

GAPS: Few-Shot Incremental Semantic Segmentation via Guided Copy-Paste Synthesis

Ri-Zhao Qiu¹, Peiyi Chen^{1*}, Wangzhe Sun², Yu-Xiong Wang¹, Kris Hauser¹

¹University of Illinois Urbana-Champaign

²Vanderbilt University

{rizhaoq2, peiyic2, yxw, kkhauser}@illinois.edu wangzhe.sun@vanderbilt.edu

Abstract

Few-shot incremental segmentation is the task of updating a segmentation model, as novel classes are introduced online over time with a small number of training images. Although incremental segmentation methods exist in the literature, they tend to fall short in the few-shot regime and when given partially-annotated training images, where only the novel class is segmented. This paper proposes a data synthesizer, Guided copy-And-Paste Synthesis (GAPS), that improves the performance of few-shot incremental segmentation in a model-agnostic fashion. Despite the great success of copy-paste synthesis in conventional offline visual recognition, we demonstrate substantially degraded performance of its naïve extension in our online scenario, due to newly encountered challenges. To this end, GAPS (i) addresses the partial-annotation problem by leveraging copy-paste to generate fully-labeled data for training, (ii) helps augment the few images of novel objects by introducing a guided sampling process, and (iii) mitigates catastrophic forgetting by employing a diverse memory-replay buffer. Compared to existing state-of-the-art methods, GAPS dramatically boosts the novel IoU of baseline methods on established few-shot incremental segmentation benchmarks by up to 80%. More notably, GAPS maintains good performance in even more impoverished annotation settings, where only single instances of novel objects are annotated.

1. Introduction

Incremental segmentation is an important capability for open-world AI systems. For example, consider a house-keeping robot that has been trained to segment common household objects, but once deployed in a user’s home it encounters a previously unseen type of furniture. For such practical applications, incremental segmentation would be

capable of expanding the set of recognized classes to contain the new object. There are a few desired properties of incremental segmentation algorithms to operate under these scenarios. First of all, the algorithm should be equipped with **few-shot learning capability**, which means that the algorithm can benefit from as few as one image provided by a user rather than requiring hundreds of images annotated offline by professional annotators. Second, providing full segmentation annotation of an image is time-consuming. To avoid causing substantial burdens for untrained users, the algorithm needs to be trainable with **partially-annotated** images where only novel classes are annotated.

A few attempts have been made by recent works [3, 5, 7, 28, 30] on *non-few-shot* incremental segmentation to investigate learning with partially-annotated images, which is termed *semantic background shift* [3]. Background shift describes a challenge unique where classes that are not in the current learning step are assigned ‘background’ labels, which prohibits direct end-to-end training. Recent work uses either modified loss [3, 30] or pseudo-labeling [5, 7, 28] as *proxies* to train on partially-annotated images. However, although these proxying methods demonstrate good performance under the non-few-shot settings, they rely on rich annotations and fall short when only a limited amount of data is presented to the model, due to a lack of diversity of data. An even more restrictive setting occurs when users label only a single instance of the novel class, where these methods fall short, due to the training containing non-annotated instances of the novel class.

To address the aforementioned challenges, we propose GAPS (Guided copy-And-Paste Synthesis), which improves the training of incremental segmentation models by synthesizing fully-annotated images from partially-annotated examples. It is *model-agnostic*, and can be inserted as a plug-and-play module into different incremental learning algorithms, e.g., standard fine-tuning or PIFS [4]. Copy-paste generates diverse training data to boost performance under few-shot settings, enables the model to learn with partially-annotated images with as few as one anno-

