

SAILS: Segment Anything with Incrementally Learned Semantics for Task-Invariant and Training-Free Continual Learning

Shishir Muralidhara¹ Didier Stricker^{1,2} René Schuster^{1,2}

¹German Research Center for Artificial Intelligence (DFKI)

²RPTU – University of Kaiserslautern-Landau, Kaiserslautern

firstname.lastname@dfki.de

Abstract—Continual learning remains constrained by the need for repeated retraining, high computational costs, and the persistent challenge of forgetting. These factors significantly limit the applicability of continual learning in real-world settings, as iterative model updates require significant computational resources and inherently exacerbate forgetting. We present SAILS – Segment Anything with Incrementally Learned Semantics, a training-free framework for Class-Incremental Semantic Segmentation (CISS) that sidesteps these challenges entirely. SAILS leverages foundational models to decouple CISS into two stages: Zero-shot region extraction using Segment Anything Model (SAM), followed by semantic association through prototypes in a fixed feature space. SAILS incorporates selective intra-class clustering, resulting in multiple prototypes per class to better model intra-class variability. Our results demonstrate that, despite requiring no incremental training, SAILS typically surpasses the performance of existing training-based approaches on standard CISS datasets, particularly in long and challenging task sequences where forgetting tends to be most severe. By avoiding parameter updates, SAILS completely eliminates forgetting and maintains consistent, task-invariant performance. Furthermore, SAILS exhibits positive backward transfer, where the introduction of new classes can enhance performance on previous classes.

I. INTRODUCTION

Continual learning is a paradigm where models continuously adapt to new data and tasks over time, unlike traditional isolated learning [1] that rely on fixed datasets and predefined objectives. However, this process introduces the challenge of catastrophic forgetting [2], where learning new tasks overwrites previously learned knowledge, degrading performance on previous tasks. Forgetting is an inherent consequence of the training process, as learning requires updating model parameters that adversely interferes with previously learned weights. The severity of forgetting largely depends on the task sequence, with longer sequences amplifying the cumulative effect of incremental updates and worsening forgetting over time. The stability-plasticity dilemma [3] also influences forgetting, highlighting the challenge of retaining prior knowledge (stability) while adapting to new information (plasticity). While mitigating forgetting remains the primary challenge in continual learning, efficiency is an equally crucial factor, particularly for real-world applications. Continual learning is inherently storage-efficient, as it works under the confines of restricted or no access to previous data. However, computational efficiency has only recently received attention, despite its

importance for practical deployment. Most CL methods assume access to extensive offline resources for continual updates, an unrealistic expectation in resource-constrained environments. Parameter-Efficient Continual Learning (PECL) addresses this by leveraging Parameter-Efficient Fine-Tuning (PEFT), updating only a small subset of parameters. This significantly reduces overhead, making continual learning more practical in resource-constrained environments. Despite its advantages, PECL still struggles with forgetting, as even minimal weight updates can interfere with previously learned tasks. Although more efficient than full fine-tuning, PECL methods still involve time-consuming training per task, limiting their use in real-world scenarios. Given these challenges, the desideratum of continual learning is to eliminate forgetting while maintaining resource-efficient adaptation. While preventing weight updates entirely is the only way to fully avoid forgetting, it limits the model’s capacity to learn new tasks. This trade-off presents a core challenge: *How can we build a continual learning system that is both resource-efficient and capable of learning new tasks without forgetting?*

In this work, we present Segment Anything with Incrementally Learned Semantics (SAILS), a novel training-free framework for Class-Incremental Semantic Segmentation (CISS). Our method leverages the powerful generalization capabilities of foundational models to segment objects and progressively learn new class semantics over time. We formulate CISS as a combination of two components: (1) Spatial segmentation and (2) Semantic association. First, we utilize the Segment Anything Model (SAM) [4] to extract object regions in a zero-shot, prompt-free manner. Then, we incrementally assign semantic meaning to these segments using class prototypes derived from a frozen network pretrained on large and diverse data. This enables continual adaptation without updating any model parameters. Our approach offers several advantages:

- **Training-Free Continual Learning:** Our framework enables continual learning without retraining, ensuring computational efficiency and near real-time adaptation.
- **Forgetting-Free Learning:** By avoiding model updates, SAILS sidesteps forgetting, maintaining stable performance across tasks without degradation over time.

- **Task-Invariant Results:** SAILS achieves consistent performance regardless of task sequence or length.
- **Positive Backward Transfer:** SAILS promotes positive backward transfer, improving performance on previous classes while adapting to new classes.

II. BACKGROUND AND RELATED WORKS

Incremental learning involves learning from a sequence of tasks over time and is commonly categorized into three settings [5]. In domain-incremental learning, the input distribution shifts while the number of classes remains fixed. In class-incremental learning, new classes from the same input distribution are introduced incrementally. Task-incremental learning can involve either domain or class shifts but assumes task IDs are available during inference. Recent works have extended these paradigms to more realistic scenarios: Evolving ontologies and semantic refinement [6], [7], incremental learning across modalities [8], [9], and simultaneous adaptation to domain and class shifts [10], [11]. In this work, we focus on class-incremental semantic segmentation, where disjoint subsets of classes are introduced sequentially.

