



SegViT v2: Exploring Efficient and Continual Semantic Segmentation with Plain Vision Transformers

Bowen Zhang¹ · Liyang Liu¹ · Minh Hieu Phan¹ · Zhi Tian³ · Chunhua Shen² · Yifan Liu¹

Received: 3 March 2023 / Accepted: 10 August 2023 / Published online: 27 October 2023
© The Author(s) 2023

Abstract

This paper investigates the capability of plain Vision Transformers (ViTs) for semantic segmentation using the encoder–decoder framework and introduce *SegViTv2*. In this study, we introduce a novel Attention-to-Mask (ATM) module to design a lightweight decoder effective for plain ViT. The proposed ATM converts the global attention map into semantic masks for high-quality segmentation results. Our decoder outperforms popular decoder UPerNet using various ViT backbones while consuming only about 5% of the computational cost. For the encoder, we address the concern of the relatively high computational cost in the ViT-based encoders and propose a *Shrunk++* structure that incorporates edge-aware query-based down-sampling (EQD) and query-based up-sampling (QU) modules. The *Shrunk++* structure reduces the computational cost of the encoder by up to 50% while maintaining competitive performance. Furthermore, we propose to adapt SegViT for continual semantic segmentation, demonstrating nearly zero forgetting of previously learned knowledge. Experiments show that our proposed SegViTv2 surpasses recent segmentation methods on three popular benchmarks including ADE20k, COCO-Stuff-10k and PASCAL-Context datasets. The code is available through the following link: <https://github.com/zbwpx/SegVit>.

Keywords Vision transformer · Incremental learning · Semantic segmentation · Continual learning

Communicated by Kaiyang Zhou.

Bowen Zhang, Liyang Liu and Minh Hieu Phan have contributed equally to this work.

✉ Yifan Liu
yifan.liu04@adelaide.edu.au

Bowen Zhang
b.zhang@adelaide.edu.au

Liyang Liu
akide.liu@adelaide.edu.au

Minh Hieu Phan
vuminhhieu.phan@adelaide.edu.au

Zhi Tian
zhi.tian@outlook.com

Chunhua Shen
chunhuashen@zju.edu.cn

¹ The University of Adelaide, Adelaide, Australia

² Zhejiang University, Hangzhou, China

³ Meituan Inc., Beijing, China

1 Introduction

Semantic segmentation is a pivotal computer vision task that aims to assign labels to every pixel on the image. Widely adopted state-of-the-art methods like Fully Convolutional Networks (FCN) (Long et al., 2015) utilize deep convolutional neural networks (ConvNet) as encoders and incorporate segmentation decoders for dense predictions. Prior works (Wang et al., 2020; Yuan et al., 2020; Chen et al., 2018) have aimed to enhance performance by augmenting contextual information or incorporating multi-scale information, leveraging the inherent multi-scale and *hierarchical* attributes of the ConvNet architectures.

The advent of the Vision Transformer (ViT) (Dosovitskiy et al., 2021) has offered a paradigm shift, serving as a robust backbone for numerous computer vision tasks. ViT, distinct from ConvNet base models, retains a plain and *non-hierarchical* architecture while preserving the resolution of the feature maps. To conveniently leverage existing segmentation decoders for dense prediction, such as U-Net (Ronneberger et al., 2015) or DeepLab (Chen et al., 2018), recent Transformer-based approaches, including Swin Trans-

former (Liu et al., 2021) and PVT (Wang et al., 2021), have developed a *hierarchical* ViT to extract hierarchical feature representations.

However, modifying the original ViT structures requires training the networks from scratch rather than using off-the-shelf plain ViT checkpoints due to the discrepancy between the hierarchical and plain architectures, such as spatial down-sampling (Xu et al., 2022). Altering the plain ViT architecture compromises the use of rich representations from vision-language pre-training methods like CLIP (Radford et al., 2021), BEiT (Bao et al., 2022), BEiT-v2 (Peng et al., 2022), MVP (Wei et al., 2022), and COTS (Lu et al., 2022).

Hence, there is a clear advantage to developing effective decoders for the original ViT structures in order to leverage those powerful representations. Previous works, such as UPerNet (Xiao et al., 2018) and DPT (Ranftl et al., 2021), have primarily focused on hierarchical feature maps and neglected the distinctive characteristics of the plain Vision Transformer. Consequently, these methods introduce computation-intensive operations while offering limited performance gains, as shown in Fig. 1.

A recent trend in several works, such as SETR (Zheng et al., 2021) or Segmenter (Strudel et al., 2021), aims to develop decoders specifically tailored for the Plain ViT architecture. However, these designs are often an extension of per-pixel classification techniques derived from traditional convolution-based decoders. For example, SETR’s decoder (Zheng et al., 2021) uses a sequence of convolutions and bilinear up-sampling to increase the ViT’s extracted feature maps gradually. It then applies a naive MLP to the extracted features to perform pixel-wise classification, which isolates the neighboring contexts surrounding the pixel. Current pixel-wise classification decoder designs overlook the

importance of contextual learning when assigning labels to each pixel.

Another prevalent issue in deep networks, including Transformer, is ‘catastrophic forgetting’ (French, 1999; Kirkpatrick et al., 2017), where the model’s performance on previously learned tasks deteriorates as it learns new ones (Shao & Feng, 2022; Wang et al., 2022, ?; Phan et al., 2022). This limitation poses significant challenges for the application of deep segmentation models in dynamic real-world environments. Recently, the rapid development of the foundation model pre-trained on large-scale data has sparked interest among researchers in studying its transferability across various downstream tasks (Ostapenko et al., 2022). These models are capable of extracting powerful and generalized representations, which has led to a growing interest in exploring their extensibility to new classes and tasks while retaining the previously learned knowledge representations (Ramasesh et al., 2022; Wu et al., 2022).

Inspired by these challenges, this paper aims to develop plain Vision Transformer-based model for effective semantic segmentation without resorting to hierarchical backbone. As self-supervision and multi-modality pre-training continue to evolve, we anticipate that the plain vision transformer will learn enhanced visual representations. Consequently, decoders for dense tasks are expected to adapt more flexibly and efficiently to these representations.

In light of these research gaps, we propose *SegViTv2*—a novel, efficient segmentation network that features a plain Vision Transformer and exhibits robustness against forgetting. We introduce a novel Attention-to-Mask (ATM) module that operates as a lightweight component for the SegViT decoder. Leveraging the non-linearity of cross-attention learning, our proposed ATM employs learnable class tokens as queries to pinpoint spatial locations that exhibit high compatibility with each class. We advocate for regions affiliated with a particular class to possess substantial similarity values that correspond to the respective class token.

As depicted in Fig. 2, the ATM generates a meaningful similarity map that accentuates regions with a strong affinity towards the ‘Table’ and ‘Chair’ categories. By simply implementing a Sigmoid operation, we can transform these similarity maps into mask-level predictions. The computation of the mask scales linearly with the number of pixels, a negligible cost that can be integrated into any backbone to bolster segmentation accuracy. Building upon this efficient ATM module, we present a novel semantic segmentation paradigm that utilizes the cost-effective structure of plain ViT, referred to as SegViT. Within this paradigm, multiple ATM modules are deployed at various layers to extract segmentation masks at different scales. The final prediction is the summation of the outputs derived from these layers.

To alleviate the computational burdens of plain Vision Transformers (ViTs), we introduce the *Shrunk* and *Shrunk++*

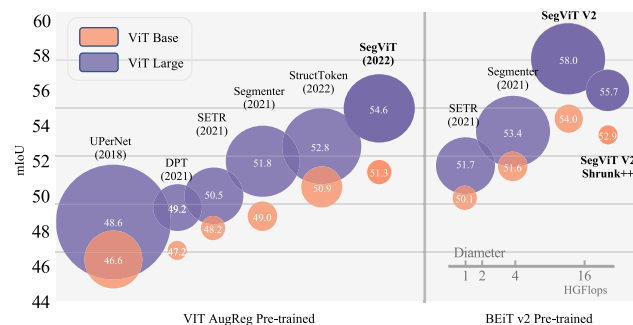


Fig. 1 Comparison with previous methods in terms of performance and efficiency on ADE20K dataset. The bubbles in the accompanying graph represent the ViT Base and ViT Large models, respectively, with the size of each bubble corresponding to the FLOPs of the variant segmentation methods. SegViT-BEiT v2 Large achieves state-of-the-art performance with a **58.0%** mIoU on the ADE20K validation set. Additionally, our efficient, optimized version, SegViT-Shrunk-BEiT v2 Large, saves half of the GFLOPs compared to UPerNet, significantly reducing computational overhead while maintaining a competitive performance of **55.7%**




Fig. 2 The overall concept of our Attention-to-Mask decoder. ATM learns the similarity map for each category by capturing the cross-attention between the class tokens and the spatial feature map (Left). Sigmoid is applied to produce category-specific masks, highlighting

structures, which incorporate query-based downsampling (QD) and query-based upsampling (QU). The proposed QD employs a 2×2 nearest neighbor downsampling technique to obtain a sparser token mesh, reducing the number of tokens involved in attention computations. In *Shrunk++*, we extend QD to edge-aware query-based downsampling (EQD). EQD selectively preserves tokens situated at object edges, as they possess more discriminative information. Consequently, QU recovers the discarded tokens within the object's homogeneous body, reconstructing high-resolution features crucial for accurately dense prediction. Integrating the *Shrunk++* structure with the ATM module as the decoder, our SegViTv2 achieves computational reductions of up to 50% while maintaining competitive performance.

We further adapt our SegViTv2 framework for continual learning. Leveraging the robust, generalized representation of the foundational model, this paper investigates its adaptability to new classes and tasks, ensuring retention of prior knowledge. Recent techniques in continual semantic segmentation (CSS) aim to replay old data (Maracani et al., 2021; Cha et al., 2021) or distill knowledge from the previous model to mitigate model divergence (Cermelli et al., 2020; Phan et al., 2022; Zhang et al., 2022). These methods fine-tune parameters related to old tasks, which can disrupt the previously learned solutions and result in forgetting. In contrast, our proposed SegViT supports learning new classes without interfering with previously acquired knowledge. We strive to establish a forget-free SegViT framework, achieved by incorporating a new ATM module dedicated to new tasks while freezing all old parameters. Consequently, the proposed SegViT architecture has the potential to eliminate the issue of forgetting.

