classes in \mathcal{Y}^t , while in overlapped, images can contain any class, even future ones.

Web Dataset We utilize the Flickr website as our primary source for web images due to its extensive collection of diverse images. However, it is important to note that if licensing poses an issue, alternative publicly available sources can be explored. For each class being learned, a set of 10K web images has been gathered using Eq. (4). For the new knowledge, we select 500 samples for training after implementing the strategies detailed in Sec. 3.2. For the old class images, we retrieved 20 images per caption (N_{CAP}) using the download caption equation (Eq. (8)). However, to ensure comparability to the competitors [37], we employed 100 rehearsal images ($M = 100$).

Implementation Details Following [7], we use DeeplabV3 [9] with ResNet-101 [17] for Pascal VOC and Wide-ResNet-38 [44] for COCO, with output stride equal to 16 and 8 respectively. The model is trained with SGD for 30 epochs at the initial step and 40 epochs for the following incremental steps. For the first 5 epochs, only the localizer is trained. Then the main model is trained with the pseudo annotation combined with localizer output and previous model. For the discriminator M_D , we adopted the lightweight EfficientNet-B0 model [41] followed by 3 fully connected layers with dimensions 1000, 256 and 2 to prevent memory burden. The backbone network has been pre-trained on ImageNet. Then, it is trained only at the initial step, with dataset images of initial classes as positive samples and the web images queried with related class names as negative samples. The training lasts for 10 epochs, with image size 224 and batch size 24, optimized by SGD. For the captioning model M_{CAP} , we employed the pre-trained VLM Openflamingo [4] as it could be applied to various vision-text tasks using a single model and is publicly open-sourced. For filtering in Sec. 3.3, we used the NLTK library [5], which is an NLP tool able to tokenize sentences and parse them in the Penn Treebank [42]. The filtering threshold is set to $T = 0.6$. One additional aspect to consider is that when dealing with both new and existing data sourced from the web, the weight of the L_{KDE} loss is reduced to 0.5. The code is available at <https://github.com/dota-109/Web-WILSS>.

Table 1: Comparison on single-step multi-class settings on PASCAL-VOC.

Method	Train	Memory	M	15-5						10-10					
				Disjoint			Overlapped			Disjoint			Overlapped		
				1-15	16-20	All	1-15	16-20	All	1-10	11-20	All	1-10	11-20	All
CAM† [52]	VOC	-	-	69.3	26.1	59.4	69.9	25.6	59.7	65.4	41.3	54.5	70.8	44.2	58.5
SEAM† [43]	VOC	-	-	71.0	33.1	62.7	68.3	31.8	60.4	65.1	53.5	60.6	67.5	55.4	62.7
SS† [3]	VOC	-	-	71.6	26.0	61.5	72.2	27.5	62.1	60.7	25.7	45.0	69.6	32.8	52.5
EPS† [20]	VOC	-	-	72.4	38.5	65.2	69.4	34.5	62.1	64.2	54.1	60.6	69.0	57.0	64.3
WILSON† [7]	VOC	-	-	73.6	43.8	67.3	74.2	41.7	67.2	64.5	<u>54.3</u>	60.8	70.4	57.1	65.0
RaSP [37]	VOC	-	-	-	-	-	76.2	47.0	70.0	-	-	-	72.3	<u>57.2</u>	<u>65.9</u>
FMWISS [49]	VOC	VOC	50	75.9	<u>50.8</u>	70.7	<u>78.4</u>	<u>54.5</u>	<u>73.3</u>	68.5	58.2	64.6	73.8	62.3	69.1
Ours	VOC	WEB	100	77.1	49.1	<u>71.1</u>	78.2	54.9	<u>73.3</u>	68.5	52.0	<u>61.7</u>	73.6	55.5	65.7
Ours	VOC	WEB	500	77.1	53.0	72.0	78.8	53.8	73.4	67.1	52.1	61.0	73.0	55.2	65.3
WILSON [7]	WEB	-	-	75.2	45.5	68.9	74.2	43.8	67.8	66.9	47.3	58.6	72.0	49.7	62.1
Ours	WEB	-	-	75.5	45.0	69.0	74.4	45.9	68.4	<u>68.3</u>	49.7	60.4	73.8	53.9	65.0
Ours	WEB	WEB	100	<u>76.9</u>	47.0	70.5	78.3	47.8	71.7	68.0	50.0	60.4	<u>73.7</u>	54.5	65.3
Ours	WEB	WEB	500	76.2	50.0	70.7	77.8	45.6	70.8	67.9	50.4	60.6	72.4	54.4	64.6