A. Class-Incremental Semantic Segmentation

ILT [12] uses knowledge distillation across output and intermediate features with a frozen encoder. MiB [13] mitigates background shift through a modified distillation loss, to compare previous background predictions with current background and new class predictions. PLOP [14] uses pseudo-labels for background pixels representing previously seen classes. MBS [15] highlights the misclassification of undetected old class pixels as background during pseudo-labeling, and proposes selective pseudo-labeling and adaptive distillation. SSUL-M [16] introduces an unknown class to separate future classes from the background, in addition to pseudo-labeling of previous classes and rehearsal with exemplar memory. ALIFE [17] replays stored feature representations instead of explicitly storing data from previous tasks. RECALL [18] retrieves or generates old class images and pseudo-labels them using previous models. SATS [19] leverages transformer self-attention maps and class-specific region pooling for both inter-class and intra-class knowledge distillation. CIT [20] introduces a class independent transformation to convert outputs into class-independent forms, enabling accumulative distillation. NeST [21] proposed a new classifier pre-tuning method that learns a transformation from old classifiers to initialize new ones.

B. Parameter-Efficient Continual Learning

Parameter-Efficient Continual Learning (PECL) enables continual learning in resource-constrained environments by leveraging Parameter-Efficient Fine-Tuning (PEFT) methods to reduce trainable parameters and computational overhead. PEFT methods [22] include additive tuning, which introduces new parameters through adapters [23], [24], or prompts [25], [26]); Partial tuning of selected weights [27], [28], [29], and reparameterization with low-rank updates [30], [31]. Building on these methods, recent

PECL techniques include: Training of task-specific LoRA models [32] and merging of them using task arithmetic [33]. CoLoR [34] uses LoRA expert models and k-means clustering to infer task-ID during inference. LAE [35] introduces a three-stage framework: online PEFT for task learning, offline PEFT for knowledge accumulation, and ensembling for inference. O-LoRA [36] incrementally adds LoRA modules while enforcing orthogonality across tasks to reduce interference. InfLoRA [37] proposes interference-free low-rank adaptation by constructing subspaces to prevent interference between tasks. CLoRA [38] uses a single LoRA module with knowledge distillation to transfer knowledge between tasks in class-incremental semantic segmentation.

C. Continual Learning with Pre-trained Networks

Recent advancements in foundation models and large-scale pretraining have significantly influenced continual learning. Pretrained models provide strong generalization capabilities, reducing the need for extensive task-specific training. However, adapting these models for continual learning introduces unique challenges, including catastrophic forgetting of pretrained knowledge, distribution shifts between pretraining and downstream tasks, and limited plasticity in frozen or partially frozen networks. Pelosin [39] presents a training-free approach that leverages pretrained models to compute class prototypes and populate a memory bank. Similarly, [40] demonstrate that a frozen pretrained feature extractor with a Nearest Mean Classifier (NMC) can outperform more complex methods, underscoring the strength of pretrained representations. Extending this idea, [41] proposes Adapt and Merge (APER), which first adapts the pretrained model on the initial training set using PEFT to bridge the domain gap. APER concatenates frozen embeddings from both models to preserve generalizability and adaptivity. Supervised Contrastive Replay [42] improves nearest-prototype classification by encouraging intra-class compactness and inter-class separation in the embedding space. RanPAC [43] inserts a frozen random projection layer between pretrained model features and the output head to improve separability between classes, ensuring better class prototypes

III. SEGMENT ANYTHING WITH INCREMENTALLY LEARNED SEMANTICS

While continual learning with pretrained models has made significant strides, the majority of these advancements have focused primarily on image classification tasks. Class-incremental learning involves incrementally learning from a sequence of tasks $T = \{0, 1, \dots, n\}$, where each task t introduces non-overlapping subsets of classes such that $C_t \subset C$ and $C_i \cap C_j = \emptyset \forall i \neq j$. In classification, each image I corresponds to a single label $c \in C_t$, and the subset of classes $C_t \subset C$ encountered across tasks T is disjoint. In contrast, semantic segmentation involves dense, pixel-wise predictions where multiple classes may exist within a single image. At each time step t , the model has access to annotations only belonging to a subset of classes C_t . Pixels belonging to classes outside C_t , which include classes

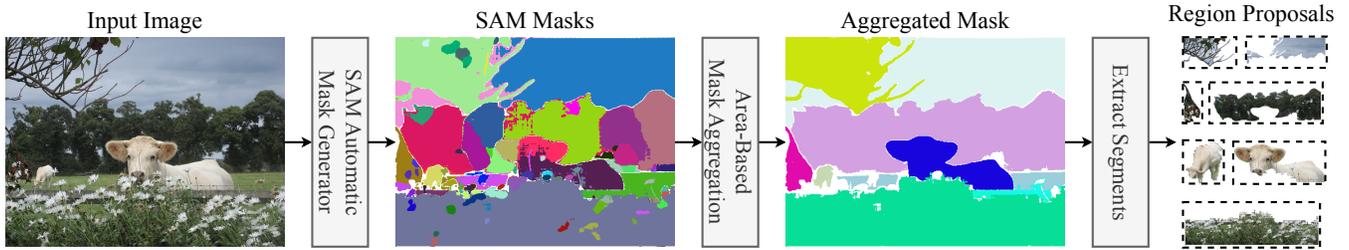


Fig. 1: The input image is first segmented using SAM, and the resulting masks are iteratively refined to produce an aggregated mask with distinct, non-overlapping regions, which are then used for extraction of region proposals from the image.

seen so far $C_{0:t-1}$ and potential future classes $C_{t+1:n}$ are labeled as background. When distilling knowledge between tasks, this results in a mismatch between the predictions of the previous and current task model, resulting in background shift and exacerbating the problem of catastrophic forgetting. Background shift necessitates relearning of the current classes, which may have been previously learned as the background class, and would require retraining. In a training-free setup, where updating previous knowledge is not possible, we must find a way to handle classes outside the current subset C_t without retraining. To address this challenge, we propose a modular decomposition of class-incremental semantic segmentation into two subproblems:

- **Spatial Segmentation** Using SAM, we generate class-agnostic region proposals $\mathcal{R} = \{r_1, r_2, \dots, r_K\}$. These proposals represent potential objects and provide a spatial prior for downstream semantic reasoning.
- **Semantic Association** As new classes C_t are introduced, we update the semantic understanding of the regions using class prototypes, ensuring a consistent and continually refined mapping from regions to classes.