Our key contributions can be summarized as follows:

- We introduce the Attention-to-Mask (ATM) decoder module, a potent and efficient tool for semantic segmentation. For the first time, we exploit spatial information present in attention maps to generate mask predictions

the area with high similarity to the corresponding class (Middle). ATM enhances the semantic representations by encouraging the feature to be similar to the target class token and dissimilar to other tokens

for each category, proposing a new paradigm for semantic segmentation.

- We present the *Shrunk++* structure, applicable to any plain ViT backbone, which alleviates the intrinsically high computational expense of the *non-hierarchical* ViT while maintaining competitive performance, as illustrated in Fig. 1. We are the first work capitalizing on edge information to decrease and restore tokens for efficient computation. Our *Shrunk++* version of SegViTv2, tested on the ADE20K dataset, achieves a mIoU of 55.7%, with a computational cost of 308.8 GFLOPs, marking a reduction of approximately 50% compared to the original SegViT (637.9 GFLOPs).
- We propose a new SegViT architecture capable of continual learning with nearly zero forgetting. To our knowledge, we are the first work seeking to completely freeze all parameters for old classes, thereby nearly obliterating the issue of catastrophic forgetting.

2 Related Work

2.1 Semantic Segmentation

Semantic segmentation aims to partition an image into regions with meaningful categories. Fully Convolutional Networks (FCNs) used to be the dominant approach to this task. To enlarge the receptive field, several approaches (Zhao et al., 2017; Chen et al., 2018) propose dilated convolutions or apply spatial pyramid pooling to capture contextual information at multiple scales. Most semantic segmentation methods aim to classify each pixel directly using a classification loss. This paradigm naturally partitions images into different classes.

Various methods have achieved significant advancements by integrating Transformers into the semantic segmentation task. Early works (Liu et al., 2021; Dong et al., 2022) directly adapt the transformer encoder, designed for classification, into semantic segmentation by fine-tuning it together with

segmentation decoders such as UPerNet (Xiao et al., 2018). Recent approaches (Xie et al., 2021; Strudel et al., 2021; Cheng et al., 2021) have focused on designing the overall segmentation framework to achieve better adaptation. For instance, SETR (Zheng et al., 2021) views semantic segmentation as a sequence-to-sequence task and proposes a pure Transformer encoder combined with a standard convolution-based decoder. SegFormer (Xie et al., 2021) employs a hierarchical encoder design to extract features from fine-to-coarse levels and a lightweight decoder design for efficient prediction. However, the SegFormer decoder adopts the pyramid structure by fusing multi-scale features, which is specialized for hierarchical ViTs such as Swin Transformer (Liu et al., 2021). The above-mentioned methods aim to design either a naive convolution-based decoder or a pyramid-structure decoder for hierarchical base models. Nonetheless, designing an effective decoder specialized for plain ViTs remains an open research question.

Recently, several segmentation methods propose a universal framework that unifies multiple tasks, including instance segmentation, semantic segmentation, and object detection. For example, Mask DINO (Li et al., 2022) extends DINO with a mask prediction branch, achieving promising results in the instance, panoptic, and semantic segmentation tasks. Mask2Former (Cheng et al., 2022) enhances MaskFormer (Cheng et al., 2021) by introducing deformable multi-scale attention in the decoder and a masked cross-attention mechanism. OneFormer (Jain et al., 2022) represents a universal image segmentation framework with a multi-task train-once design, outperforming specialized models in various tasks.

Recent methods (Cheng et al., 2021; Strudel et al., 2021; Zhang et al., 2021) propose decoupling the per-pixel classification into image partitioning and region classification. For image partitioning, they use learnable tokens as mask embeddings and associate them with the extracted feature map to generate object masks. For region classification, the learnable tokens are fed to a classifier to predict the class corresponding to each mask. This paradigm enables global segmentation and alleviates the burden on the decoder to perform per-pixel classification, resulting in state-of-the-art performance (Cheng et al., 2021). While previous works use generic tokens for mask generation, this work explicitly utilizes class-specific tokens to enhance the semantics of mask embeddings, thereby improving segmentation accuracy.

2.2 Mask-Oriented Segmentation

Compared to previous mask-oriented segmentation techniques such as MaskFormer (Cheng et al., 2021) and Mask2Former (Cheng et al., 2022), our method presents several novel conceptual differences and advantages. Specifically, our approach is tailored to address semantic segmentation problems by assigning each class to a fixed token

and generating the corresponding mask directly. In contrast, MaskFormer relies on Hungarian matching, with each learnable query corresponding to spatial information instead of category information. Our Attention-to-Mask (ATM) approach eliminates the need for positional embedding, as we utilize the attention map between the class token and the feature map. Our overarching goal is to adapt Plain Vision Transformers for dense prediction, as recent studies have demonstrated that self-supervised learning (He et al., 2022; Chen et al., 2022; Touvron et al., 2022; Peng et al., 2022) and multimodal learning (Radford et al., 2021) are enhanced by hierarchical ViT structures. Our approach enhances the representation ability of class tokens by applying transformer blocks.

Previous CNN-based decoders, such as OCRNet (Yuan et al., 2019) and K-Net (Zhang et al., 2021), have demonstrated the effectiveness of the attention mechanism in modeling contextual information. For example, K-Net utilizes semantic kernels (one kernel for each class) and performs convolution operations to generate the semantic mask. In contrast, our proposed ATM module integrates cross-attention mechanisms, allowing for more effective contextual learning. While OCRNet (Yuan et al., 2019) applies cross-attention from the class token to the feature map to enhance feature representations, it still employs a standard linear predictor in the decoder to produce the segmentation map. On the other hand, our proposed ATM module is specifically designed for generating segmentation outputs, paving the way for future research on effective decoders for plain ViT. Additionally, existing convolution-based attention networks such as OCRNet (Yuan et al., 2019), K-Net (Zhang et al., 2021), and DANet (Fu et al., 2019) adopt the traditional per-pixel classification framework for segmentation generation. In contrast, our proposed SegViT decouples segmentation into mask prediction and classification, which proves advantageous for establishing connections between the class proxy and language representations (Zhou et al., 2022), as well as facilitating continual learning.

2.3 Transformers for Vision

In the realm of image classification tasks, attention-based transformer models have emerged as powerful alternatives to standard convolution-based networks. The original ViT (Dosovitskiy et al., 2021) represents a plain, non-hierarchical architecture. However, there have been several advancements in the field of hierarchical transformers, such as PVT (Wang et al., 2021), Swin Transformer (Liu et al., 2021), Twins (Chu et al., 2021), SegFormer (Xie et al., 2021), and P2T (Wu et al., 2022). These hierarchical transformer models inherit certain design elements from convolution-based networks, including hierarchical structures, pooling, and downsampling with convolutions. Consequently, they can be seamlessly employed

as direct replacements for convolutional-based networks and can be coupled with existing decoder heads for tasks such as semantic segmentation.

2.4 Self-Supervised Vision Transformers

Self-supervised learning has emerged as a powerful technique for pretraining visual models, eliminating the need for labeled data. One notable self-supervised method is MAE (He et al., 2022) (Masked Autoencoder), which trains a vision transformer to reconstruct masked regions of input images. This approach results in a high generalization capacity. Another significant method is CLIP (Radford et al., 2021) (Contrastive Language-Image Pre-Training), which involves joint training of a vision transformer and a language model on a large corpus of text and images, leading to the creation of a comprehensive knowledge store. CAE (Chen et al., 2022) aims to learn image representations that are invariant to context changes and effectively capture underlying semantic content. Furthermore, iBot (Zhou et al., 2022) performs masked visual learning using an online tokenizer and self-distillation mechanism, facilitating semantic representation learning. In our approach, we leverage attention to masks to optimize the extraction of dense hidden representations, thereby enhancing the segmentation capability of our model.

2.5 Plain-Backbone Decoders

For dense prediction tasks, such as semantic segmentation, the high-resolution feature maps produced by the backbone are vital for preserving spatial details. In typical hierarchical transformer models, techniques such as FPN (Lin et al., 2017) or dilated backbone are employed to generate high-resolution feature maps by merging features from different levels. However, when it comes to a plain, non-hierarchical transformer backbone, the resolution remains the same across all layers. SETR (Zheng et al., 2021) proposed a straightforward approach to address segmentation tasks by treating transformer outputs from the base model in a sequence-to-sequence perspective. Segmenter (Strudel et al., 2021) combines class embeddings and transformer patch embeddings and applies several self-attention layers on the combined tokens to learn discriminative embeddings. In their approach, the class tokens are used as input to the ViT backbone, resulting in increased computational complexity. In contrast, our SegViT introduces the class tokens as input to the ATM, the Attention-to-Mask module, thereby reducing computational costs while still benefiting from the integration of class tokens.

2.6 Continual Learning

Continual learning (CL) aims to address the issue of forgetting, ensuring consistent performance on previously learned classes while adapting to new ones (Chen and Liu, 2016). Most CL methods propose regularization techniques for convolution-based networks (Li & Hoiem, 2018; Douillard et al., 2020; Kang et al., 2022; Peng et al., 2021) or expand the network architectures to accommodate new tasks (Yan et al., 2021), thereby avoiding the need to store and replay old data. In recent years, efforts have also emerged to prevent forgetting in Transformer models. Dytox (Douillard et al., 2022) dynamically learns new task tokens, which are then utilized to make the learned embeddings more relevant to the specific task. Lifelong ViT (Wang et al., 2022) and contrastive ViT (Wang et al., 2022) introduce cross-attention mechanisms between tasks through external key vectors, and they slow down the changes to these keys to mitigate forgetting. Despite the use of complex mechanisms to prevent forgetting, these methods still require fine-tuning of the network for new classes, which can result in interference with previously learned knowledge.

In the field of semantic segmentation, recent research has been devoted to addressing the forgetting issue in continual learning. However, in addition to forgetting, continual semantic segmentation (CSS) also encounters the problem of "background shift." This refers to the situation where foreground object classes from previous tasks are mistakenly classified as background in the current task (Cermelli et al., 2020). REMINDER (Phan et al., 2022) tackles forgetting in CSS by utilizing class similarity to identify the classes that are more likely to be forgotten. It then focuses on revising those specific classes to mitigate the forgetting problem. RCIL (Zhang et al., 2022) introduces a two-branch convolutional network, with one branch frozen and the other trained to prevent forgetting. At the end of each learning step, the trainable branch is merged with the frozen branch, which can introduce model interference. However, it is worth noting that existing CSS and CL techniques typically involve fine-tuning certain parts of the network dedicated to the old tasks. Unfortunately, this fine-tuning process can lead to forgetting as the model diverges from the previously learned solution.