4.1 Experimental Comparison

Single-Step Multi-class Results The results for **15-5** and **10-10** on VOC (i.e., a single incremental step with 5 and 10 classes) are shown in Table 1. For comparison purposes, we consider all other methods within the WILSS framework using the VOC dataset for training, while we compared our approach with WILSON [7] also in the new setting, as it serves as the baseline framework. From the table, several key points can be highlighted:

1. Our rehearsal strategy, utilizing web data when training with the original dataset, outperforms all previous methods that do not store samples from the original dataset. In particular, in the 15-5 setting our approach outperforms all competitors and achieves an improvement of 5% in the disjoint setup and of 6% in the overlapped one w.r.t [7]. Furthermore, unlike FMWISS [49], we did not rely on pixel-level supervision from pre-trained models, but solely on image-level supervision; On the 10-10 setup competitors are closer to our approach, however except [49] that uses VOC supervision, we are the best in the disjoint setup and second best in the overlapped.
2. In scenarios where no data from VOC is available, leveraging web data leads to comparable results to the state-of-the-art (up to 73.4% in the 15-5 and 65.3% in the 10-10 setup), with further improvements observed when utilizing rehearsal data (in this case we get very close to the case where VOC data is used, especially in the disjoint setting);
3. Both of our strategies aimed at enhancing web data for old and new classes demonstrate improvements over the baseline [7];
4. We observed that employing a larger number of images (i.e., 500 as suggested in [27]) leads to better results when training with data from the original distribution, but yields worse results when learning from the web.

Table 2: Comparison on 60-20 setting on COCO-to-VOC. †: from [7].

Method	Train		Memory	COCO			VOC
	1-60	61-80	Source M	1-60	61-80	All	61-80
CAM† [52]	COCO	VOC	-	-	30.7	20.3	28.1
SEAM† [43]	COCO	VOC	-	-	31.2	28.2	30.5
SS† [3]	COCO	VOC	-	-	35.1	36.9	35.5
EPST† [20]	COCO	VOC	-	-	34.9	38.4	35.8
WILSON [7]	COCO	VOC	-	-	39.8	<u>41.0</u>	<u>40.6</u>
RaSP [37]	COCO	VOC	-	-	41.1	<u>41.0</u>	41.6
FMWISS [49]	COCO	VOC	COCO 50	39.9	44.7	41.6	63.6
Ours	COCO	VOC	WEB 100	<u>40.7</u>	33.0	39.3	<u>56.7</u>
WILSON [7]	COCO	WEB	-	-	<u>40.0</u>	<u>33.2</u>	<u>38.8</u>
Ours	COCO	WEB	-	-	40.2	35.6	39.6
WILSON	COCO	WEB	WEB 100	<u>40.6</u>	<u>28.1</u>	<u>38.0</u>	<u>41.3</u>
Ours	COCO	WEB	WEB 100	40.9	35.8	40.2	49.5

Table 3: Comparison on single-step multi-class settings on PASCAL-VOC.

Method	Train	Memory	M	15-1			10-1		
				1-15	16-20	All	1-10	11-20	All
WILSON [7]	VOC	-	-	0.0	2.3	0.6	0.0	0.2	0.1
RaSP [37]	VOC	-	-	17.7	0.9	13.2	2.0	0.7	1.3
RaSP [37]	VOC	VOC	100	63.3	28.3	56.0	38.9	30.9	36.9
WILSON [7]	VOC	ImageNet	100	75.7	32.9	65.9	66.8	34.9	<u>52.3</u>
RaSP [37]	VOC	ImageNet	100	75.7	35.2	66.6	66.8	39.1	54.4
WILSON	VOC	WEB	100	73.9	29.9	63.6	63.0	35.6	49.9
Ours	VOC	WEB	100	73.6	36.2	65.1	<u>63.9</u>	<u>39.2</u>	52.0
WILSON	WEB	WEB	100	74.4	41.6	<u>67.2</u>	61.0	28.3	42.9
Ours	WEB	WEB	100	<u>74.5</u>	41.9	67.4	63.7	42.4	54.4

Results on the COCO-to-VOC benchmark are shown in Tab. 2. Our replay strategy achieves a mIoU of 56.7% on VOC, outperforming exemplar-free competitors, second only to [49] which however stores COCO samples. We demonstrate learning from web data as a viable alternative to controlled data collection, outperforming models trained on original datasets. In settings using web data for the incremental step, we achieve accuracies on VOC around 50% without exposure to VOC data distribution. Incorporating web data for both old and new classes increases COCO mIoU, though this improvement does not fully transfer to VOC. However, results surpass WILSON trained with the same data on both datasets. Please note that our method focuses on obtaining web data matching the original distribution, i.e., aligned with COCO, for both old and new classes. Future investigations could explore optimal strategies to balance between matching the original distribution and generalization to unseen domains. Nonetheless, new web images allow incremental training, while old web images remain advantageous, particularly in multi-step single-class learning scenarios, where competitors necessitate the incorporation of old classes’ samples.