A. Spatial Segmentation

We first segment the input images using the automatic mask generator from Segment Anything Model (SAM) [4]. SAM provides the flexibility to generate segmentation masks without any task-specific training, making it suitable for continual learning. However, integrating SAM into class-incremental semantic segmentation poses two key challenges: mask granularity and missing semantic context.

1) *Mask Granularity.*: The output from SAM typically consists of a large number of masks, many of which are overlapping and redundant. These masks tend to be fine-grained, often representing sub-parts or fragments of objects. While this granularity can be valuable for certain applications, it poses a challenge for semantic segmentation. These fragmented regions do not adequately capture the holistic characteristics of objects, which is crucial for subsequent semantic learning. To address these challenges, we implement a refinement step that consolidates the initial SAM output into a single mask per image. By prioritizing large, non-overlapping masks and discarding redundant and partial masks, our approach produces object-level segments.

Let I be the input image, and $\mathcal{M} = \{M_1, M_2, \dots, M_N\}$

denote the set of binary masks generated by SAM, where each mask $M_i \in \{0, 1\}^{H \times W}$ represents a region and is accompanied by metadata, including its area. We first remove masks that trivially cover most of the image, as they are unlikely to represent meaningful object instances. We retain masks whose area is less than a threshold fraction $\tau_{area} = 0.9$ of the total image size:

$$\mathcal{M}' = \{M_i \in \mathcal{M} \mid \text{area}_i < \tau_{area} \cdot (H \cdot W)\} \quad (1)$$

Next, we sort the filtered masks in descending order by area. Larger masks are more likely to represent entire objects, whereas smaller masks often correspond to object parts. Sorting ensures that larger, more comprehensive regions are prioritized during mask consolidation:

$$\mathcal{M}' = \{M_1, M_2, \dots, M_K\} \text{ with } A_1 \geq \dots \geq A_K \quad (2)$$

where A_i denotes the area of mask M_i .

We initialize an empty composite mask $M_{\text{agg}} \in \{1, 2, \dots\}^{H \times W}$, where each pixel holds an integer label corresponding to a region r_k . Initially, all values in M_{agg} are zero, indicating that no region has been assigned. We then iterate over each sorted mask $M_i \in \mathcal{M}'$. For each mask, we check if it contains foreground pixels not yet labeled in M_{agg} . If such pixels exist, we assign a new label c to these pixels, incrementing c sequentially for each new region. The update rule for each pixel (x, y) is:

$$M_{\text{agg}}(x, y) = c \text{ if } M_i(x, y) = 1 \text{ and } M_{\text{agg}}(x, y) = 0 \quad (3)$$

We repeat this process for all masks in \mathcal{M}' , resulting in the composite mask M_{agg} that contains unique and non-overlapping region proposals $\mathcal{R} = \{r_1, r_2, \dots, r_K\}$. These regions form the basis for the subsequent incremental semantic learning. Fig. 1 illustrates this aggregation process, where the output from SAM is consolidated into M_{agg} .

B. Semantic Association

SAM produces class-agnostic masks by segmenting image regions purely based on visual and spatial characteristics, without any semantic information. Following the initial segmentation step using SAM, the resulting mask segments or region proposals are used for extracting regions of interest (RoIs) from the image. The key challenge that follows is to incrementally assign semantic labels to these segments