3 Method

In this section, we first introduce the overall architecture of our proposed SegViT model for semantic segmentation. Then, we discuss the *Shrunk* and *Shrunk++* architectures

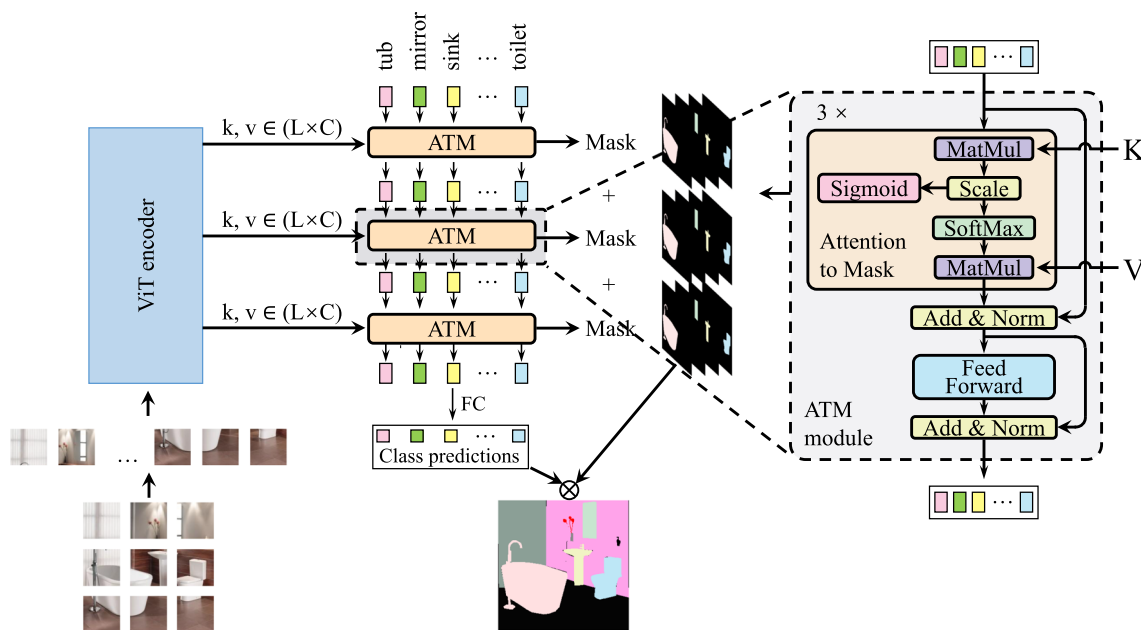


Fig. 3 The overall SegViT structure with the ATM module. The Attention-to-Mask (ATM) module inherits the typical transformer decoder structure. It takes in randomly initialized class embeddings as queries and the feature maps from the ViT backbone to generate keys and values. The outputs of the ATM module are used as the input queries for the next layer. The ATM module is carried out sequentially with inputs from different layers of the backbone as keys and values in

a cascade manner. A linear transform is then applied to the output of the ATM module to produce the class predictions for each token. The mask for the corresponding class is transferred from the similarities between queries and keys in the ATM module. We have removed the self-attention mechanism in ATM decoder layers further improve the efficiency while maintaining the performance

designed to reduce the model’s computational cost. Lastly, we explore the adaptation of our SegViT model for the context of continual semantic segmentation to minimize forgetting.

3.1 Overall SegViT Architecture

SegViT comprises a ViT-based encoder responsible for feature extraction and a decoder used to learn the segmentation map. For the encoder, we designed the ‘Shrunk’ structure to reduce the computational overhead associated with the plain ViT. Regarding the decoder, we introduce a novel lightweight module named Attention-to-Mask (ATM). This module generates class-specific masks denoted as M and class predictions denoted as P , which determine the presence of a particular class in the image. The mask outputs from a stack of ATM modules are combined and then multiplied by the class predictions to obtain the final segmentation output. Figure 3 illustrates the overall architecture of our proposed SegViT.

3.1.1 Encoder

Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, the plain vision transformer backbone reshapes it into a sequence of tokens $\mathcal{F}_0 \in$

$\mathbb{R}^{L \times C}$, where $L = \frac{HW}{p^2}$, P is the patch size, and C is the number of channels. To capture positional information, learnable position embeddings of the same size as \mathcal{F}_0 are added. Subsequently, the token sequence \mathcal{F}_0 is processed by m transformer layers to produce the output. The output tokens for each layer are defined as $[\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_m] \in \mathbb{R}^{L \times C}$. For a plain vision transformer like ViT, the number of tokens are high and remains constant for each layer. Processing a substantial number of tokens for every layer results in elevated computational costs for plain ViT. We denote a plain ViT-based encoder as the ‘Single’ structure. To mitigate computational costs, we introduce the *Shrunk* and *Shrunk++* structures, tailored to create a more computationally efficient ViT-based encoder. Further details regarding the *Shrunk* structure can be found in Sect. 3.2.

3.1.2 Decoder

Attention-to-Mask (ATM) Cross-attention can be described as the mapping between two sequences of tokens, denoted as $\{\mathbf{v}_1, \mathbf{v}_2\}$. In our case, we define two token sequences: $\mathcal{G} \in \mathbb{R}^{N \times C}$ with a length N equal to the number of classes, and $\mathcal{F}_i \in \mathbb{R}^{L \times C}$. To enable cross-attention, linear transformations are applied to each token sequence, resulting in the query (Q), key (K), and value (V) representations. This process is described by Eq. (1).

$$\begin{aligned} Q &= \phi_q(\mathcal{G}) \in \mathbb{R}^{N \times C}, \\ K &= \phi_k(\mathcal{F}_i) \in \mathbb{R}^{L \times C}, \\ V &= \phi_v(\mathcal{F}_i) \in \mathbb{R}^{L \times C}. \end{aligned} \quad (1)$$

The similarity map is calculated by computing the dot product between the query and key representations. Following the scaled dot-product attention mechanism, the similarity map and attention map are calculated as follows:

$$S(Q, K) = \frac{QK^T}{\sqrt{d_k}} \in \mathbb{R}^{N \times L}, \quad (2)$$

$$\text{Attention}(\mathcal{G}, \mathcal{F}_i) = \text{Softmax}(S(Q, K))V \in \mathbb{R}^{N \times C},$$

where $\sqrt{d_k}$ is a scaling factor with d_k equals to the dimension of the keys.

The shape of the similarity map $S(Q, K)$ is determined by the lengths of the two token sequences, N and L . The attention mechanism updates \mathcal{G} by performing a weighted sum of V , where the weights are derived from the similarity map after applying the softmax function along the L dimension.

In dot-product attention, the softmax function is used to concentrate attention exclusively on the token with the highest similarity. However, we believe that tokens other than those with maximum similarity also carry meaningful information. Based on this intuition, we have designed a lightweight module that generates semantic predictions more directly. To this end, we assign \mathcal{G} as the class embeddings for the segmentation task, and \mathcal{F}_i as the output of layer i of the ViT backbone. A semantic mask is paired with each token in \mathcal{G} to represent the semantic prediction for each class. The binary mask M is defined as follows:

$$\text{Mask}(\mathcal{G}, \mathcal{F}_i) = \text{Sigmoid}(S(Q, K)) \in \mathbb{R}^{N \times L}. \quad (3)$$

The masks have a shape of $N \times L$, which can be reshaped to $N \times \frac{H}{P} \times \frac{W}{P}$ and bilinearly upsampled to the original image size $N \times H \times W$. As depicted in the right section of Fig. 3, the ATM mechanism produces masks as an intermediate output during cross-attention.

The final output tokens $Z \in \mathbb{R}^{L \times C}$ from the ATM module are utilized for classification. A fully connected layer (FC) parameterized by $W \in \mathbb{R}^{C \times 2}$ followed by the Softmax function is used to predict whether the object class is present in the image or not. The class predictions $\mathcal{P} \in \mathbb{R}^{N \times 2}$ are formally defined as:

$$\mathcal{P} = \text{Softmax}(WZ). \quad (4)$$

Here, $P_{c,1}$ indicates the likelihood of class c appearing in the image. For simplicity, we refer to P_c as the probability score for class c .

The output segmentation map for class $O_s \in \mathbb{R}^{H \times W}$ is obtained by element-wise multiplication of the reshaped class-specific mask M_c and its corresponding prediction score P_c : $O_c = P_c \odot M_c$. During inference, the label is assigned to each pixel i by selecting the class with the highest score using $\text{argmax}_c O_{i,c}$.

Indeed, plain base models like ViT do not inherently possess multiple stages with features of different scales. Consequently, structures such as Feature Pyramid Networks (FPN) that merge features from multiple scales are not applicable to them.

Nevertheless, features from layers other than the last one in ViT contain valuable low-level semantic information, which can contribute to improving performance. In SegViT, we have developed a structure that leverages feature maps from different layers of ViT to enrich the feature representations. This allows us to incorporate and benefit from the rich low-level semantic information present in those feature maps.

SegViT is trained via the classification loss and the binary mask loss. The classification loss (\mathcal{L}_{cls}) minimizes cross-entropy between the class prediction and the actual target. The mask loss ($\mathcal{L}_{\text{mask}}$) consists of a focal loss (Lin et al., 2017) and a dice loss (Milletari et al., 2016) for optimizing the segmentation accuracy and addressing sample imbalance issues in mask prediction. The dice loss and focal loss respectively minimize the dice and focal scores between the predicted masks and the ground-truth segmentation. The final loss is the combination of each loss, formally defined as:

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \lambda_{\text{focal}} \mathcal{L}_{\text{focal}} + \lambda_{\text{dice}} \mathcal{L}_{\text{dice}} \quad (5)$$

where λ_{focal} and λ_{dice} are hyperparameters that control the strength of each loss function. Previous mask transformer methods such as MaskFormer (Cheng et al., 2021) and DETR (Carion et al., 2020) have adopted the binary mask loss and fine-tuned their hyperparameters through empirical experiments. Hence, for consistency, we directly use the same values as MaskFormer and DETR for the loss hyperparameters: $\lambda_{\text{focal}} = 20.0$ and $\lambda_{\text{dice}} = 1.0$.