*: Work done while at UIUC

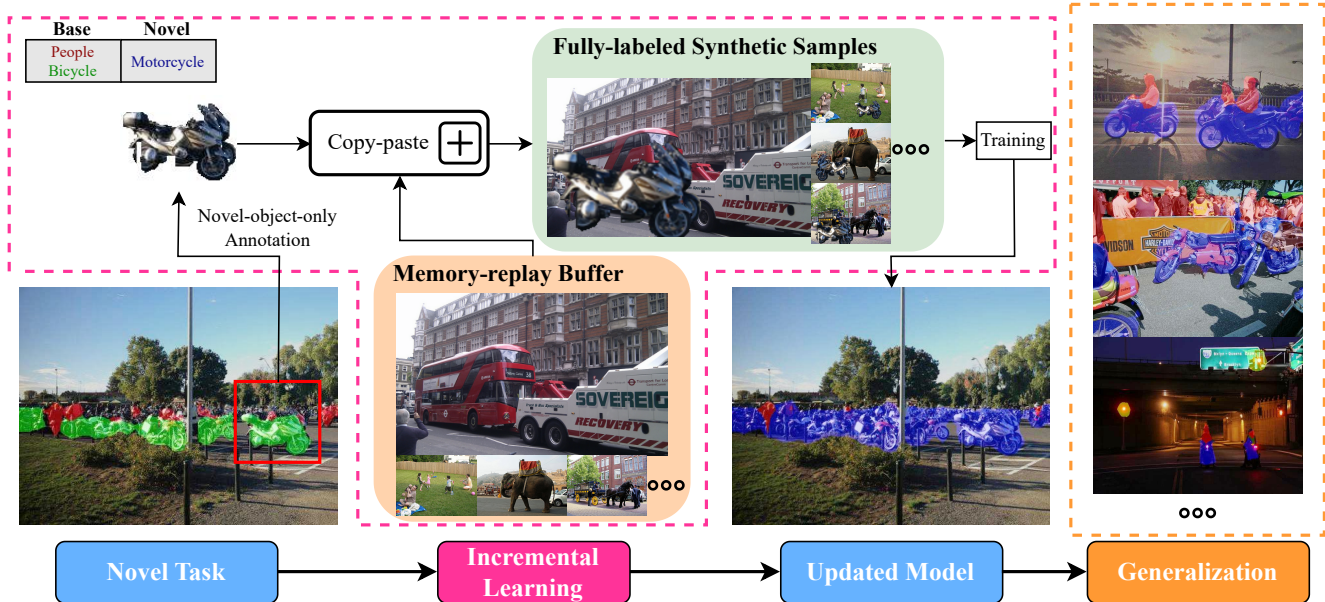


Figure 1. Our proposed method utilizes guided copy-paste augmentation to synthesize diverse training data, using as few as **one single novel instance** for training. For example, the model encounters an image of many motorcycles, which is novel to the model. As a result, the model incorrectly assigns learned bicycle labels to these pixels and therefore needs to be updated. Our proposed method can adapt to the novel motorcycle class with an annotation of a single motorcycle, which can be efficiently annotated; whereas previous work [3, 4] require time-consuming annotation of all instances of motorcycles or even the entire image. Best view in color.

tated novel instance out of many novel instances in an image (e.g., as illustrated at the lower left part of Fig. 1), which is a *stricter* setting than semantic background shift [3].

To the best of our knowledge, we are the *first* to introduce copy-paste as a synthesis technique to create a diverse data source for few-shot incremental segmentation. Although copy-paste [12] has been shown to be an effective data augmentation technique for offline visual recognition tasks, we identify new key technical challenges to adapting it to few-shot incremental settings. First, how should the synthesizer pick representative samples from the base dataset to construct a *diverse* pool of fully-annotated base scenes? Second, given the constructed pool of fully-annotated images, how should it select the most *suitable* base images to be pasted on? Third, after an informative image is selected, from what distribution should it sample current and previously learned novel objects to *balance* sample frequency and avoid over-sampling or under-sampling? Our GAPS method differs from a naïve (e.g., uniform random sampling) copy-paste process by a *guided* strategy that considers diversity of the memory-replay buffer, imbalanced class frequencies between base classes and novel classes, and contextual similarity of images.

In summary, our contributions are as follow:

1. We are the first to introduce copy-paste as a synthesis technique to address partially-labeled images for incremental segmentation.

2. To address the gaps between copy-paste under the offline setting as an augmentation technique and under the online setting as a synthesis technique, we design a guided copy-paste process that improves the distribution of synthesized images by enforcing diversity of the memory-replay buffer, exploiting contextual information, and balancing class frequencies.
3. The proposed GAPS technique consistently boosts the performance of a variety of incremental learning algorithms from simple fine-tuning to sophisticated state-of-the-arts under the few-shot setting. Furthermore, we demonstrate the strength of GAPS to cope with a more challenging task setting where only one instance out of many novel instances in an image is annotated, which highlights copy-paste as a better alternative to pseudo-labeling or modified loss for practical incremental segmentation applications.

2. Related Work

Incremental Learning for Semantic Segmentation. It is known that many learning-based models suffer from catastrophic forgetting [19], a phenomenon that causes models to perform significantly worse on old tasks when they are fine-tuned to adapt to new tasks. *Incremental learning* studies how to enable models to adapt to new classes while mitigating catastrophic forgetting without accessing

the old dataset or full-scale re-training. This problem has been studied extensively in image classification [1, 2, 16–18, 22, 29]; whilst relatively less work have been done to study incremental learning under the task setting of semantic segmentation [3, 5, 7, 20, 30]. Noticeably, a few attempts have been made by recent work to address the semantic background shift problem proposed by [3] via either pseudo-labeling [5, 7, 28] or modified loss [3, 30] to train on partially-annotated images of novel classes. However, existing work relies on rich annotations and tends to fail when only a limited amount of data is available. In contrast, our work enables incremental segmentation learning with *few data* via a guided copy-paste process, which demonstrates promising performance under the few-shot and more impoverished single-instance setting. Furthermore, GAPS is a *model-agnostic data pre-processor*, which is orthogonal to incremental learning techniques such as regularization [16].

Few-Shot Semantic Segmentation. *Few-shot semantic segmentation* methods predict segmentation masks of novel classes using only a few training examples of the novel class. Many meta-learning-based methods [23, 26, 27, 31] and even specialized datasets [15] have been proposed to address such a problem. However, few-shot semantic segmentation methods produce novel-class-only binary foreground-background segmentation. In comparison, our proposed method works in a more challenging and realistic setting where both base classes and novel classes need to be segmented.

Few-Shot Incremental Segmentation. While there are many works in few-shot incremental image classification [6, 24], relatively fewer works have been done to investigate few-shot incremental segmentation [4, 11, 25]. [25] designs a meta-learning-based classifier that adjusts learned prototypes by modeling interaction between base classes and incoming novel class. Unlike [25], which only performs a single update of weights in the classifier, PIFS [4] apply regularization techniques to allow fine-tuning of the entire network, achieving state-of-the-art result in few-shot incremental semantic segmentation. However, PIFS [4] is fine-tuned on only a small number of samples, which leads to sub-optimal performance due to overfitting. In addition, PIFS requires fully-annotated images as input, which hinders its potential for practical applications.

Copy-Paste Augmentation. Copy-and-paste is an augmentation technique that copies a subset of objects from one image and pastes onto the other image using their segmentation masks. Many works [8, 9, 12] have been done to investigate how copy-and-paste augmentation can help with various visual tasks. Dvornik *et al.* apply copy-paste augmentation in object detection by designing a neural network to consider context and guide copy-paste. However, the context guidance method proposed by Dvornik *et al.* can not be trivially applied to our application since it requires abun-

dant fully-annotated training data. More recently, Ghiasi *et al.* conduct extensive experiments to demonstrate the effectiveness of simple copy-paste in the instance segmentation problem. We extend the augmentation strategy from [12] and construct an intuitive baseline called Naïve copy-Paste Synthesis (NPS) to adapt it to our online task setting. However, as we will demonstrate in the ablation study, such naïve adaptation gives unsatisfactory performance in our task setting because of *gaps* between the static offline learning and continual online learning. In our work, we propose a series of techniques to guide the copy-paste synthesizer to address these gaps, whose effectiveness is evident from the significant improvement from NPS.