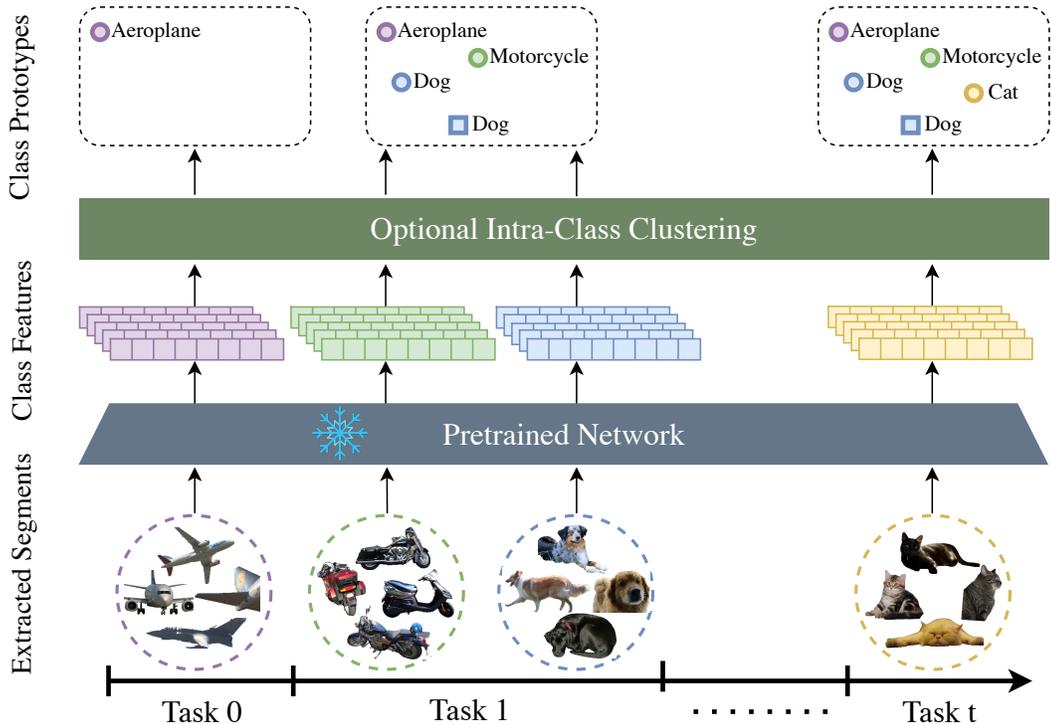


Fig. 2: Overview of the incremental semantic learning process. Regions of interest (ROIs) corresponding to classes introduced in the current task are embedded using a frozen pretrained backbone. For each class, either a single prototype or multiple sub-prototypes are computed based on the intra-class variability.

as new classes are introduced over time. This requires a continual learning method that sequentially integrates new classes with unlabeled regions while preserving previously learned knowledge and preventing forgetting. This reframes continual semantic segmentation as a continual classification problem of the segmented regions. Although existing continual learning methods for image classification could be adapted for this purpose, they rely on training and updating the model, which brings us back to the core motivation of this paper: Training is computationally expensive, forgetting becomes inevitable, and severe performance degradation often occurs over long task sequences.

To overcome these challenges, we propose a prototype-based classifier built upon frozen, pretrained vision backbones that requires no additional training or parameter updates. Our approach consists of the following steps:

- **Feature Extraction:** Given an input image, we extract Regions of Interest (ROIs) using the aggregated SAM mask. Each ROI r is a rectangular crop around the object, with the background masked and is passed through a frozen network $f(\cdot)$ to obtain a feature embedding:

$$\mathbf{z}_r = \text{norm}(f(r)) \in \mathbb{R}^d \quad (4)$$

where d is the feature dimension.

- **Prototype Computation:** For each new class $c \in C_t$, we compute a class prototype \mathbf{p}_c by averaging the

embeddings of all segments belonging to class c :

$$\mathbf{p}_c = \text{norm}\left(\frac{1}{N_c} \sum_{i=1}^{N_c} \mathbf{z}_{r_i}\right), \quad (5)$$

where N_c is the number of segments labeled as class c .

- **Classification by Prototype Similarity:** To classify a new segment r , we compute the cosine similarity between its embedding \mathbf{z}_r and each stored prototype for all classes $C_{0:t}$ encountered up to task t . The segment is assigned the class with the highest similarity score exceeding a predefined threshold τ_{sim} :

$$y_r = \begin{cases} \arg \max_{c \in C_{0:t}} \frac{\mathbf{z}_r \cdot \mathbf{p}_c}{\|\mathbf{z}_r\| \|\mathbf{p}_c\|}, & \text{if } sim \geq \tau_{sim} \\ \text{BG}, & \text{otherwise} \end{cases} \quad (6)$$

To address the background shift problem in class-incremental semantic segmentation, we incorporate a similarity threshold $\tau_{sim} = 0.5$ during classification. Some ROIs may correspond to future classes or background regions that are not semantically relevant. As both cases lack valid labels in the current task, we assign such ROIs to the background (BG) class if their similarity to all known class prototypes falls below the threshold. This prevents misclassification and avoids assigning incorrect semantics to ambiguous or currently irrelevant regions.

1) *Intra-Class Variance:* A key limitation of relying on a single prototype per class is its inability to capture high intra-class variability effectively. Classes like dogs with

varied breeds or vehicles of different types exhibit significant visual diversity that a single prototype cannot fully represent. Averaging widely dispersed embeddings can cause the class prototype to fall into semantically ambiguous regions of the feature space, failing to align well with any specific subgroup. Addressing this limitation, we introduce an intra-class clustering mechanism for classes with high variance. To identify classes with high variance, we compute a variance score based on cosine distance. For each class, we calculate the mean feature vector μ_c and measure the average cosine distance to all feature vectors in the class. This score reflects how tightly the features are grouped, with higher values indicating more variance. Classes with variance scores above 0.4 are chosen for sub-clustering.

$$\text{Variance}_c = \frac{1}{N_c} \sum_{i=1}^{N_c} \text{cosine_distance}(\mathbf{x}_i, \mu_c) \quad (7)$$

By freezing the backbone and only incrementally adding new class prototypes, our approach remains training-free, efficient, and inherently resistant to forgetting as it avoids interference with previous class prototypes. Fig. 2 illustrates this process: RoIs corresponding to classes introduced in the current task are embedded using a frozen pretrained backbone, and prototypes are computed for each class. During inference, the RoI embedding is matched against all prototypes and sub-prototypes across classes and assigned the class with the highest similarity.

IV. RESULTS

In this section, we present experiments and results that demonstrate the effectiveness of our approach. We begin by describing the datasets used, followed by the baseline methods and implementation details. We then provide a detailed analysis of our approach’s performance across incremental tasks, highlighting its advantages over existing methods.