Multi-Step Single-class Results Results for the **15-1** and **10-1** overlapped settings are in Table 3. Notably, among competitors, only [37] addresses these settings. However, it is noteworthy that [37] still relied on handcrafted datasets (i.e., VOC or ImageNet), rather than unsupervised data. We can observe that:

1. Leveraging web data as rehearsal enables WILSON to learn within the multi-step single-class scenario;
2. Our approach outperforms the standard strategy based on [7] of around 1.5% in the 15-1 and 2% in the 10-10 setup;
3. Although remembering past classes from human polished and large curated datasets is still the best option, our method manages to narrow the performance gap getting very close to competitors exploiting ImageNet rehearsal;
4. The most notable advantage of our approach is demonstrated in the final setting of Table 3, where learning new knowledge from the web achieves superior performance compared to training from the original dataset and




Fig. 4: Qualitative results in the overlapped 10-1 setting.

Table 4: Ablation study on the web learning strategies: Pascal-VOC Overlapped. **Table 5:** Ablation study on the rehearsal strategies on Pascal-VOC Overlapped.

Learning from Web		10-10			15-5		
Domain filtering	Caption Labeling	1-10	11-20	All	1-15	16-20	All
		72.0	49.7	62.1	74.2	43.8	67.8
✓		72.0	51.9	63.1	73.7	46.3	68.0
	✓	72.6	52.1	63.6	73.8	46.7	68.2
✓	✓	73.8	53.9	65.0	74.4	45.9	68.4

Rehearsal Strategy		10-10			15-5		
Caption query	Caption Filtering	1-10	11-20	All	1-15	16-20	All
		71.6	53.1	63.6	76.0	47.4	69.8
✓		72.2	55.1	64.9	77.6	53.9	72.6
	✓	71.5	52.9	63.5	78.8	47.8	71.9
✓	✓	73.0	55.2	65.3	78.7	53.7	73.4

allows us to outperform competitors in the 15-1 setting and to get exactly the same performances of [37] in the 10-1 but using web data instead of VOC and ImageNet for the competitor. This is very relevant as some classes in VOC are short in data, suggesting the potential viability of our setup.

Figure 4 shows some visual examples for the most challenging 10-1 setting. Notice how they confirm the numerical evaluation with our approach being the only one able to properly recognize the main objects in the scene while WILSON show a large number of wrong detections. The shape of the objects is good even if not extremely precise - yet recall that there is no segmentation ground truth in the incremental steps - and in any case much better than the main competitor.

Ablation Results In order to analyze the impact of our web rehearsal strategies we performed some ablation studies. Table 4 shows the impact of the two newly introduced modules for the web learning strategy proposed in this paper. The domain discriminator in the Fourier domain allows for a gain of around 1% in the 10-10 setup, while the new caption labeling strategy for the localizer allows for a bit larger improvement of 1.5% in the same setup. Notice how the two strategies are complementary and their combined usage leads to an improvement of around 3%.

Table 5 shows instead the impact of the caption querying and caption filtering proposed in this paper. Using just class names as queries we achieve a mIoU

of 63.6% in the overlapped 10-10 setting and 69.8% in the 15-5. Using the captions for web querying allows for improved performances in both 10-10 and 15-5 setting. Generating new captions to filter the images leads to an improvement in the 15-5 setup while the 10-10 performances remain roughly the same. Note that, when no queried caption is available, the filter strategy solely relies on the class name, reducing the strength of the text description. Finally, the potential for gathering more reliable data is achieved when combining the two strategies, leading to an improvement of 1.7% in the 10-10 scenario and 3.6% in the 15-5 scenario.

5 Limitations

The present framework relies on supervising through a foundation model, similar to other setups employing such models. While it is reasonable to assume that a large model encompasses knowledge of concepts connected to future classes one aims to learn, this assumption may not hold. Subsequent research could address this concern. A practical implementation of the system should download web data on the fly while the learning goes on, how to optimize the bandwidth usage and the download time requirements is also left as future research.