3.2 Shrunk Structure for Efficient Plain ViT Encoder

Recent efforts, such as DynamicViT (Rao et al., 2021), TokenLearner (Ryoo et al., 2021), and SPViT (Kong et al., 2022), propose token pruning techniques to accelerate vision transformers. However, most of these approaches are specifically designed for image classification tasks and, as a result, discard valuable information. However, when adapting these techniques to semantic segmentation tasks, they may fail to preserve high-resolution features that are necessary for accurate dense prediction tasks.

In this paper, we introduce the *Shrunk* structure. This method employs query-based down-sampling (QD) to prune

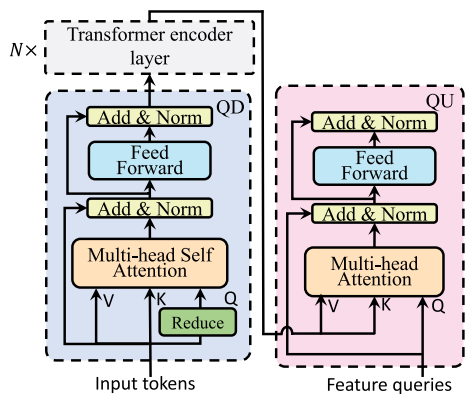


Fig. 4 Architecture of the proposed query-downsampling (QD) layer (blue block) and the query-up-sampling (QU) layer (block). The QD layer uses an efficient down-sampling technique (green block) and removes less informative input tokens used for the query. The QU layer takes a set of trainable query tokens and learns to recover the discarded tokens using multi-head attention (Color figure online)

the input token sequence \mathcal{F}_i and uses query up-sampling (QU) to retrieve the discarded tokens, ensuring preservation of fine-detail features vital for semantic segmentation. The overall architecture of QD and QU is illustrated in Fig. 4.

For QD, we have re-designed the Transformer encoder block (Vaswani et al., 2017) and incorporated efficient down-sampling operations to specifically reduce the number of query tokens. In a Transformer encoder layer, the computational cost is directly influenced by the number of query tokens, and the output size is determined by the query token size. To mitigate the computational burden while maintaining information integrity, a viable strategy is to selectively reduce the number of query tokens while preserving the key

and value tokens. This approach allows for an effective reduction in the output size of the current layer, leading to reduced computational costs for subsequent layers.

For QU, we perform up-sampling using a token sequence - either predefined or inherited - that has a higher resolution than the query tokens. The key and value tokens are taken from the token sequence obtained from the backbone, which typically has a lower resolution. The output size is dictated by the query tokens with higher resolution. Through the cross-attention mechanism, information from the key and value tokens is integrated into the output. This process facilitates a non-linear merging of information and demonstrates an upsampling behavior, effectively increasing the resolution of the output.

As illustrated in Fig. 5, our proposed *Shrunk* structure incorporates the QD and QU modules. Specifically, we integrate a QD operation at the middle depth of the ViT backbone, precisely at the 8th layer of a 24-layer backbone. The QD operation downsamples the query tokens using a 2×2 nearest neighbor downsampling operation, resulting in a feature map size reduction to $1/32$. However, such downsampling can potentially cause information loss and performance degradation. To mitigate this issue, prior to applying the QD operation, we employ a QU operation to the feature map. This involves initializing a set of query tokens with a resolution of $1/16$ to store the information. Subsequently, as the downsampled feature map progresses through the remaining backbone layers, it is merged and upsampled using another QU operation alongside the previously stored $1/16$ high-resolution feature map. This iterative process ultimately generates a

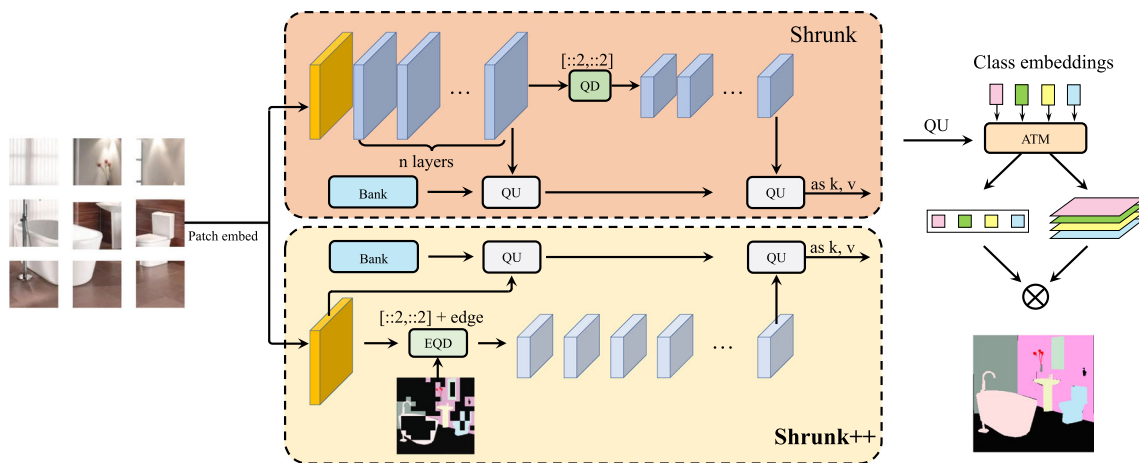


Fig. 5 Illustrations of the *Shrunk* and *Shrunk++*. In the diagram, the and boxes respectively refer to the transformer encoder block and the patch embedding block. In SegVit (Zhang et al., 2022), the proposed *Shrunk* structure employs query downsampling (QD) on the middle-level features to preserve the information. In the new *Shrunk++* architecture, we introduce the Edged Query Downsampling (EQD) tech-

nique which consolidates every four adjacent tokens into one token and additionally includes the tokens that contain edges. This enhancement enables downsampling operations to take place before the first layer without significant performance degradation, offering computational savings for the initial layers of the *Shrunk* model. The edge information is extracted using a lightweight parallel edge detection head

1/16 high-resolution feature map enriched with semantic information processed by the backbone.

Despite the effectiveness of the proposed *Shrunk* approach in maintaining performance, it requires the integration of the QD operation within the intermediate layers of the backbone. This necessity arises due to the fact that shallow layers primarily capture low-level features, and applying downsampling to these layers would result in significant information loss. Consequently, these low-level layers continue to be computed at a higher resolution, limiting the potential reduction in computational cost.

To address this limitation and further optimize the backbone, we introduce SegViTv2 using a novel architecture called *Shrunk++*. In this architecture, we incorporate an edge detection module in the QD section and introduce an Edged Query Downsampling (EQD) technique to update the QD process. In addition to the 2×2 nearest downsampling operation that eliminates every 4 consecutive tokens, our approach aims to retain tokens that contain multiple categories, specifically tokens that contain an edge. By preserving the 2×2 sparse tokens, we retain important semantic information, while also preserving the edge tokens to retain detailed spatial information. By retaining both types of information, we minimize the loss of valuable information and overcome the limitations associated with low-level layers. To extract edges, we add a separate branch using a lightweight multilayer perceptron (MLP) termed as the edge detection head that learns to detect edges from the input image. The edge detection head operates as an auxiliary branch, trained simultaneously with the main ATM decoder. This head processes the input image, which has the same dimensions as the backbone. Let the input image have C channels, aligned with the backbone. The Multi-Layer Perceptron (MLP) in this head consists of three layers, with dimensions C , $C/2$, and 2, respectively. Let I represent the input image, and the output of the MLP can be defined as $E = \text{MLP}(I; W_1, W_2, W_3)$, where W_1, W_2, W_3 are the weights for the three layers. The output E is then passed through a softmax activation function, resulting in $S = \text{Softmax}(E)$. To determine the confidence level of a token belonging to an edge, we apply a threshold τ . In our implementation, we set τ to 0.7. To obtain the ground-truth (GT) edge, we perform post-processing on the GT segmentation map Y . Since the input has been tokenized with a patch size of P , we tokenize the GT and reshape it into a sequence of tokens denoted as $Y \in \mathbb{R}^{(HW/P^2) \times P \times P}$, where the last two dimensions correspond to the patch dimensions. We consider a patch to contain an edge if there exists any edge pixel within the patch. We define the edge mask $Mask_i$ as follows:

$$Mask_i = \begin{cases} 1 & \text{if } \sum_{j,k} Y_{i,j,k} > 0, \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

For each element s_i in S , we create a binary edge mask M_i : $M_i = 1$, if $s_i \geq \tau$. The cross-entropy loss is computed between the generated edge mask M_i and the ground-truth edge mask Y_i : $\mathcal{L}_{\text{edge}} = - \sum_i Y_i \log(M_i) + (1 - Y_i) \log(1 - M_i)$. By incorporating the Edge Detection head as an auxiliary branch, the *Shrunk++* architecture effectively retains detailed spatial contexts throughout the query downsampling process, forming an Edge Query Downsampling (EQD) structure. This EQD structure effectively captures and retains edge information during sparse downsampling, significantly reducing computational overhead while maintaining performance. The integration of EQD enables the *Shrunk++* architecture to strike a remarkable balance between computational efficiency and maintaining high-performance levels.

3.3 Exploration on Continual Semantic Segmentation

Continual semantic segmentation aims to train a segmentation model in T steps without forgetting. At step t , we are given a dataset \mathcal{D}^t which comprises a set of pairs (X^t, Y^t) , where X^t is an image of size $H \times W$ and Y^t is the ground-truth segmentation map. Here, Y^t only consists of labels in current classes \mathcal{C}^t , while all other classes (i.e., old classes $\mathcal{C}^{1:t-1}$ or future classes $\mathcal{C}^{t+1:T}$) are assigned to the background. In continual learning, the model at step t should be able to predict all classes $\mathcal{C}^{1:t}$ in history.