A. Datasets

We evaluate our approach on two standard benchmarks for class-incremental semantic segmentation (CISS):

- **PASCAL VOC 2012** [44] consists of 21 semantic classes with 1,464 images in the training set. Most CISS methods use the augmented set with over 10k images, while our training-free approach achieves similar results using only the original training split, eliminating the need for large-scale annotated data.
- **Cityscapes** [45] is a challenging urban scene understanding dataset with 19 semantic classes and 2,975 training images. Our results on Cityscapes further highlight the effectiveness and suitability of our approach for real-world applications like autonomous driving.

B. Baselines and Implementation

For all comparisons, we adopt the SegFormer-B2 architecture from the SATS framework [19], while for SSUL, we use its corresponding framework. We perform a comprehensive comparison against joint training (JT) and several state-of-the-art CISS methods, including ILT [12], MiB [13],

TABLE I: Results (in mIoU) on PASCAL VOC [44] dataset after learning all steps. SAILS* denotes intra-class clustering.

Method	1-1	2-1	5-1	2-2
	(20 Steps)	(19 Steps)	(16 Steps)	(10 Steps)
JT	80.12	80.12	80.12	80.12
ILT [12]	05.55	05.95	06.25	10.72
MiB [13]	17.03	25.74	45.63	51.83
PLOP [14]	07.55	08.14	06.94	11.32
SATS [19]	17.97	05.89	37.86	39.74
SSUL [16]	29.15	38.32	48.65	45.31
CLIP	35.97	35.97	35.97	35.97
SAILS	54.43	54.43	54.43	54.43
SAILS*	57.96	57.96	57.96	57.96

PLOP [14], SSUL [16], and SATS [19]. The joint training baseline trains the model on all classes simultaneously in a single task. As a result, it does not experience catastrophic forgetting and serves as the upper-bound for comparison. We also compare against SAM+CLIP, which combines SAM and the default CLIP [46] ViT-B model with the officially published code for zero-shot prediction. While it shares the training-free benefits, we highlight its challenges and limitations. For feature extraction in our approach, we use a Swin-B [47] network pretrained on ImageNet. To handle classes with high variance, we apply k-means clustering with $k = 5$ within each class to obtain multiple fine-grained sub-prototypes that better capture semantic diversity.

C. PASCAL VOC

We evaluate robustness under long task sequences where conventional methods typically struggle due to the cumulative effect of forgetting. CIL settings use *init-inc* notation, where *init* is the initial number of classes and *inc* the classes added incrementally, both of which influence the severity of forgetting. We conduct experiments under the *1-1*, *2-1*, *5-1* and *2-2* task splits, which lead to a large number of incremental steps. The performance of our approach is compared against state-of-the-art (SotA) baselines, as summarized in Tab. I. Performance of existing approaches degrades significantly as the number of steps increases, highlighting their limitations in handling long task sequences. In contrast, our training-free framework, SAILS, maintains consistent performance across all task configurations, showcasing its robustness and task-invariant results regardless of sequence length. Furthermore, with intra-class clustering, we observe improved performance across tasks.

D. Cityscapes

To further validate the generalizability of our approach, we evaluate SAILS on the Cityscapes [45] dataset under similar challenging task splits: *1-1*, *2-1*, *3-2*, and *5-1*. We compare SAILS against SotA baselines under these extreme task splits and the results are presented in Tab. II. Forgetting is amplified by the underrepresented classes such as *motorcycle* and *bicycle*, which are encountered toward

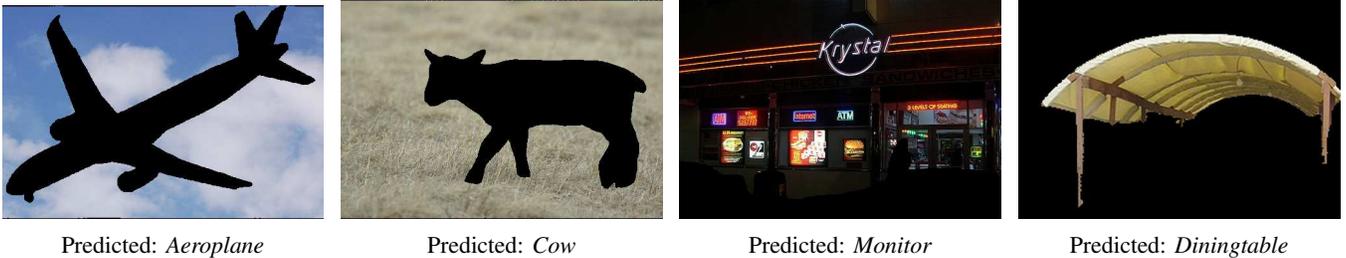


Fig. 3: CLIP misclassifications of region segments due to shortcut learning and ambiguous regions.

TABLE II: Results (in mIoU) on Cityscapes [45] dataset after learning all tasks. SAILS* denotes intra-class clustering.

Method	1-1 (19 Steps)	2-1 (18 Steps)	5-1 (15 Steps)	3-2 (9 Steps)
JT	59.06	59.06	59.06	59.06
ILT [12]	00.12	00.22	01.52	00.81
MiB [13]	29.71	31.03	35.35	35.06
PLOP [14]	11.48	11.78	16.79	22.10
SATS [19]	15.16	21.26	24.12	33.63
CLIP	15.48	15.48	15.48	15.48
SAILS	36.51	36.51	36.51	36.51
SAILS*	45.06	45.06	45.06	45.06

TABLE III: Comparison of segment quality across methods.