3. Method

Problem Setup. Let $\mathcal{X} \subset \mathbb{R}^{H \times W \times 3}$ be a set of RGB images with size $H \times W$, $\mathcal{C} \subset \mathbb{N}$ be a set of category labels, and $\mathcal{Y}^{\mathcal{C}} \subset \mathbb{R}^{H \times W \times |\mathcal{C}|}$ be a set of label masks (*i.e.*, per-pixel category labels in \mathcal{C}). In semantic segmentation, we aim to learn a model ϕ that maps an image $x \in \mathcal{X}$ to a segmentation mask $y \in \mathcal{Y}^{\mathcal{C}}$. Different from standard semantic segmentation, in few-shot incremental segmentation, \mathcal{C} is expanded over time through two stages. During the *base learning stage*, the model is provided with a base dataset $\mathcal{D}_0 = \{(x_i, y_i) | x_i \in \mathcal{X}, y_i \in \mathcal{Y}^{\mathcal{C}_0}\}$, where \mathcal{C}_0 is a set of classes in the base dataset. \mathcal{D}_0 generally contains many fully-annotated image-mask pairs and is used to train the model $\phi_0 : \mathcal{X} \rightarrow \mathcal{Y}^{\mathcal{C}_0}$ from scratch.

During the *incremental learning stage*, a sequence of tasks $\{D_1, D_2, \dots\}$ with novel categories is presented to the model, where $D_j = \{(x_i, y_i) | x_i \in \mathcal{X}, y_i \in \mathcal{Y}^{\mathcal{C}_j}\}$ and \mathcal{C}_j is a set of classes for task D_j . In few-shot learning, the size of the training sets for the novel tasks is small, *i.e.*, $|D_j| \ll |\mathcal{D}_0|$. After adapting to task D_j , the model is updated as $\phi_j : \mathcal{X} \rightarrow \mathcal{Y}^{\cup_{i=0, \dots, j} \mathcal{C}_i}$. The goal of incremental learning is to optimize the model performance jointly on both previous tasks and the current task. To enforce the partially-annotated image setting, we follow Cermelli *et al.* and assume that only novel classes are annotated, *i.e.*, $\mathcal{C}_i \cap \mathcal{C}_j = \emptyset$ for all $i \neq j$.

Method overview. Fig. 2 illustrates our proposed Guided copy-Paste Synthesis (GAPS) framework for few-shot incremental segmentation. It is a *generic and model-agnostic* data synthesis framework that generates fully-labeled scenes from partially-annotated images of novel objects as a preprocessor to the underlying segmentation model. After the standard base learning stage with base dataset \mathcal{D}_0 , we build a memory-replay buffer $\hat{\mathcal{D}}_0$ using an *diversity-guided exemplar selection strategy* (Section 3.2). During the incremental learning stage, fully-labeled samples are synthesized by copying from the masked novel objects in D_1, \dots, D_j and pasting onto base exemplars from the replay buffer $\hat{\mathcal{D}}_0$. The strategy by which we choose

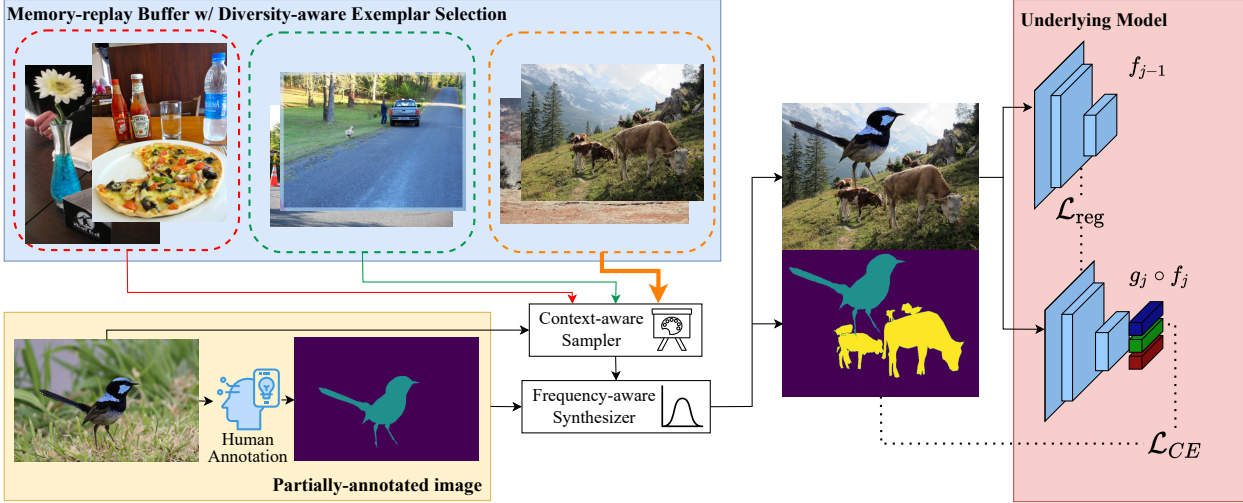


Figure 2. Overview of GAPS. During the incremental learning stage, GAPS takes in as few as one annotated instance of a single image. It is more probable for GAPS to select a scene contextually similar to the provided image from memory-replay buffer \hat{D}_0 . The image is then probabilistically pasted to generate synthetic fully-labeled scenes. Note that GAPS is model-agnostic, and here we use PIFS [4] as an example for the underlying segmentation model to illustrate how GAPS is applied as a pre-processor. Best seen in color.

base exemplars and novel segments is *context-guided* (Section 3.3) and *class-frequency-guided* (Section 3.4).