3.3.1 SegViT for Continual Learning

Existing continual semantic segmentation methods (Zhang et al., 2022; Phan et al., 2022) propose regularization algorithms to preserve the past knowledge of a specific architecture, DeepLabV3. These methods focus on continual semantic segmentation for DeepLabV3 with a ResNet backbone, which has a less robust visual representation for distinguishing between different categories. Consequently, these methods require fine-tuning model parameters to learn new classes while attempting to retain knowledge of old classes. Unfortunately, adapting the old parameters dedicated to the previous task inevitably interferes with past knowledge, leading to catastrophic forgetting. In contrast, our proposed SegViT decouples class prediction from mask segmentation, making it inherently suitable for a continual learning setting. By leveraging the powerful representation capability of the plain vision transformer, we can learn new classes by solely fine-tuning the class proxy (i.e., the class token) while keeping the old parameters frozen. This approach eliminates the need for fine-tuning old parameters when learning new tasks, effectively addressing the issue of catastrophic forgetting.

During training on the current task t , we add a new sequence of learnable tokens $\mathcal{G}^t \in \mathbb{R}^{|\mathcal{C}^t| \times C}$, where $|\mathcal{C}^t|$ is the number of classes in the current task. To learn new classes,




Fig. 6 Overview of SegViT adapted for continual semantic segmentation. When learning a new task t , we grow and train a separate ATM and fully-connected layer to produce mask and class prediction. All the parameters dedicated to the old task $t - 1$, including ATM, FC layers, and the ViT encoder, are frozen. This prevents interfering with the old knowledge, which guarantees no forgetting

we grow and train new ATM modules and a fully-connected layer for mask prediction and mask classification. For simplicity, we ignore the parallel structure of ATM modules. A single ATM module refers to multiple ATM modules. Let A^t and W^t denote the ATM module and the weights of the fully connected (FC) layer for task t . All parameters for prior tasks, including the ViT encoder, the ATM module, and the FC layer, are completely frozen. Figure 6 illustrates the overview of our SegViT architecture adapted for continual semantic segmentation.

Given the encoder extracted features \mathcal{F}_T and the class tokens \mathcal{G}^t , the ATM produces the mask predictions M^t and the output tokens Z^t corresponding to the mask:

$$M^t, Z^t = \text{ATM}(\mathcal{G}^t, \mathcal{F}^T). \tag{7}$$

Based on Eq. (4), the class prediction \mathcal{P} is obtained by applying FC on the class token Z^t .

The prediction score S_c^t for each class c is multiplied by the corresponding mask M_c^t to get the segmentation map O_c^t for class c :

$$O_c^t = S_c^t \odot M_c^t, \tag{8}$$

where \odot denotes the element-wise multiplication. The segmentation \hat{O}^t is obtained by taking the class c having the highest score in every pixel, defined as

$$\hat{O}^t = \underset{c \in \mathcal{C}^t}{\text{argmax}} O_{i,c}^t \tag{9}$$

Based on the ground truth Y^t for task t , SegViT is trained using the loss function defined in Eq. (5). To produce the final segmentation across all tasks, we concatenate the individual outputs O^t from each task.

4 Experiments

4.1 Datasets

ADE20K (Zhou et al., 2017) is a challenging scene parsing dataset which contains 20, 210 images as the training set and 2, 000 images as the validation set with 150 semantic classes. *COCO-Stuff-10K* (Caesar et al., 2018) is a scene parsing benchmark with 9, 000 training images and 1, 000 test images. Even though the dataset contains 182 categories, not all categories exist in the test split. We follow the implementation of mmsegmentation (MMSegmentation, 2020) with 171 categories to conduct the experiments.

PASCAL-Context (Mottaghi et al., 2014) is a dataset with 4, 996 images in the training set and 5, 104 images in the validation set. There are 60 semantic classes in total, including a class representing ‘background’.

4.2 Implementation Details

4.2.1 Transformer Backbone

We employ the naive ViT (Dosovitskiy et al., 2021) as the backbone for our method. For our ablation studies, we primarily utilize the ‘Base’ variation, while also presenting results based on the ‘Large’ variant. Notably, variations in performance can arise due to different pre-trained weights, as indicated by Segmenter (Strudel et al., 2021). To ensure equitable comparisons, we adopt the pre-trained weights provided by Augreg (Steiner et al., 2021), aligning with practices employed in Strudel (Strudel et al., 2021) and StructToken (Lin et al., 2022). These weights stem from training on ImageNet-21k with strong data augmentation and regularization techniques (Steiner et al., 2021). To explore the maximum capacity and assess the upper bound of our method, we also conduct experiments using stronger base models such as DEiT v3 (Touvron et al., 2022) and BEiT v2 (Peng et al., 2022).

4.2.2 Training Settings

We use MMSegmentation (MMSegmentation, 2020) and follow the commonly used training settings. During training, we apply sequential data augmentation techniques, including random horizontal flipping, random resizing within a ratio of 0.5 to 2.0, and random cropping. For most settings, the cropping dimensions are set to 512×512 , except for PASCAL-Context where we use 480×480 , and for ViT-large backbone on ADE20K where we use 640×640 . The batch size is set to 16 for all datasets with a total iteration of 160k, 80k, and 80k for ADE20k, COCO-Stuff-10k, and PASCAL-Context respectively.

Table 1 Experiment results on the ADE20K val. split

Method	Backbone	Crop size	GFLOPs	mIoU (ss)	mIoU (ms)	Inf time (fps)
UPerNet (Xiao et al., 2018)	ViT-Base	512 × 512	443.9	46.6	47.5	16.07
DPT* (Ranftl et al., 2021)	ViT-Base	512 × 512	219.8	47.2	47.9	23.63
SETR-MLA* (Zheng et al., 2021)	ViT-Base	512 × 512	113.5	48.2	49.3	–
Segmenter* (Strudel et al., 2021)	ViT-Base	512 × 512	129.6	49.0	50.0	20.46
StructToken (Lin et al., 2022)	ViT-Base	512 × 512	171.5	50.9	51.8	14.22
MaskFormer (Cheng et al., 2021)	Swin-B(21K)	640 × 640	198.3	52.7	53.9	–
Mask2Former (Cheng et al., 2022)	Swin-B(21K)	640 × 640	223.4	53.9	55.1	12.43
SegViT (ours)	ViT-Base	512 × 512	120.9	51.3	53.0	31.52
SegViT (<i>Shrunk++</i> , Ours)	BEiTv2-Base	512 × 512	74.4	52.9	53.3	25.03
SegViT (ours)	BEiTv2-Base	512 × 512	120.9	54.0	54.9	23.59
DPT* (Ranftl et al., 2021)	ViT-Large [†]	640 × 640	800.0	49.2	49.5	9.38
UPerNet (Xiao et al., 2018)	ViT-Large [†]	640 × 640	1993.9	48.6	50.0	3.88
SETR-MLA (Zheng et al., 2021)	ViT-Large	512 × 512	368.6	48.6	50.3	5.17
MCIBI (Jin et al., 2021)	ViT-Large	512 × 512	> 400	–	50.8	–
Segmenter (Strudel et al., 2021)	ViT-Large [†]	640 × 640	671.8	51.8	53.6	4.73
StructToken (Lin et al., 2022)	ViT-Large [†]	640 × 640	774.6	52.8	54.2	4.1
KNet+UPerNet (Zhang et al., 2021)	Swin-L(21K)	640 × 640	659.3	52.2	53.3	11.28
MaskFormer (Cheng et al., 2021)	Swin-L(21K)	640 × 640	378.1	54.1	55.6	10.21
Mask2Former (Cheng et al., 2022)	Swin-L(21K)	640 × 640	402.7	56.1	57.3	8.81
SegViT (ours)	ViT-Large [†]	640 × 640	637.9	54.6	55.2	9.37
SegViT (<i>Shrunk</i> , ours)	ViT-Large [†]	640 × 640	373.5	53.9	55.1	10.18
SegViT (<i>Shrunk++</i> , ours)	ViT-Large [†]	640 × 640	209.1	53.0	54.9	10.26
SegViT (<i>Shrunk++</i> , ours)	BEiTv2-Large	512 × 512	210.3	55.1	56.1	9.82
SegViT (ours)	BEiTv2-Large	512 × 512	374.0	56.5	58.0	9.39
SegViT (<i>Shrunk++</i> , ours)	BEiTv2-Large	640 × 640	308.8	55.7	57.0	9.38
SegViT (ours)	BEiTv2-Large	640 × 640	637.9	58.0	58.2	6.25

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. ‘ms’ means that mIoU is calculated using multi-scale inference. ‘†’ means the models use the backbone weights pre-trained by AugReg (Steiner et al., 2021). ‘*’ represents the model reproduced under the same settings as the official repo. The GFLOPs are measured at single-scale inference with the given crop size. We report inference speed for our SegViT and reproduce previous methods in terms of Frame Per Second (FPS) on a single A100 device

4.2.3 Evaluation Metric

We use the mean Intersection over Union (mIoU) as the metric to evaluate the performance. ‘ss’ means single-scale testing and ‘ms’ test time augmentation with multi-scaled (0.5, 0.75, 1.0, 1.25, 1.5, 1.75) inputs. All reported mIoU scores are in a percentage format. All reported computational costs in GFLOPs are measured using the fvcore¹ library.

4.3 Comparisons with the State-of-the-Art Methods


4.3.1 Results on ADE20K

Table 1 reports the comparison with the state-of-the-art methods on ADE20K validation set using ViT backbone. The SegViT uses the ATM module with multi-layer inputs from

the original ViT backbone, while the *Shrunk* is the one that conducts QD to the ViT backbone and saves 50% of the computational cost without sacrificing too much performance. Our approach achieves a state-of-the-art mIoU of 58.2% (MS) with the BEiTv2 Large backbone. To ensure a fair comparison, we evaluate our SegViT module with the BEiT-v2 large backbone on a crop size of 512 × 512, which consumes 374.0 GFLOPs. Our approach achieves a slightly better performance of 56.5% mIoU compared to Mask2former-Swin-L, which achieves 56.1% with 402.7 GFLOPs on a crop size of 640 × 640. Additionally, our *Shrunk* version offers around a 50% reduction in computational cost (308.8 GFLOPs), while delivering competitive performance with a mIoU of 57.0% (MS). Optimizing SegViT with ViT-Large using the proposed *Shrunk++* reduces the computational cost of *Shrunk* by 3.05 times, while preserving the mIoU. Figure 7 shows the visual results of different segmentation methods. In contrast to other methods that often confuse

¹ <https://github.com/facebookresearch/fvcore>

Fig. 7 Visuals results of different segmentation networks and plain ViT backbones on the ADE20K validation set (Zhou et al., 2017). It includes the following models: **a** Segmenter (Strudel et al., 2021) with ViT large, **b** StructToken (Lin et al., 2022) with ViT large, **c** UPerNet (Xiao et al., 2018) with BEiT large, and **d** SegViT V2 with BEiTv2 large. The results demonstrate that our methods effectively generate accurate segmentation masks and unlock the potential of plain ViT. Zoom in for a better view



(a) Segmenter ViT-Large(51.8)

(b) StructToken ViT-Large(52.8)

(c) UPerNet BEiT-Large(56.3)

(d) SegViT V2 BEiT2-Large(58.0)

similar classes and misclassify related concepts, our SegViT stands out by more precise object boundary delineation and achieving accurate segmentation of complete objects, even in cluttered scenes.