Segments	mIoU	Precision	Recall
Ground Truth	56.18	80.35	71.20
Ground Truth (Non-BG)	70.87	98.49	71.64
SAM [4]	27.52	47.67	37.67
Mask2Former [48]	40.51	64.71	48.14
SAM-Agg (Ours)	54.43	79.25	63.67

the end of the task sequence in Cityscapes, and hinders effective learning in later stages. In contrast, SAILS exhibits stable performance across all tasks, mirroring the robustness observed on PASCAL VOC. Our training-free approach circumvents the cumulative performance degradation observed in methods that rely on model updates.

E. Intra-Class Clustering

While our training-free approach ensures efficiency and avoids forgetting, it relies on frozen features, limiting adaptability to intra-class variability. To address this, we introduce *intra-class clustering*, allowing multiple prototypes per class to better capture feature diversity. As shown in Tab. IV, intra-class clustering leads to improved segmentation performance for most of these classes. However, for visually similar classes like *chair* and *sofa*, performance slightly drops. This is likely due to increased confusion between these classes, as the additional prototypes can represent similar visual features. For *potted plant*, clustering improves recall but reduces precision. *Background* segments like trees and vegetation are misclassified as potted plants due to their visual similarity.

F. Positive Backward Transfer

While most continual learning methods have primarily focused on mitigating forgetting, SAILS goes further by exhibiting positive backward transfer, improving performance on previous tasks as new classes are added. This effect is more pronounced when the newly added classes are visually similar to the previously learned classes. The new semantic context helps disambiguate earlier misclassifications between visually similar classes. We illustrate this phenomenon in Tab. V using the 15-5 task setting on PASCAL VOC. After learning the initial 15 classes in step 0, we observe that their performance improves following the introduction of the final 5 classes in step 1. For instance, introducing *train* enhances the previously learned *bus* by offering finer semantic contrast. Similarly, the addition of *sofa* improves the visually similar *chair* class.

G. Zero-Shot Segmentation using SAM and CLIP

While CLIP [46] is effective for zero-shot classification, applying it to region-level segmentation reveals several limitations. It exhibits shortcut learning, where regions are misclassified based on scene context rather than object-specific features. In Fig. 3, the sky surrounding the plane is misclassified as *aeroplane*, likely due to the plane’s silhouette. Similarly, the grass around the *sheep* is labeled as *cow*, likely due to the contextual similarity between the two scenes. CLIP also struggles with partial or ambiguous region segments such as windows and signboards being classified as *monitor* based on shape similarity. More critically, CLIP struggles in continual learning settings that require recognizing novel, fine-grained classes not encountered during its pretraining. Without task-specific fine-tuning, its utility for incremental learning is limited.

V. DISCUSSION

In this section, we discuss key design choices and limitations of our framework, and outline promising directions for future work. Our goal is to highlight the trade-offs underlying our approach and identify areas where further research could extend its applicability.

A. Leveraging SAM for Segmentation

A key design choice in our framework is the use of SAM [4] for universal region segmentation. While computationally heavier, SAM provides dense and high-quality masks and crucially, generalizes across domains without retraining.

TABLE IV: Task-invariant class-wise IoU on PASCAL VOC [44]. SAILS* denotes our approach with intra-class clustering. Classes represented with intra-class clustering are highlighted in **bold**.

Method	background	airplane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	diningtable	dog	horse	motorbike	person	pottedplant	sheep	sofa	train	tvmonitor
SAILS	83.13	32.13	42.63	34.53	52.88	52.40	88.16	60.24	83.59	34.65	77.49	01.86	17.48	87.28	78.60	32.76	34.64	70.81	25.94	81.68	70.17
SAILS*	83.16	32.13	42.63	41.20	64.84	58.37	88.16	73.09	83.59	24.12	77.46	10.51	42.63	87.28	78.82	55.12	25.78	70.81	25.65	81.68	70.17

TABLE V: Class-wise IoU without intra-class clustering on PASCAL VOC [44] (15-5 setting), highlighting positive backward transfer on the initial 15 classes after learning the final 5 classes in step 1.

Step	background	airplane	bicycle	bird	boat	bottle	bus	car	cat	chair	cow	diningtable	dog	horse	motorbike	person	pottedplant	sheep	sofa	train	tvmonitor
0	86.16	32.13	42.63	34.53	52.88	52.40	86.47	60.24	83.59	33.06	77.11	01.89	17.47	87.28	78.60	32.75					
1	83.13	32.13	42.63	34.53	52.88	52.40	88.16	60.24	83.59	34.65	77.49	01.86	17.48	87.28	78.60	32.76	34.64	70.81	25.94	81.68	70.17

In contrast, salient object detection covers only dominant foreground objects, and panoptic segmentation models are dataset and domain specific, leading to degraded performance when applied to new domains. To quantify segmentation quality, we compare SAM segments, our aggregated object-level segments (SAM-Agg), and ground truth (GT). We include proposals from [49] generated using Mask2Former [48], a panoptic segmentation model trained on MS-COCO [50]. We assign classes using the same training-free, prototype-based classification, and the results are presented in Tab. III.

Ground-truth segments represent the upper bound of segmentation performance, providing precise and semantically consistent masks, serving both as a reference for ideal segmentation and evaluating our prototype-based classification. Excluding background segments further improves accuracy and precision while maintaining high recall. SAM’s fine-grained masks capture detailed part-level structures but often over-segment objects into fragmented components such as parts or limbs. Our refinement (SAM-Agg) consolidates these into coherent object-level masks, leading to an improvement in precision and recall while also reducing the number of segments and thus computational cost. Mask2Former produces stronger segments than SAM, but this advantage is largely due to its training on COCO [50], which closely resembles PASCAL VOC [44] highlighting SAM’s strength as a training-free, domain-agnostic alternative. More computationally efficient alternatives to SAM, can help alleviate the computational demands and lower resource requirements.