3.1. Few-Shot Incremental Segmentation Model

In principle, GAPS is model-agnostic, which means that it can work with many incremental segmentation models as a diverse data source to improve their performance. Here we adopt PIFS [4] as the main baseline underlying segmentation model for its state-of-the-art performance on few-shot incremental semantic segmentation and support for end-to-end training. The PIFS segmentation model ϕ is composed of a convolution-based feature extractor f and a per-pixel classification layer g using prototypical representation – g is configured to classify the pixels into n classes, so it is parameterized with prototypes $W = [w_1, w_2, \dots, w_n]$. Intuitively, f maps every pixel in an input image onto the unit hyper-sphere in a high-dimensional representation space. g then generates probability prediction by comparing cosine similarity of feature vectors with learned class prototypes w_i in the representation space and applying softmax of the resulting similarities. Following PIFS, given a previously unseen class $n + 1$ from task $D_{n+1} = \{(x_i, y_i) | x_i \in \mathcal{X}, y_i \in \mathcal{Y}^{n+1}\}$, instead of randomly initializing the prototype w_{n+1} , we apply the MAP (Maksed Average Pooling) function [4, 27] to estimate the prototype,

$$w_{n+1} = \frac{1}{|D_{n+1}|} \text{MAP}(x_i, y_i, n + 1), i \in D_{n+1} \quad (1)$$

$$= \frac{1}{|D_{n+1}|} \sum_{(x_i, y_i) \in D_{n+1}} \frac{M_{n+1}(y_i, j) \frac{f_n^j(x_i)}{\|f_n^j(x_i)\|}}{M_{n+1}(y_i, j)} \quad (2)$$

where $M_{n+1}(y_i, j)$ is a binary function that returns 1 if the j -th pixel in mask y_i is class $n + 1$, and 0 otherwise. $f_n^j(x_i)$ denotes the feature vector at the j -th pixel of $f_n(x_i)$.

In addition, we want to note that our re-implementation of PIFS [4] uses \mathcal{L}_2 regularization rather than the prototype distillation loss proposed by Cermelli *et al.* We found experimentally that when a diverse data source is used (i.e., our proposed GAPS), \mathcal{L}_2 regularization works better.

To be more precise, we construct a penalization term \mathcal{L}_{REG} to regularize the output before the classifier. For incremental learning task D_j with image-mask pairs (x, y) , we have

$$\mathcal{L}_{REG} = \|f_j(x) - f_{j-1}(x)\|_2. \quad (3)$$

The final training loss is given by

$$\mathcal{L}(x, y) = \mathcal{L}_{CE}(\phi_j(x), y) + \lambda \mathcal{L}_{REG}, \quad (4)$$

where \mathcal{L}_{CE} is either the standard cross-entropy loss or the modified cross-entropy loss from [3]. λ is a hyper-parameter used to weight the regularization loss. All other components are the same as in [4]. We denote our re-implementation of PIFS with \mathcal{L}_2 regularization loss as PIFS(\mathcal{L}_2).

3.2. Diversity-guided Exemplar Selection with Learned Prototypes

For methods with memory-replaying (e.g., SSUL [5]), GAPS can work directly on top of their constructed buffers with minimal modification. For other methods such as PIFS [4], we propose a diversity-guided exemplar selection process that builds a small yet diverse memory-replay buffer

\hat{D}_0 from \mathcal{D}_0 to mitigate catastrophic forgetting. Selecting diverse examples that are representative of the base dataset helps mitigate catastrophic forgetting, as suggested by [22]. Inspired by Bang et al. [1], we select samples distributed uniformly along a spectrum from easy to hard for diversity.

Here, we present an algorithm (Algorithm 1) to construct \hat{D}_0 by exploiting the Masked Average Pooling (MAP) function from [4]. Intuitively, we approximate the difficulty of every sample by their similarity between the estimated prototype with learned prototypes. Estimated prototypes that are close to the learned prototype are considered easy samples and vice versa. After building a list of base samples sorted by difficulties, we select samples from equally-spaced intervals to ensure samples of all difficulties are selected for diversity.

During the incremental learning stage, we select at most k samples for each novel class using the same algorithm to memorize novel classes. To maintain the size of the memory-replay buffer, we remove old samples from the memory-replay buffer but keep at least 80% of the samples to be fully-annotated samples, so that we have diverse base images for copy-pasting.

3.3. Context-guided Sampling

We hypothesize that synthesizing novel objects onto contextually consistent base images would result in an improved learning process. For example, a TV is more likely to appear in an apartment rather than in the middle of traffic on streets, and thus a TV object should more likely be pasted onto an image of another apartment rather than an outdoor landscape. To guide the copy-pasting process, we

design a context-guided sampling algorithm to select images from \hat{D}_0 that are contextually similar to the provided partially-labeled images.

One way of estimating pairwise contextual similarities between two images is to design a mapping $h : \mathcal{X} \rightarrow \mathbb{R}^m$ that maps an image into a metric space, where the metrics serve as a proxy of the contextual similarity between two images. Here, in GAPS, we extract the knowledge of the learned feature extractor. In incremental learning task D_i , scene embedding of image i is estimated by,

$$h_i = \frac{\text{GAP}(f_{i-1}(h_i))}{\|\text{GAP}(f_{i-1}(h_i))\|_2} \quad (5)$$

where $\text{GAP} : \mathbb{R}^{B \times C \times H \times W} \rightarrow \mathbb{R}^{B \times C}$ is the commonly used global average pooling function. Two scene embedding vectors h_i and h_j can then be compared by cosine similarity.