6 Conclusions

This paper introduces a novel approach to address incremental learning in semantic segmentation by leveraging web data instead of manually selected and annotated datasets. We examine both the existing WILSS setup from literature and a novel approach where supervised data is utilized only in the initial step, followed by learning driven solely by language cues and web images. Our method utilizes a VLM to extract meaningful captions for data samples, which are then used to guide web querying and select suitable web images. By integrating this with state-of-the-art continual learning strategies, we achieve remarkable performance, demonstrating the effectiveness of web data in replacing manually curated datasets in continual learning tasks. Future research will address the pointed out limitations and consider the extension of the approach to other segmentation architectures or to tasks behind semantic segmentation.

Acknowledgements: This work was partially supported by the National Natural Science Foundation of China (62371356), the European Union, Italian National Recovery, Resilience Plan of NextGenerationEU, partnership “Telecommunications of the Future” (PE00000001- “RESTART”), Fundamental Research Funds for the Central Universities of China (ZYTS24005) and the National Defense Basic Scientific Research Program of China (JCKY2021413B005).

References

1. Ahn, J., Kwak, S.: Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4981–4990 (2018)
2. Alayrac, J.B., Donahue, J., Luc, P., Miech, A., Barr, I., Hasson, Y., Lenc, K., Mensch, A., Millican, K., Reynolds, M., Ring, R., Rutherford, E., Cabi, S., Han, T., Gong, Z., Samangooei, S., Monteiro, M., Menick, J., Borgeaud, S., Brock, A., Nematzadeh, A., Sharifzadeh, S., Binkowski, M., Barreira, R., Vinyals, O., Zisserman, A., Simonyan, K.: Flamingo: a visual language model for few-shot learning. ArXiv **abs/2204.14198** (2022)
3. Araslanov, N., Roth, S.: Single-stage semantic segmentation from image labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4253–4262 (2020)
4. Awadalla, A., Gao, I., Gardner, J., Hessel, J., Hanafy, Y., Zhu, W., Marathe, K., Bitton, Y., Gadre, S., Sagawa, S., Jitsev, J., Kornblith, S., Koh, P.W., Ilharco, G., Wortsman, M., Schmidt, L.: Openflamingo: An open-source framework for training large autoregressive vision-language models. arXiv preprint arXiv:2308.01390 (2023)
5. Bird, S., Klein, E., Loper, E.: Natural language processing with Python: analyzing text with the natural language toolkit. " O'Reilly Media, Inc." (2009)
6. Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9650–9660 (2021)
7. Cermelli, F., Fontanel, D., Tavera, A., Ciccone, M., Caputo, B.: Incremental learning in semantic segmentation from image labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4371–4381 (2022)
8. Cermelli, F., Mancini, M., Bulò, S.R., Ricci, E., Caputo, B.: Modeling the background for incremental learning in semantic segmentation. In: CVPR (2020)
9. 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)
10. Chen, X., Gupta, A.: Webly supervised learning of convolutional networks. In: ICCV. pp. 1431–1439 (2015)
11. Dai, J., He, K., Sun, J.: Boxsup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In: Proceedings of the IEEE international conference on computer vision. pp. 1635–1643 (2015)
12. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
13. Douillard, A., Chen, Y., Dapogny, A., Cord, M.: Plop: Learning without forgetting for continual semantic segmentation. In: CVPR (2021)
14. Duan, H., Zhao, Y., Xiong, Y., Liu, W., Lin, D.: Omni-sourced webly-supervised learning for video recognition. In: European Conference on Computer Vision. pp. 670–688. Springer (2020)
15. Everingham, M., Van Gool, L., Williams, C.K., Winn, J., Zisserman, A.: The pascal visual object classes (voc) challenge. International journal of computer vision **88**, 303–338 (2010)
16. Fellbaum, C.: WordNet: An electronic lexical database. MIT press (1998)

17. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
18. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
19. Klingner, M., Bär, A., Donn, P., Fingscheidt, T.: Class-incremental learning for semantic segmentation re-using neither old data nor old labels. International Conference on Intelligent Transportation Systems (2020)
20. Lee, S., Lee, M., Lee, J., Shim, H.: Railroad is not a train: Saliency as pseudo-pixel supervision for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5495–5505 (2021)
21. Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: Gligen: Open-set grounded text-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22511–22521 (2023)
22. Lin, D., Dai, J., Jia, J., He, K., Sun, J.: Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3159–3167 (2016)
23. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13. pp. 740–755. Springer (2014)
24. Liu, C., Rizzoli, G., Barbato, F., Michieli, U., Niu, Y., Zanuttigh, P.: Recall+: Adversarial web-based replay for continual learning in semantic segmentation. arXiv preprint arXiv:2309.10479 (2023)
25. Lukasik, M., Bhojanapalli, S., Menon, A., Kumar, S.: Does label smoothing mitigate label noise? In: International Conference on Machine Learning. pp. 6448–6458. PMLR (2020)
26. Luo, A., Li, X., Yang, F., Jiao, Z., Cheng, H.: Webly-supervised learning for salient object detection. Pattern Recognition **103**, 107308 (2020)
27. 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)
28. McEver, R.A., Manjunath, B.: Pcams: Weakly supervised semantic segmentation using point supervision. arXiv preprint arXiv:2007.05615 (2020)
29. Michieli, U., Zanuttigh, P.: Incremental Learning Techniques for Semantic Segmentation (2019)
30. 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)
31. Niu, L., Veeraraghavan, A., Sabharwal, A.: Webly supervised learning meets zero-shot learning: A hybrid approach for fine-grained classification. In: CVPR. pp. 7171–7180 (2018)
32. Phan, M.H., Phung, S.L., Tran-Thanh, L., Bouzerdoun, 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)

33. Qin, Z., Xu, Z., Zhou, Y., Zheng, Z., Cheng, Z., Tang, H., Shang, L., Sun, B., Peng, X., Timofte, R., Yao, H., Wang, K., You, Y.: Dataset growth (2024), <https://arxiv.org/abs/2405.18347>
34. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., Krueger, G., Sutskever, I.: Learning transferable visual models from natural language supervision. In: Meila, M., Zhang, T. (eds.) Proceedings of the 38th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/radford21a.html>
35. Ravi, S., Chinchure, A., Sigal, L., Liao, R., Shwartz, V.: Vlc-bert: visual question answering with contextualized commonsense knowledge. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 1155–1165 (2023)
36. Rizzoli, G., Shenaj, D., Zanuttigh, P.: Source-free domain adaptation for rgb-d semantic segmentation with vision transformers. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 615–624 (2024)
37. Roy, S., Volpi, R., Csurka, G., Larlus, D.: Rasp: Relation-aware semantic prior for weakly supervised incremental segmentation. arXiv preprint arXiv:2305.19879 (2023)
38. Shenaj, D., Rizzoli, G., Zanuttigh, P.: Federated learning in computer vision. IEEE Access (2023)
39. Song, C., Huang, Y., Ouyang, W., Wang, L.: Box-driven class-wise region masking and filling rate guided loss for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3136–3145 (2019)
40. Sun, H., He, X., Peng, Y.: Hcl: Hierarchical consistency learning for webly supervised fine-grained recognition. IEEE Transactions on Multimedia (2023)
41. Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International conference on machine learning. pp. 6105–6114. PMLR (2019)
42. Taylor, A., Marcus, M., Santorini, B.: The penn treebank: an overview. Treebanks: Building and using parsed corpora pp. 5–22 (2003)
43. Wang, Y., Zhang, J., Kan, M., Shan, S., Chen, X.: Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 12275–12284 (2020)
44. Wu, Z., Shen, C., Van Den Hengel, A.: Wider or deeper: Revisiting the resnet model for visual recognition. Pattern Recognition **90**, 119–133 (2019)
45. Xiao, J.W., Zhang, C.B., Feng, J., Liu, X., van de Weijer, J., Cheng, M.M.: End-points weight fusion for class incremental semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7204–7213 (2023)
46. 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 (2022)
47. Yang, J., Feng, L., Chen, W., Yan, X., Zheng, H., Luo, P., Zhang, W.: Webly supervised image classification with self-contained confidence. In: European Conference on Computer Vision. pp. 779–795. Springer (2020)
48. Yang, Y., Soatto, S.: Fda: Fourier domain adaptation for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4085–4095 (2020)

49. Yu, C., Zhou, Q., Li, J., Yuan, J., Wang, Z., Wang, F.: Foundation model drives weakly incremental learning for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 23685–23694 (2023)
50. 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)
51. 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)
52. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A.: Learning deep features for discriminative localization. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2921–2929 (2016)
53. Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from CLIP. In: European Conference on Computer Vision. pp. 696–712. Springer (2022)