B. Frozen Pretrained Representations

Since our framework uses a frozen feature extractor, its adaptability is ultimately constrained by the discriminative capacity of the pretrained backbone. To mitigate this limitation, we employ selective intra-class clustering to identify subgroups within each class. However, distinguishing between visually or semantically similar classes remains challenging. Classes such as *sofa* and *chair* share many visual characteristics, leading to frequent misclassifications.

Using a frozen backbone without task-specific retraining ensures task-invariance and prevents catastrophic forgetting, but it also constrains adaptability. Future work could explore methods to improve downstream adaptation while retaining efficiency through PECL. Additionally, methods built on top of the frozen features, such as contrastive learning, can enhance separability between classes while maintaining the benefits of a frozen, stable backbone.

VI. CONCLUSION

Continual learning is often hindered in real-world settings due to the need for repeated retraining, high computational costs, and catastrophic forgetting. Addressing these challenges, we present SAILS, a training-free framework for class-incremental semantic segmentation that avoids retraining and forgetting by leveraging SAM for zero-shot region extraction and prototype-based semantic association. With selective intra-class clustering, SAILS effectively handles class variability and consistently outperforms training-based methods, especially in long task sequences. The training-free approach ensures task-invariant performance and even positive backward transfer, making it a robust and efficient solution for continual semantic segmentation.

ACKNOWLEDGMENTS

This work was partially funded by the German Federal Ministry of Research, Technology, and Space under the project COPPER (16IW24009).

REFERENCES

- [1] B. Liu, “Lifelong machine learning: a paradigm for continuous learning,” *Frontiers of Computer Science*, 2017.
- [2] R. M. French, “Catastrophic forgetting in connectionist networks,” *Trends in cognitive sciences*, 1999.
- [3] M. Mermillod, A. Bugaiska, and P. Bonin, “The stability-plasticity dilemma: Investigating the continuum from catastrophic forgetting to age-limited learning effects,” 2013.
- [4] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, “Segment anything,” in *International Conference on Computer Vision (ICCV)*, 2023.
- [5] G. M. Van de Ven and A. S. Tolias, “Three scenarios for continual learning,” *arXiv preprint arXiv:1904.07734*, 2019.
- [6] Z. Lin, D. Pathak, Y.-X. Wang, D. Ramanan, and S. Kong, “Continual learning with evolving class ontologies,” *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