To find contextually similar base images to each novel image, we evaluate the cosine similarity of the novel image to each of the examples in \hat{D}_0 , and construct a contextually similar subset \mathbb{S} with $|\hat{D}_0|/10$ most contextually similar examples (If interested, visualization of query samples and their contextually-similar counterparts are included in Fig. 4 in the supplementary material). When there are multiple novel images, we take a union of selected examples. To allow other base scenes to be sampled to mitigate catastrophic forgetting, we sample from \mathbb{S} with a probability of α , and sample from \hat{D}_0 with a probability of $1 - \alpha$, where α is a hyperparameter set to 0.9 in our implementation. Note that we only need to compute scene embedding once for every image in \hat{D}_0 and incoming partially-annotated images. Hence, the context-guided sampling algorithm poses only minor computational overhead to GAPS.

3.4. Class-frequency-guided Probabilistic Synthesis

Now the final question is, given a fully-annotated image x_B and an image of a novel object x_N , how frequent should we apply copy-paste? There is a trade-off between oversampling and undersampling. As one extreme, one can follow [12] and always apply copy-paste augmentation to paste novel objects onto every base image. However, this will lead to oversampling of novel categories in the current task, which we found to hurt the performance of existing classes. On the other hand, rarely pasting novel instances would lead to undersampling of the novel class. Therefore, to guide copy-paste in the online setting, we design a synthesis strategy called vRFS based on RFS (Repeat Factor Sampling) described by [13] to perform synthesis.

To apply vRFS, we first need to compute category-wise sampling factor r_c for every c as in RFS. If $c \in \mathcal{C}_0$, we set $r_c = 1$ as since during the construction of \hat{D}_0 we already consider class balance by class-wise uniformly sampling. If $c \in \mathcal{C}_j$ with $j \geq 1$, we first compute its class

Algorithm 1 Construct Memory-replay Buffer

Require: number of exemplars n
 $k \leftarrow \text{FLOOR}(n/|Y_0|)$ // Sample per class
for c from 1 to $|Y_0|$ **do**
 $S_c \leftarrow \{(x_i, y_i) \in D_0, c \in y_i\}$
 for $(x_i, y_i) \in S_c$ **do**
 $p_i \leftarrow \text{MAP}(x_i, y_i, c)$ // Pred. Proto.
 $s_i \leftarrow \text{COSINESIMILARITY}(p_i, w_c)$
 end for
 Sort S_c by similarity score s_i
 $ES_c \leftarrow \{\}$ // final exemplar set of class c
 for $j = 1, 2, \dots, k$ **do**
 $L_{idx} \leftarrow j \cdot |S_c|/k$
 $U_{idx} \leftarrow \text{MIN}(L_{idx} + |S_c|/k, |S_c|)$
 $(x, y) \leftarrow \text{SAMPLE}(S_c[L_{idx} : U_{idx}])$
 $ES_c \leftarrow ES_c \cup (x, y)$
 end for
end for
 $\hat{D}_0 \leftarrow \text{UNIFORMSAMPLE}(\bigcup_{i=1, \dots, |Y_0|} ES_i, n)$

frequency by $f_c = nShot/|\hat{D}_0|$, where $nShot$ denotes the number of images in D_j with at least one pixel of c . Then, the category-wise sampling factor for c is given by $r_c = \text{MAX}(1, \sqrt{t/f_c})$. Note that in [13], t is chosen as a hyperparameter to be tuned. However, we empirically found that setting t to be the multiplicative inverse of total number of classes, or $t = 1/|\cup_{1,\dots,j} Y_j|$, is enough to yield stable results across different datasets and under different few-shot settings. This eliminates the need to search a hyperparameter for different settings and make our proposed method more robust towards different task settings.

During the synthesis process, we first randomly select a novel class c_N from \mathcal{C}_j , and another class c_o from $\cup_{1,\dots,j} \mathcal{C}_j \setminus \{c_N\}$. We first decide if c_o should be pasted onto x_B . To apply vRFS resampling, we hallucinate two *virtual* samples: in the first sample where copy-paste would not be applied, the image-level sampling factor is given by 1. In the second sample where copy-paste synthesis were to be performed, we would obtain a sample with image-level sampling factor of $r_i = \text{MAX}_{c \in i} r_c = r_{c_o}$. Thus, the probability to synthesize class c_o onto x_B is given by $r_{c_o}/(1 + r_{c_o})$. We then repeat the process for the novel class c_N . Note that vRFS synthesis is applied twice for every class, resulting in up to two pasted instances of c_N in the final image.

4. Experiments

4.1. Datasets

We follow literature in few-shot segmentation and few-shot incremental segmentation [4, 21, 23, 25] and evaluate our model on the PASCAL-5ⁱ dataset [23] and the COCO-20ⁱ dataset [21]. PASCAL-5ⁱ is artificially built from the PASCAL VOC 2012 Semantic Segmentation dataset [10] with additional annotations from the SBD [14] dataset. The original VOC segmentation dataset provides segmentation annotations for 20 object categories. The PASCAL-5ⁱ dataset manually splits the original dataset into 4 folds for cross-validation. For each fold, 5 categories are selected as novel categories, while the remaining 15 categories are regarded as base categories. In our experiments, images containing at least one pixel of the novel categories are excluded from the base dataset. The construction of the COCO-20ⁱ dataset handles the 80 thing classes in COCO in a similar manner, where the dataset is split into 4 folds and each fold contains 20 categories. The rest of the process to construct the base dataset and the novel dataset in COCO-20ⁱ is same as the PASCAL-5ⁱ dataset.

4.2. Evaluation Protocols

In the base learning stage, the model is trained using the entire base dataset. In incremental learning stages, sequences of tasks are presented to the model. We use the

same evaluation protocol as proposed in [4] for fair comparisons, where 5 incremental learning tasks are used for PASCAL-5ⁱ and each task contains 1 class from the novel split. On the COCO-20ⁱ dataset, there are 4 incremental learning tasks, and each task contains 5 classes from the novel split.