4.3.2 Results on COCO-Stuff-10K

Table 2 shows the result on the COCO-Stuff-10K dataset. Our method achieves 50.3% which is higher than the previous state-to-the-art StrucToken by 1.2% with less computational cost. Our *Shrunk* version achieves 49.4% mIoU with 224.8

GFLOPs, which is similar to the computational cost of a dilated ResNet-101 backbone but with much higher performance. By extending SegViT with the effective *Shrunk++*, we significantly decrease its GFLOPs by 1.82 times, while retaining a competitive mIoU.

4.3.3 Results on PASCAL-Context

Table 3 shows the results on the PASCAL-Context dataset. We follow HRNet (Sun et al., 2019) to evaluate our method and report the results under 59 classes (without background)

Table 2 Experiment results on the COCO-Stuff-10K test . split

Method	Backbone	GFLOPs	mIoU (ms)
DANet (Fu et al., 2019)	Dilated-ResNet-101	289.3	39.7
MaskFormer (Cheng et al., 2021)	ResNet-101-fpn	81.7	39.8
EMANet (Li et al., 2019)	Dilated-ResNet-101	247.4	39.9
SpyGR (Li et al., 2020)	ResNet-101-fpn	>80	39.9
OCRNet (Yuan et al., 2020)	HRNetV2-W48	167.9	40.5
GINet (Wu et al., 2020)	JPU-ResNet-101	>200	40.6
RecoNet (Chen et al., 2020)	Dilated-ResNet-101	>200	41.5
ISNet (Jin et al., 2021)	Dilated-ResNeSt-101	228.3	42.1
MCIBI (Jin et al., 2021)	ViT-Large	>380	44.9
StructToken (Lin et al., 2022)	ViT-Large	>400	49.1
SenFormer (Bousselham et al., 2021)	Swin-Large	>400	50.1
SegViT (<i>Shrunk</i> , ours)	ViT-Large	224.8	49.40
SegViT (ours)	ViT-Large	383.9	50.30
SegViT (<i>Shrunk++</i> , ours)	BEiTv2-Large	213.3	50.54
SegViT (ours)	BEiTv2-Large	388.2	53.46

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. Following published methods, we report the results with multi-scale inference (denoted by ‘ms’). The GFLOPs is measured at single scale inference with a crop size of 512×512

Table 3 Experimental results on the PASCAL-Context val . split

Method	Backbone	GFLOPs	mIoU ₅₉ (ms)	mIoU ₆₀ (ms)
RefineNet (Lin et al., 2017)	ResNet-152	–	–	47.3
UNet++ (Zhou et al., 2018)	ResNet-101	–	47.7	–
PSPNet (Zhao et al., 2017)	Dilated-ResNet-101	157.0	47.8	–
Ding et al. (Ding et al., 2018)	ResNet-101	–	51.6	–
EncNet (Zhang et al., 2018)	Dilated-ResNet-101	192.1	52.6	–
HRNet (Sun et al., 2019)	HRNetV2-W48	82.7	54.0	48.3
NRD (Zhang et al., 2021)	ResNet-101	42.9	54.1	49.0
GFFNet (Li et al., 2020)	Dilated-ResNet-101	–	54.3	–
EfficientFCN (Liu et al., 2020)	ResNet-101	52.8	55.3	–
OCRNet (Yuan et al., 2020)	HRNetV2-W48	143.9	56.2	–
SETR-MLA (Zheng et al., 2021)	ViT-Large	318.5	–	55.8
Segmenter (Strudel et al., 2021)	ViT-Large	346.2	–	59.0
SenFormer (Bousselham et al., 2021)	Swin-Large	–	64.0	–
SegViT (<i>Shrunk</i> , ours)	ViT-Large	186.9	62.3	57.40
SegViT (ours)	ViT-Large	321.6	65.3	59.30
SegViT (<i>Shrunk++</i> , ours)	BEiTv2-Large	179.3	64.91	59.92
SegViT (ours)	BEiTv2-Large	329.7	67.14	61.63

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. Following published methods, we report the results with multi-scale inference (denoted by ‘ms’). mIoU₅₉: mIoU averaged over 59 classes (without background). mIoU₆₀: mIoU averaged over 60 classes (59 classes plus background). Both metrics were used in the literature, and we report for the 60 classes. The GFLOPs are measured at single scale inference with a crop size of 480×480

and 60 classes (with background). Using full SegViT structure without adopting *Shrunk* or *Shrunk++*, we reach mIoU of 67.14% and 61.63% respectively for those two metrics, outperforming the state-of-the-art methods using the ViT backbones with less computational cost. By applying

Shrunk and *Shrunk++* architecture, the computational cost in terms of GLOPs is reduced by 42% and 45%, respectively. Among all approaches evaluated on the PASCAL-Context dataset, SegViTv2 with *Shrunk++* achieves the best trade-off between accuracy and efficiency.

Table 4 Comparisons between our proposed ATM module with SETR (Zheng et al., 2021)

Decoder	Loss	mIoU (ss)
SETR	CE loss	46.5
ATM	CE loss	47.0 (+ 0.5)
ATM	\mathcal{L}_{mask} loss	49.6 (+ 3.1)

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. ‘CE loss’ indicates the cross-entropy loss commonly used in semantic segmentation. The experiments on the ADE20k dataset are carried out using the ViT-Base backbone

4.4 Ablation Study

In this section, we conduct extensive ablation studies to show the effectiveness of our proposed methods.

4.4.1 Effect of the ATM Module

We conducted an analysis to evaluate the impact of using the proposed ATM module as an encoder. The results are summarized in Table 4. To establish a baseline for comparison, we introduced SETR-naive, which utilizes two 1×1 convolutions to directly derive per-pixel classifications from the final layer of the ViT-Base transformer output. From the results, it is evident that applying the ATM module under the supervision of a conventional cross-entropy loss leads to a performance improvement of 0.5%. However, the performance gains become much more substantial when we decouple the classification and mask prediction processes, supervising each separately. This approach results in a significant performance boost of 3.1%, highlighting the efficacy of the ATM module in enhancing semantic segmentation performance.

4.4.2 Ablation of the Feature Levels

The effects of using multiple-layer inputs from the backbone to the ATM modules are presented in Table 5. The incorporation of feature maps from lower layers leads to a notable performance improvement of 1.3%. We further investigated the impact of including more layers of features and observed additional gains in performance. After empirical testing, we determined that utilizing three layers yielded optimal results, resulting in an overall mIoU boost of 1.7%. These ablation studies confirm the effectiveness of our proposed ATM decoder and highlight the advantage of incorporating multi-layer features into the segmentation structure. This integration significantly enhances the performance of semantic segmentation tasks.

Table 5 Results of using different layer inputs to the SegViT structure on ADE20K dataset using ViT-Base as the backbone

	Used layers	mIoU (ss)
Single layer	[12]	49.6
Cascade	[6, 12]	50.9 (+ 1.3)
Cascade	[6, 8, 12]	51.3 (+ 1.7)
Cascade	[3, 6, 9, 12]	51.2 (+ 1.6)

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. Involving multi-layer features leads to obvious performance gains

Table 6 The experiments use the Swin-Tiny (Liu et al., 2021) backbone and are carried out on the ADE20K dataset

Method	mIoU (ss)	GFLOPs
Maskformer (Cheng et al., 2021)	46.7	57.3
Mask2former (Cheng et al., 2022)	47.7	73.7
SegViT (ours)	47.1	48.0

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. The GFLOPs are measured at single-scale inference with a crop size of 512×512

4.4.3 SegViT on Hierarchical Base Models

We conducted an analysis to evaluate the performance of SegViT on hierarchical base models. For comparison, we selected two competitive methods, Maskformer (Cheng et al., 2021) and Mask2former (Cheng et al., 2022). The results presented in Table 6 indicate that, even though our method was not specifically designed for hierarchical base models, we are still able to achieve competitive performance while maintaining computational efficiency. This demonstrates the applicability of our SegViT approach to various types of ViT-Base models.

4.4.4 Ablation of *Shrunk* and *Shrunk++* Strategies

In this section, we analyze the effectiveness of the different SegViT structures. Table 7 presents the effects of various techniques employed in each SegViT structure, including query upsampling (QU), query downsampling (QD), token-squeezing (TS) techniques, and segmentation heads. Applying the ATM head to the ‘Single’ structure yields a notable performance improvement of 6.67% compared with using the SETR head. This demonstrates the effectiveness of the ATM head in enhancing the performance of the baseline structure. However, applying QD to the ‘Single’ structure with the ATM head leads to a performance drop of 2.7%, suggesting the occurrence of information loss during the downsampling

Table 7 Ablation results of *Shrunk* and *Shrunk++* version on the ADE20K dataset

Structure	QD	QU	QD_{layer}	QD_{method}	Head	mIoU (ss)	GFLOPs
Single	–	–	–	–	SETR	46.5	107.3
Single	–	–	–	–	ATM	49.6 (+ 3.1)	115.8
Naive <i>Shrunk</i>	✓	–	6	2×2	ATM	46.9 (+ 0.4)	74.1
<i>Shrunk</i>	✓	✓	6	2×2	ATM	50.0 (+ 3.5)	97.1
Nearest-TS	✓	✓	0	3×3	ATM	38.9(– 7.6)	32.8
Nearest-TS	✓	✓	0	2×2	ATM	43.3(–3.2)	46.1
Shrunk++	✓	✓	0	3×3 -Edge	ATM	47.9(+ 1.4)	69.3
Shrunk++	✓	✓	0	2×2 -Edge	ATM	49.9(+ 3.4)	74.6

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. We explored various shrink strategies. The GFLOPs are measured at single-scale inference with a crop size of 512×512 on the ViT-Base backbone. QD: query-based downsampling. QU: query-based upsampling. QD_{layer} indicates which layer to apply the QD. QD_{method} indicates the downsampling method for QD

Table 8 Ablation results of different decoder methods with their corresponding feature merge types and loss types

Decoder	Multi-level Features		Loss Types			mIoU (ss)
	FPN	Token Merge	Pixel level	Dot product	Attention Mask	
SETR-MLA (Zheng et al., 2021)	✓		✓			48.2
Segmenter (Strudel et al., 2021)			✓			49.0
MaskFormer (Cheng et al., 2021)	✓			✓		46.7
Ours-Variant 1					✓	49.6
Ours-Variant 2		✓		✓		50.6
Ours		✓			✓	51.2

ViT-Base is employed as the backbone for all the variants

phase. Importantly, incorporating QU restores the performance. QU helps recover the discarded information from QD and reconstructs the high-resolution feature map, which is crucial for dense prediction tasks. Jointly leveraging QU and QD, the *Shrunk* architecture achieves optimal performance while reducing computational costs by 16.15% in comparison to the ‘Single’ structure.