- [7] S. Muralidhara, S. Bukhari, G. Schneider, D. Stricker, and R. Schuster, "Cleo: Continual learning of evolving ontologies," in *European Conference on Computer Vision (ECCV)*, 2024.
- [8] N. Hegde, S. Muralidhara, R. Schuster, and D. Stricker, "Modality-incremental learning with disjoint relevance mapping networks for image-based semantic segmentation," in *Winter Conference on Applications of Computer Vision (WACV)*, 2025.
- [9] Y. Song, X. Yang, D. Jiang, Y. Wang, and C. Xu, "Harmony: A unified framework for modality incremental learning," *arXiv:2504.13218*, 2025.
- [10] M. Toldo, U. Michieli, and P. Zanuttigh, "Learning with style: Continual semantic segmentation across tasks and domains," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- [11] M.-Y. Park, J.-H. Lee, and G.-M. Park, "Versatile incremental learning: Towards class and domain-agnostic incremental learning," in *European Conference on Computer Vision*, pp. 271–288, Springer, 2024.
- [12] U. Michieli and P. Zanuttigh, "Incremental learning techniques for semantic segmentation," in *International Conference on Computer Vision Workshops (ICCVW)*, 2019.
- [13] F. Cermelli, M. Mancini, S. R. Bulò, E. Ricci, and B. Caputo, "Modeling the background for incremental learning in semantic segmentation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [14] A. Douillard, Y. Chen, A. Dapogny, and M. Cord, "Plop: Learning without forgetting for continual semantic segmentation," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [15] G. Park, W. Moon, S. Lee, T.-Y. Kim, and J.-P. Heo, "Mitigating background shift in class-incremental semantic segmentation," in *European Conference on Computer Vision (ECCV)*, 2024.
- [16] S. Cha, Y. Yoo, T. Moon, *et al.*, "Ssul: Semantic segmentation with unknown label for exemplar-based class-incremental learning," *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.
- [17] Y. Oh, D. Baek, and B. Ham, "Alife: Adaptive logit regularizer and feature replay for incremental semantic segmentation," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [18] A. Maracani, U. Michieli, M. Toldo, and P. Zanuttigh, "Recall: Replay-based continual learning in semantic segmentation," in *International Conference on Computer Vision (ICCV)*, 2021.
- [19] Y. Qiu, Y. Shen, Z. Sun, Y. Zheng, X. Chang, W. Zheng, and R. Wang, "Sats: Self-attention transfer for continual semantic segmentation," *Pattern Recognition*, 2023.
- [20] J. Ge, B. Zhang, A. Liu, V. M. H. Phan, Q. Chen, Y. Shu, and Y. Zhao, "Cit: Rethinking class-incremental semantic segmentation with a class independent transformation," *Pattern Recognition*, 2025.
- [21] Z. Xie, H. Lu, J.-w. Xiao, E. Wang, L. Zhang, and X. Liu, "Early preparation pays off: New classifier pre-tuning for class incremental semantic segmentation," in *European Conference on Computer Vision (ECCV)*, 2024.
- [22] L. Xu, H. Xie, S.-Z. J. Qin, X. Tao, and F. L. Wang, "Parameter-efficient fine-tuning methods for pretrained language models: A critical review and assessment," *arXiv:2312.12148*, 2023.
- [23] A. Rücklé, G. Geigle, M. Glockner, T. Beck, J. Pfeiffer, N. Reimers, and I. Gurevych, "Adapterdrop: On the efficiency of adapters in transformers," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [24] H. Liu, D. Tam, M. Muqeeth, J. Mohta, T. Huang, M. Bansal, and C. A. Raffel, "Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning," *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [25] B. Lester, R. Al-Rfou, and N. Constant, "The power of scale for parameter-efficient prompt tuning," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2021.
- [26] X. L. Li and P. Liang, "Prefix-tuning: Optimizing continuous prompts for generation," *arXiv:2101.00190*, 2021.
- [27] E. B. Zaken, S. Ravfogel, and Y. Goldberg, "Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models," *arXiv:2106.10199*, 2021.
- [28] M. Zhao, T. Lin, F. Mi, M. Jaggi, and H. Schütze, "Masking as an efficient alternative to finetuning for pretrained language models," in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [29] D. Guo, A. M. Rush, and Y. Kim, "Parameter-efficient transfer learning with diff pruning," *arXiv:2012.0746*, 2020.
- [30] E. J. Hu, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, W. Chen, *et al.*, "Lora: Low-rank adaptation of large language models," in *International Conference on Learning Representations (ICLR)*, 2021.
- [31] M. Valipour, M. Rezagholizadeh, I. Kobyzev, and A. Ghodsi, "Dylora: Parameter efficient tuning of pre-trained models using dynamic search-free low-rank adaptation," *arXiv:2210.07558*, 2022.
- [32] R. Chitale, A. Vaitya, A. Kane, and A. Ghotkar, "Task arithmetic with lora for continual learning," *arXiv:2311.02428*, 2023.
- [33] G. Ilharco, M. T. Ribeiro, M. Wortsman, S. Gururangan, L. Schmidt, H. Hajishirzi, and A. Farhadi, "Editing models with task arithmetic," in *International Conference on Learning Representations (ICLR)*, 2023.
- [34] M. Wistuba, P. T. Sivaprasad, L. Balles, and G. Zappella, "Continual learning with low rank adaptation," in *Workshops of Advances in Neural Information Processing (NeurIPS)*, 2023.
- [35] Q. Gao, C. Zhao, Y. Sun, T. Xi, G. Zhang, B. Ghanem, and J. Zhang, "A unified continual learning framework with general parameter-efficient tuning," in *International Conference on Computer Vision (ICCV)*, 2023.
- [36] X. Wang, T. Chen, Q. Ge, H. Xia, R. Bao, R. Zheng, Q. Zhang, T. Gui, and X. Huang, "Orthogonal subspace learning for language model continual learning," *arXiv:2310.14152*, 2023.
- [37] Y.-S. Liang and W.-J. Li, "Inflora: Interference-free low-rank adaptation for continual learning," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [38] S. Muralidhara, D. Stricker, and R. Schuster, "Clora: Parameter-efficient continual learning with low-rank adaptation," *arXiv:2507.19887*, 2025.
- [39] F. Pelosin, "Simpler is better: off-the-shelf continual learning through pretrained backbones," *arXiv:2205.01586*, 2022.
- [40] P. Janson, W. Zhang, R. Aljundi, and M. Elhoseiny, "A simple baseline that questions the use of pretrained-models in continual learning," *arXiv:2210.04428*, 2022.
- [41] D.-W. Zhou, Z.-W. Cai, H.-J. Ye, D.-C. Zhan, and Z. Liu, "Revisiting class-incremental learning with pre-trained models: Generalizability and adaptivity are all you need," *International Journal of Computer Vision (IJCV)*, 2025.
- [42] Z. Mai, R. Li, H. Kim, and S. Sanner, "Supervised contrastive replay: Revisiting the nearest class mean classifier in online class-incremental continual learning," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [43] M. D. McDonnell, D. Gong, A. Parvaneh, E. Abbasnejad, and A. Van den Hengel, "Ranpac: Random projections and pre-trained models for continual learning," *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [44] M. Everingham, L. van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International Journal of Computer Vision (IJCV)*, 2010.
- [45] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [46] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, *et al.*, "Learning transferable visual models from natural language supervision," in *International Conference on Machine Learning (ICML)*, 2021.
- [47] Z. Liu, Y. Lin, Y. Cao, H. Hu, Y. Wei, Z. Zhang, S. Lin, and B. Guo, "Swin transformer: Hierarchical vision transformer using shifted windows," in *International Conference on Computer Vision (ICCV)*, 2021.
- [48] B. Cheng, I. Misra, A. G. Schwing, A. Kirillov, and R. Girdhar, "Masked-attention mask transformer for universal image segmentation," in *Conference on Computer Vision and Pattern Recognition*, 2022.
- [49] Z. Zhang, G. Gao, Z. Fang, J. Jiao, and Y. Wei, "Mining unseen classes via regional objectness: A simple baseline for incremental segmentation," *Advances in neural information processing systems*, 2022.
- [50] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick, "Microsoft coco: Common objects in context," in *European Conference on Computer Vision (ECCV)*, 2014.