We evaluate the performance of the model on the entire validation set of the corresponding dataset after every step. For fair comparisons with our main baseline PIFS [4], we average results across different steps and exclude completely unseen classes from evaluation of current step. We use three different metrics to evaluate the performance of the model: mean Intersection-over-Union (mIoU) over base categories, mIoU over novel categories, and harmonic mean of the base mIoU and the novel mIoU. Unless otherwise noted, the numbers are computed by averaging results over splits in a cross-validating fashion.

To average out randomness due to few training samples, we also average results over multiple runs with different set of few-shot training samples. For experiments on splits on PASCAL-5ⁱ, we found that averaging results over 10 runs with randomly sampled few-shot novel images yields stable results. For COCO-20ⁱ, we found that averaging results over 5 runs is enough to yield stable results.

4.3. Main Results

In Table 1, we evaluate various incremental segmentation methods on the PASCAL-5ⁱ dataset and the COCO-20ⁱ dataset, and combine them with GAPS where appropriate.

Baselines. There are two main baselines we are comparing to. The first one is SSUL [5], which is the state-of-the-art method in non-few-shot incremental segmentation. The second one is PIFS [4], for it is the state-of-the-art method in few-shot incremental semantic segmentation. We also report the performance of two variants of *PIFS*: one is our re-implementation $\text{PIFS}(\mathcal{L}_2)$ described in Sec. 3.1. The other variant is $\text{PIFS}(\mathcal{L}_2)+\text{MEM}$, which uses a memory-replay buffer of the same size of *GAPS*. To handle partially-labeled images, we follow SSUL and perform pseudo-labeling on partially annotated samples before adding them to memory for $\text{PIFS}(\mathcal{L}_2)+\text{MEM}$. In addition, we also evaluate simple fine-tuning and MiB [3].

GAPS consistently increases performance under few-shot settings. Methods combined with our proposed data source, GAPS, consistently outperform their un-augmented counterpart on both the base and novel categories' performance. It is worth noting that GAPS substantially boosts the performance of methods that originally require fully-annotated training images (i.e., fine-tuning and PIFS), despite using only partially-annotated images now. Even for methods that do not carry out end-to-end training and update only the classifier (i.e., SSUL), GAPS still steadily increases performance on novel categories. Compared to our

METHOD	BASE	NOVEL	HM	BASE	NOVEL	HM
	PASCAL-5 ¹ 1-SHOT			PASCAL-5 ¹ 5-SHOT		
MIB [3]	43.9	2.6	4.9	60.9	5.8	10.5
FINETUNE*	47.2	3.9	7.2	58.7	7.7	13.6
FINETUNE+GAPS	64.2(+17.0)	16.2(+12.3)	25.9(+18.7)	66.8(+8.1)	38.1(+30.4)	48.5(+34.9)
SSUL [5]	73.9	16.4	26.8	74.8	27.8	40.5
SSUL+GAPS	74.0(+0.1)	19.9(+3.5)	31.3(+4.5)	74.9(+0.1)	30.0(+2.2)	42.8(+2.3)
PIFS* [4]	64.1	16.9	26.7	64.5	27.5	38.6
PIFS(\mathcal{L}_2)* ¹	64.6	19.7	30.2	57.7	24.5	34.4
PIFS(\mathcal{L}_2)+MEM	68.1	17.4	27.8	69.3	39.7	50.5
PIFS(\mathcal{L}_2)+GAPS	66.8(+2.2)	23.6(+3.9)	34.9(+4.7)	68.2(+10.5)	43.9(+19.4)	53.4(+19.0)
COCO-20 ¹ 1-SHOT			COCO-20 ¹ 5-SHOT			
MIB [3]	40.4	3.1	5.8	43.8	11.5	18.2
FINETUNE*	38.5	4.8	8.5	39.5	11.5	17.8
FINETUNE+GAPS	44.5(+6.0)	11.0(+6.2)	17.7(+9.5)	46.4(+6.9)	24.9(+13.4)	32.4(+14.6)
SSUL [5]	51.0	6.3	11.3	51.6	15.0	23.2
SSUL+GAPS	50.8(+0.2)	11.0(+4.7)	18.1(+6.8)	51.9(+0.3)	17.1(+2.1)	25.7(+2.5)
PIFS* [4]	40.4	10.4	16.5	41.1	18.3	25.3
PIFS(\mathcal{L}_2)* ¹	45.7	10.3	16.8	46.2	20.2	28.1
PIFS(\mathcal{L}_2)+MEM	47.8	11.2	18.1	46.8	22.0	29.9
PIFS(\mathcal{L}_2)+GAPS	46.8(+1.1)	12.7(+2.4)	20.0(+3.2)	49.1(+2.9)	25.8(+5.6)	33.8(+5.7)

Table 1. Methods augmented with our proposed GAPS consistently outperform their un-augmented counterparts in terms of IoU across different few-shot settings on COCO-20¹ and PASCAL-5¹. Methods noted with * are privileged and use fully-annotated images, others use images with novel-class-only partial annotation. ¹: our re-implementation using \mathcal{L}_2 regularization. Highest results are colored **red** and the second highest results are colored **blue**. HM stands for harmonic mean. (Best view in color).

implemented variant PIFS(\mathcal{L}_2)+MEM, our method demonstrates considerable relative improvement on novel categories but performs slightly worse on base categories due to the introduction of pseudo-labeling in PIFS(\mathcal{L}_2)+MEM, which has been shown in previous work [5] to have regularization effects on base classes.

4.4. Ablation Study

In Table 2, we ablate guidance designs in GAPS to illustrate how different types of guidance contribute to the final incremental learning performance than naïve copy-paste synthesis. Though GAPS is a synthesis method that applies to many base learning algorithms, due to the highest harmonic mean of PIFS(\mathcal{L}_2)+GAPS on all settings, here we use PIFS(\mathcal{L}_2)+GAPS for the ablation study.