In the proposed *Shrunk++* structure, we analyze the performance of two main token-squeezing techniques: nearest downsampling and edge-aware downsampling. It is important to note that token squeezing is directly applied to the first layer of the network for optimal computational efficiency. Applying naive nearest downsampling with a 3×3 kernel reduces the GFLOPs of the *Shrunk* structure without token-squeezing by a factor of 2.97. However, reducing the computational cost with 3×3 and 2×2 nearest downsampling leads to a performance drop of 13%. In contrast, by incorporating an additional edge extractor into our *Shrunk++* architecture, we significantly improve the mIoU, achieving performance on par with *Shrunk*, i.e., 49.9% mIoU, with a minor increase in computational cost to 74.6 GFLOPs. The edge-aware downsampling technique preserves the edge details, thereby preserving discriminative features for dense predictions. Among the different settings, the 2×2 + Naive MLP Edge setting achieves an optimal balance between performance and efficiency.

4.4.5 Ablation Studies on Decoder Variances

Different decoder methods are associated with specific feature merge types and loss types. In Table 8, we compare the designs of various decoders on a plain ViT backbone. For hierarchical base models like Swin, the resolution of the feature maps in each stage is reduced. Consequently, the adoption of a Feature Pyramid Network (FPN) is necessary to obtain feature maps with larger resolutions and rich semantic information. However, in Table 8, we observe that the FPN structure does not perform well with plain vision transformers. With plain ViT base models, the resolution remains constant, and the feature map of the final layer encapsulates the most comprehensive semantic information. Hence, our proposed method, which utilizes tokens to merge features from different levels, achieves superior performance. By simply replacing the FPN structure with the ATM-based token merge, we improve the performance from 46.7% to 50.6%. Regarding the loss type, the pixel-level loss refers to the conventional cross-entropy loss applied to the feature map. The dot product loss corresponds to the loss utilized in Carion et al. (2020) and Cheng et al. (2021). Attention mask loss indicates the direct application of mask supervision to the similarity map generated by the ATM during attention calculation. Incorporating loss supervision on the attention mask, as in our method, leads to a performance improvement of 0.6%.

Table 9 Ablation of the QD module in terms of the targets and methods to down-sample

Applied to	Methods	mIoU (ss)
Q	Conv	44.5
Q, K, V	Nearest	52.6
Q	Nearest	53.9

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. The experiments are carried out on the ViT-Large backbone of ADE20K dataset

4.4.6 Ablation for the QD Module

The motivation behind using QD is to leverage the pre-trained weights of the backbone. As shown in Table 9, using a stride-2 convolution with learnable parameters to down-sample the query will disturb the pre-trained weights, leading to a notable decline in performance. Applying down-sampling to both the query and the key-value pairs would inevitably lead to information loss during the down-sampling process, which is evident in the lower performance. Our results show that applying 2×2 nearest down-sampling exclusively to the query in the QD module yields better results. This approach allows us to preserve the pre-trained weights of the backbone while achieving the desired down-sampling effect.

4.5 Application 1: A Better Indicator for Feature Representation Learning

4.5.1 Background

Semantic segmentation serves as a fundamental vision task that has been extensively employed in previous research to assess the representation learning capabilities of weakly, fully, and self-supervised base models (He et al., 2022; Chen et al., 2022; Touvron et al., 2022; Peng et al., 2022). In prior

work, the UPerNet decoder structure has been commonly used for semantic segmentation. However, the UPerNet decoder may not be a suitable indicator for evaluating the feature representation ability of the base model. This is primarily due to its heavier computational requirements and slower convergence rate. Additionally, variations in feature representation acquired by the base model can be substantial due to diverse training strategies during the fine-tuning process on semantic segmentation datasets. Consequently, the task of semantic segmentation may not adequately evaluate the feature representation ability of pre-trained models.

4.5.2 Experiment Settings

In this section, we extensively evaluate our proposed SegViT across diverse weakly, fully, and self-supervised vision transformers, including those proposed by He et al. (2022), Chen et al. (2022), Touvron et al. (2022), and the BEiT model Peng et al. (2022). We demonstrate that our method outperforms UPerNet Xiao et al. (2018) in both self-supervised and multi-modality base models, achieving state-of-the-art performance. Notably, our approach achieves superior performance to UPerNet while utilizing only 5% of the computational cost in terms of the decoder head. Table 10 illustrates that our proposed SegViT head consistently outperforms UPerNet across all base models. For the ViT-Base, our method improves the performance of UPerNet on the CLIP model by 1.16% while significantly reducing the computational cost. Similar findings are evident for ViT-Large base models. Furthermore, compared to UPerNet, our proposed SegViT's decoder head exhibits a better alignment between the growth trend of segmentation accuracy and the classification accuracy on ImageNet. This clearly demonstrates the superior efficiency of our SegViT

Table 10 Comparisons for various ViT pre-training schedules on the validation set of ADE20K

Backbone	SegViT mIoU	Head FLOPs	UPerNet mIoU	Head FLOPs	ImageNet Acc
MAE Base (He et al., 2022)	49.22 (▲ 1.12)	6.89 (▼ 329.73)	48.1	336.62	83.66
CLIP Base (Radford et al., 2021)	50.76 (▲ 1.16)	6.89 (▼ 329.73)	49.6	336.62	80.20
CAE Base (Chen et al., 2022)	50.42 (▲ 0.22)	6.89 (▼ 329.73)	50.2	336.62	83.90
iBot Base (Zhou et al., 2022)	50.58 (▲ 0.58)	6.89 (▼ 329.73)	50.0	336.62	84.00
Augreg Base* [†] (Steiner et al., 2021)	51.30 (▲ 2.66)	6.89 (▼ 329.73)	48.6	336.62	85.49
DEiT v3 Base [†] (Touvron et al., 2022)	52.40 (▲ 0.60)	6.89 (▼ 329.73)	51.8	336.62	85.70
BEiT v2 Base [†] (Peng et al., 2022)	53.97 (▲ 0.47)	6.89 (▼ 329.73)	53.5	336.62	86.50
Augreg Large* [†] (Steiner et al., 2021)	54.60 (▲ 2.50)	16.36 (▼ 1,366.33)	52.1	1382.69	85.59
DEiT v3 Large* [†] (Touvron et al., 2022)	55.81 (▲ 1.21)	16.36 (▼ 1,366.33)	54.6	1382.69	87.70
BEiT v2 Large [†] (Peng et al., 2022)	58.00 (▲ 1.30)	16.36 (▼ 868.28)	56.7	884.64	87.30

All results are reported in single-scale inference. The default configuration for these base models is pre-trained on ImageNet-1K with 224×224 resolutions. ‘*’ means the models use the backbone weights pre-trained with 384×384 resolutions. ‘[†]’ means the base models pre-trained on imagenet-21K. The proposed SegViT head has a less computational cost and performs better than UPerNet among all pre-training variants

head compared to UPerNet, making it a more suitable indicator for representation learning in base models.

4.6 Application 2: Continual Semantic Segmentation

The decoupling of class prediction and mask segmentation in our proposed SegViT decoder makes it inherently well-suited for continual learning settings. This characteristic allows us to learn new classes by solely fine-tuning the class proxy (the class token), leveraging the powerful representation ability of the plain vision transformer while keeping the old parameters frozen. To validate the effectiveness of this new approach to continual learning, we conducted experiments following standard settings adopted by prior studies.

4.6.1 Experiment Settings

Continual Semantic Segmentation (CSS) has two settings (Cermelli et al., 2020; Douillard et al., 2021): disjoint and overlapped. In the disjoint setup, all pixels in the images at each step belong to the previous classes or the current class. In the overlapped setting, the dataset of each step contains all the images that have pixels of at least one current class, and all pixels from previous and future tasks are labeled as background. The overlapped setting is more realistic and challenging, thus we evaluate the performance of the overlapped setup on the ADE20k dataset.

Following prior studies (Phan et al., 2022; Cermelli et al., 2020; Douillard et al., 2021), we perform three experiments: adding 50 classes after training with 100 classes (100–50 setting with 2 steps), adding 50 classes each time after training with 50 classes (50–50 setting with 3 steps), adding 10 classes each time sequentially after training with 100 classes (100–10 setting with 6 steps).

4.6.2 Baselines

We conducted a comprehensive comparison of our proposed method against state-of-the-art Continual Semantic Segmentation (CSS) techniques, including RCIL (Zhang et al., 2022), PLOP (Douillard et al., 2021), REMINDER (Phan et al., 2022), SDR (Michieli & Zanuttigh, 2021), and MiB (Cermelli et al., 2020). To ensure fair comparisons, existing methods were evaluated using DeepLabV3 (Chen et al., 2017) with ResNet101 and ViT-Base backbones that were pre-trained on ImageNet-21k. The reported results for PLOP, RCIL, and REMINDER were obtained based on the codebases provided by the respective authors. Furthermore, we included the performance of the Oracle model, which represents the upper bound achieved by jointly training on all available data, serving as a benchmark for each method.