Our diversity-guided exemplar selection method consistently increases performance on base categories, which suggests that it is capable of choosing diverse samples to construct a representative memory-replay buffer and mitigate catastrophic forgetting to improve performance on base classes after sequential adaptations.

Context-guided sampling steadily improves performance on novel classes, which is consistent with findings in previous work [8] that background context is an important factor to consider in copy-paste synthesis. (Example visualization is available in Fig. 4 in the supplementary ma-

terial).

Frequency-guided probabilistic synthesis boosts results on novel classes. On the other hand, it does not influence the performance of base categories in a statistically significant manner. We take a closer look at step-wise performance and found that the reason is due to unguided copy-paste’s oversampling of novel classes that are being adapted, and forgetting of classes learned in the previous incremental learning stage and not in the memory-replay buffer.

4.5. More Challenging Single-Instance Experiment

Though the semantic background shift proposed by [3] relaxes the requirement to provide full segmentation annotations, it still requires *all novel instances* in images to be annotated, which can be time-consuming to obtain in cluttered scenes and hinder potential applications (e.g., the cluttered motorcycle image in Figure 1). Here we consider a more challenging task setting, which we term *single-instance incremental learning*. Namely, for training images provided in incremental learning stages, if there are multiple instances of a novel class in the image, we assume that only *one instance* will be annotated for the model.

To simulate this setting, we use the instance-level segmentation annotation provided by the COCO dataset to enforce only annotation of one novel instance in every im-

MEM	COPY-PASTE	F-GUIDE	D-GUIDE	C-GUIDE	BASE	NOVEL	HM
—	—*	—	—	—	46.2(± 0.3)	20.2(± 0.7)	28.1(± 0.3)
✓	—*	—	—	—	49.3(± 0.2)	19.4(± 0.7)	27.9(± 0.3)
✓	✓	—	—	—	47.0(± 0.2)	19.8(± 0.6)	27.8(± 0.3)
✓	✓	✓	—	—	47.2(± 0.2)	25.2(± 0.6)	32.9(± 0.3)
✓	✓	✓	✓	—	48.2(± 0.2)	25.0(± 0.7)	32.9(± 0.3)
✓	✓	✓	✓	✓	49.1(± 0.2)	25.8(± 0.6)	33.8(± 0.3)

Table 2. Ablation study of components in GAPS on PIFS(\mathcal{L}_2) on the COCO-20ⁱ dataset under 5-shot setting. Note that when only combined with the memory-replay buffer, the base IoU is higher because model has access to additional full annotations. When diversity guidance (D-guide) is disabled, \hat{D}_0 consists of random examples from the base dataset, resulting in worse base performance. When context guidance (C-guide) is disabled, a base image is uniformly sampled. When frequency guidance (F-guide) is disabled, a novel instance is sampled uniformly and is always pasted onto the base image. 95% confidence intervals over 20 trials are reported assuming that trial results are normally distributed. *: privileged. use fully-annotated masks when copy-paste is turned off.

METHOD	BASE	NOVEL	HM	BASE	NOVEL	HM
	ALL INSTANCES			SINGLE-INSTANCE ONLY		
PIFS(\mathcal{L}_2) [†]	46.2	20.2	28.1	46.1 (-0.2%)	17.6 (-12.9%)	25.4 (-9.6%)
PIFS(\mathcal{L}_2)+GAPS	49.1	25.8	33.8	49.2 (+0.2%)	25.1 (-2.7%)	33.2 (-1.8%)

Table 3. Performance of pseudo-labeling methods and GAPS under the more challenging single-instance learning setting on COCO-20ⁱ 5-shot. Only 1 novel instance out of potentially many instances in individual training images is annotated. The pseudo-labeling baseline, PIFS(\mathcal{L}_2)[†], yields substantially worse performance; whereas PIFS(\mathcal{L}_2)+GAPS has only minor performance decreases.

age is available to the model. Since state-of-the-art incremental segmentation approaches use pseudo-labeling [5], we design a method PIFS(\mathcal{L}_2)[†], which simulates combining PIFS(\mathcal{L}_2) with pseudo-labeling to cope with partially-annotated sample. Here we allow PIFS(\mathcal{L}_2)[†] to be privileged and have access to additional information – the annotation of other non-novel background pixels – to simulate an oracle pseudo-labeling model which perfectly segments learned classes but recognize unseen novel classes as background.

The results are given in Table 3. We can observe that the pseudo-labeling baseline, PIFS(\mathcal{L}_2)[†], yields substantially worse performance when the model receives single-instance annotations despite having privileged access. We reason this is due to noisy labels generated by the pseudo-labeling process, where novel instances are incorrectly labeled as background. On the contrary, PIFS(\mathcal{L}_2)+GAPS shows only a minor performance decrease with single instances. This highlights the potential of copy-paste synthesis as an alternative to the existing pseudo-labeling paradigm to cope with the more realistic single-instance setting and robustness against false negative annotations.

More results and visualization. Due to space limits, we kindly refer readers to the supplementary material for more quantitative results that justify our design choices such as memory-replay buffer strategies and vRFS over other baselines. In addition, visualized qualitative results of sample segmentation and contextually-similar set construction can also be found in the supplementary material.

5. Conclusion and Discussion

In this paper, we demonstrate how the judicious use of copy-paste dramatically boosts the performance of incremental segmentation methods under the few-shot setting and enables learning with partially-annotated images. Our proposed GAPS technique selects representative exemplars in the memory-replay buffer and addresses the problems of class imbalance and contextual mismatch in synthesis.

In future work, we are interested in further application of copy-paste as a synthesis technique to cope with the background shifting problem for incremental segmentation. We believe that copy-paste can serve as a promising alternative to pseudo-labeling and modified loss to enable learning on partially-annotated images. We also believe that further optimizing exemplar selection and sampling strategies can lead to better guidance and lead to even better performance. Finally, the ability to learn with as few as one annotated instance in an image raises several intriguing possibilities. For example, integrating our work with learning-based interactive segmentation will enable human operators to continually and adaptively teach novel classes and correct failed predictions. This workflow has many interesting applications such as robot teleoperation where sparse annotations are preferable. Learning with weaker annotations, like bounding boxes or single clicks, and even self-supervision, is also an interesting direction to explore.

References

- [1] Jihwan Bang, Heesu Kim, YoungJoon Yoo, Jung-Woo Ha, and Jonghyun Choi. Rainbow memory: Continual learning with a memory of diverse samples. In *CVPR*, 2021. 3, 5
- [2] Francisco M Castro, Manuel J Marín-Jiménez, Nicolás Guil, Cordelia Schmid, and Karteek Alahari. End-to-end incremental learning. In *ECCV*, 2018. 3
- [3] Fabio Cermelli, Massimiliano Mancini, Samuel Rota Bulo, Elisa Ricci, and Barbara Caputo. Modeling the background for incremental learning in semantic segmentation. In *CVPR*, 2020. 1, 2, 3, 4, 6, 7
- [4] Fabio Cermelli, Massimiliano Mancini, Yongqin Xian, Zeynep Akata, and Barbara Caputo. Prototype-based incremental few-shot semantic segmentation. *arXiv preprint arXiv:2012.01415*, 2021. 1, 2, 3, 4, 5, 6, 7
- [5] Sungmin Cha, YoungJoon Yoo, Taesup Moon, et al. Ssul: Semantic segmentation with unknown label for exemplar-based class-incremental learning. *NeurIPS*, 2021. 1, 3, 4, 6, 7, 8
- [6] Ali Cheraghian, Shafin Rahman, Sameera Ramasinghe, Pengfei Fang, Christian Simon, Lars Petersson, and Mehrtash Harandi. Synthesized feature based few-shot class-incremental learning on a mixture of subspaces. In *ICCV*, 2021. 3
- [7] Arthur Douillard, Yifu Chen, Arnaud Dapogny, and Matthieu Cord. Plop: Learning without forgetting for continual semantic segmentation. In *CVPR*, 2021. 1, 3
- [8] Nikita Dvornik, Julien Mairal, and Cordelia Schmid. Modeling visual context is key to augmenting object detection datasets. In *ECCV*, 2018. 3, 7
- [9] Debidatta Dwibedi, Ishan Misra, and Martial Hebert. Cut, paste and learn: Surprisingly easy synthesis for instance detection. In *ICCV*, 2017. 3
- [10] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 6
- [11] Dan Andrei Ganea, Bas Boom, and Ronald Poppe. Incremental few-shot instance segmentation. In *CVPR*, 2021. 3
- [12] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *CVPR*, 2021. 2, 3, 5
- [13] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 5, 6
- [14] Bharath Hariharan, Pablo Arbelaez, Lubomir Bourdev, Subhransu Maji, and Jitendra Malik. Semantic contours from inverse detectors. In *ICCV*, 2011. 6
- [15] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *CVPR*, 2020. 3
- [16] Zhizhong Li and Derek Hoiem. Learning without forgetting. *TPAMI*, 2017. 3
- [17] David Lopez-Paz and Marc’Aurelio Ranzato. Gradient episodic memory for continual learning. *NeurIPS*, 2017. 3
- [18] Arun Mallya, Dillon Davis, and Svetlana Lazebnik. Piggy-back: Adapting a single network to multiple tasks by learning to mask weights. In *ECCV*, 2018. 3
- [19] Michael McCloskey and Neal J Cohen. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*. Elsevier, 1989. 2
- [20] Umberto Michieli and Pietro Zanuttigh. Incremental learning techniques for semantic segmentation. In *ICCVW*, 2019. 3
- [21] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *ICCV*, 2019. 6
- [22] Sylvestre-Alvise Rebuffi, Alexander Kolesnikov, Georg Sperl, and Christoph H Lampert. icarl: Incremental classifier and representation learning. In *CVPR*, 2017. 3, 5
- [23] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *British Machine Vision Conference (BMVC)*, 2017. 3, 6
- [24] Xiaoyu Tao, Xiaopeng Hong, Xinyuan Chang, Songlin Dong, Xing Wei, and Yihong Gong. Few-shot class-incremental learning. In *CVPR*, 2020. 3
- [25] Zhuotao Tian, Xin Lai, Li Jiang, Michelle Shu, Hengshuang Zhao, and Jiaya Jia. Generalized few-shot semantic segmentation. *arXiv preprint arXiv:2010.05210*, 2020. 3, 6
- [26] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *TPAMI*, 2020. 3
- [27] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *ICCV*, 2019. 3, 4
- [28] Shipeng Yan, Jiale Zhou, Jiangwei Xie, Songyang Zhang, and Xuming He. An em framework for online incremental learning of semantic segmentation. In *Proceedings of the 29th ACM International Conference on Multimedia*, 2021. 1, 3
- [29] Jaehong Yoon, Eunho Yang, Jeongtae Lee, and Sung Ju Hwang. Lifelong learning with dynamically expandable networks. In *ICLR*, 2018. 3
- [30] Chang-Bin Zhang, Jia-Wen Xiao, Xialei Liu, Ying-Cong Chen, and Ming-Ming Cheng. Representation compensation networks for continual semantic segmentation. In *CVPR*, 2022. 1, 3
- [31] Xiaolin Zhang, Yunchao Wei, Yi Yang, and Thomas S Huang. Sg-one: Similarity guidance network for one-shot semantic segmentation. *IEEE transactions on cybernetics*, 2020. 3