4.6.3 Metrics

We evaluate the model performance by five mIoU metrics. First, we compute mIoU for the base classes C^0 , which reflects model rigidity: the model's resilience to catastrophic forgetting. Second, we compute mIoU for all incremented classes $C^{1:T}$, which measures plasticity: the model capacity in learning new tasks. Third, we compute the mIoU of all classes in $C^{0:T}$ (*all*), which shows the overall performance of models. Fourth, we report the average of mIoU (*avg*) measured step after step as proposed by Douillard et al. (2021), which evaluates performance over the entire continual learning process. To ensure fair comparisons, we evaluate the relative performance of each CSS method in terms of relative mIoU reduction compared with its Oracle model, jointly trained on all data.

5 Results and Discussion

Table 11 shows the results of different CSS methods on ADE20k. Our SegViT-CL consistently outperforms existing methods in *all* mIoU for both settings. In terms of mIoU reduction, the proposed SegViT-CL only decreases the mIoU of the Oracle model by 2.2% on the 100–50 setting, which is two times better than the second-best method, RCIL with ResNet backbone with 4.6% reduction. This substantial enhancement over existing methods underlines the effectiveness of our proposed method in the continual semantic segmentation paradigm. On a long CL setting 100–10 with 6 tasks, ours is almost forgetting-free with a marginal mIoU reduction of 0.3%, while recent CSS methods significantly suffer from forgetting with at least 5.4% mIoU reduction. Using the ViT backbone, existing methods including MiB, REMINDER, and PLOP still suffer from high mIoU reductions. Compared with the Oracle, MiB (Cermelli et al., 2020), PLOP (Douillard et al., 2021), and REMINDER (Phan et al., 2022) decrease the mIoU by 8.6%, 6.5% and 5.6% respectively on the 100–10 setting, demonstrating the sub-optimal performance of current CSS methods for ViT architecture. This highlights the need for developing a specialized ViT architecture that is robust to forgetting.

To evaluate the forgetting of every task on the 100–10 setting, we compute the performance drop at the last step compared with its initial mIoU when the model first learns the task. For example, the initial mIoU of task 2 is the mIoU of class 101–110 evaluated at step 2. Similarly, that of task 3 is the mIoU of class 111–120 reported at step 3. Table 12 shows the performance drop at the last step compared with the initial mIoU of each task. Averaged across 5 tasks, the mIoU only drops by 0.45%, which shows that SegViT is robust to forgetting across all tasks on the 100–10 setting. Figure 8 shows the mIoU on the base classes after incrementally training on

Table 11 CSS results on ADE20k in mIoU (%) on 100–50 and 100–10 settings

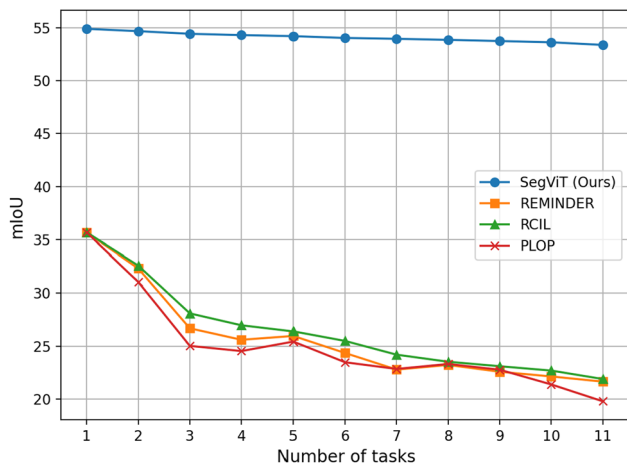
Method	100–50 (2 tasks)		101–150		all	Avg	100–10 (6 tasks)		101–150	All	Avg
	0–100	101–150	0–100	101–150			0–100	101–150			
ILT (Michieli & Zanuttigh, 2019)	18.29 (▼ 26.1)	14.40 (▼ 13.8)	17.00 (▼ 22.0)	29.42	0.11 (▼ 44.2)	3.06 (▼ 25.1)	1.09 (▼ 37.9)	12.56			
MiB (Cermelli et al., 2020)	40.52 (▼ 3.9)	17.17 (▼ 11.0)	32.79 (▼ 6.2)	37.31	38.21 (▼ 6.1)	11.12 (▼ 17.1)	29.24 (▼ 9.8)	35.12			
SDR (Michieli & Zanuttigh, 2021)	40.52 (▼ 3.8)	17.17 (▼ 11.0)	32.79 (▼ 6.2)	37.31	37.26 (▼ 7.1)	12.13 (▼ 16.1)	28.94 (▼ 10.1)	34.48			
PLOP (Douillard et al., 2021)	41.76 (▼ 2.6)	14.52 (▼ 13.7)	32.74 (▼ 6.3)	37.73	38.59 (▼ 5.8)	14.21 (▼ 14.0)	30.52 (▼ 8.5)	34.48			
REMINDER (Phan et al., 2022)	41.55 (▼ 2.8)	19.16 (▼ 9.0)	34.14 (▼ 4.9)	38.43	38.96 (▼ 5.4)	21.28 (▼ 6.9)	33.11 (▼ 5.9)	37.47			
RCIL (Zhang et al., 2022)	42.35 (▼ 2.0)	18.47 (▼ 9.7)	34.45 (▼ 4.6)	38.48	29.42 (▼ 15.0)	13.49 (▼ 14.0)	28.36 (▼ 10.0)	29.93			
Oracle-ResNet backbone	44.34	28.21	39.00	–	44.34	28.21	39.00	–			
MiB (Cermelli et al., 2020)	43.43 (▼ 3.2)	30.63 (▼ 4.3)	39.19 (▼ 3.6)	38.66	39.15 (▼ 7.5)	20.37 (▼ 14.5)	34.17 (▼ 8.6)	39.53			
PLOP (Douillard et al., 2021)	43.82 (▼ 2.8)	26.23 (▼ 8.7)	37.99 (▼ 4.8)	38.06	43.25 (▼ 3.4)	24.13 (▼ 10.8)	36.25 (▼ 6.5)	40.28			
REMINDER (Phan et al., 2022)	44.66 (▼ 2.0)	26.76 (▼ 8.1)	38.73 (▼ 4.0)	38.43	43.28 (▼ 3.4)	24.33 (▼ 10.6)	37.10 (▼ 5.6)	41.76			
Oracle-ViT backbone	46.63	34.90	42.75	–	46.63	34.90	42.75	–			
SegViT-CL (ours)	53.64 (▼ 0.5)	40.00 (▼ 5.6)	49.09 (▼ 2.2)	46.82	53.77 (▼ 0.3)	35.54 (▼ 10.0)	47.70 (▼ 3.6)	50.59			
Oracle	54.11	45.60	51.28	–	54.11	45.60	51.28	–			

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. The relative mIoU reduction compared with the joint training for each method is reported

Table 12 Performance drop (degree of forgetting) of all classes grouped by tasks on the 100–10 setting

Tasks	101–110	111–120	121–130	131–140	141–150	Avg
First time	34.93	39.78	41.10	36.22	27.95	35.99
Last time	34.51	39.30	40.86	35.09	27.95	35.54
Forgetting	▼ 0.42	▼ 0.48	▼ 0.24	▼ 1.12	▼ 0	▼ 0.45

We have utilized bold text in the tables to highlight the best or state-of-the-art (SOTA) benchmarks. We report the class mIoU when the model first learns the task, and the mIoU when the model last learns it

**Fig. 8** mIoU of recent CSS methods on the first 100 base classes after incrementally learning new tasks on 100–5 settings with 11 tasks

many tasks in 100–5, which is a long continual learning setting with 11 tasks. Overall, our SegViT achieves nearly zero forgetting for almost all tasks at the last step. In contrast to previous CSS methods which require partial fine-tuning, the proposed SegViT supports completely freezing old parameters, effectively eliminating any interference with previously acquired knowledge.

6 Conclusion

This paper presents SegViTv2, a novel approach for semantic segmentation using plain ViT transformer base models. The proposed method introduces a lightweight decoder head that incorporates the Attention-to-mask (ATM) module. Additionally, a *Shrunk++* structure is proposed to reduce the computational cost of the ViT encoder by 50% while maintaining competitive segmentation accuracy. Moreover, this work extends the SegViT framework to address the challenge of continual semantic segmentation, aiming to achieve nearly zero forgetting. By protecting the parameters of old tasks, SegViT effectively mitigates the impact of catastrophic forgetting. Extensive experimental evaluations conducted on various benchmarks demonstrate the superiority of SegViT over UPerNet, while significantly reducing computational costs. The introduced decoder head provides a robust and

cost-effective avenue for future research in the field of ViT-based semantic segmentation.

Acknowledgements This work was in part supported by the National Key R&D Program of China (No. 2022ZD0118700). Y. Liu's participation was in part supported by the start-up funding of The University of Adelaide. We express our gratitude to The University of Adelaide High-Performance Computing Services for providing the GPU Compute Resources, and to Mr. Wang Hui and Dr. Fabien Voisin for their valuable technical support for the training infrastructure.

Funding Open Access funding enabled and organized by CAUL and its Member Institutions.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Bao, H., Dong, L., Piao, S., Wei, F. (2022). BEiT: BERT pre-training of image transformers, in *International conference on learning representations*, [Online]. Available: <https://openreview.net/forum?id=p-BhZSz59o4>
- Bousselham, W., Thibault, G., Pagano, L., Machireddy, A., Gray, J., Chang, Y. H., Song, X. (2021). Efficient self-ensemble framework for semantic segmentation, *arXiv preprint arXiv:2111.13280*
- Caesar, H., Uijlings, J., Ferrari, V. (2018). Coco-stuff: Thing and stuff classes in context, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1209–1218.
- Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S. (2020). End-to-end object detection with transformers, in *Proceedings European conference on computer vision* (pp. 213–229), Springer.
- Cermelli, F., Mancini, M., Bulò, S. R., Ricci, E., Caputo, B. (2020). Modeling the background for incremental learning in semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9230–9239.
- Cha, S., Yoo, Y., Moon, T., et al. (2021). Ssul: Semantic segmentation with unknown label for exemplar-based class-incremental learning, in *Proceedings of the advances in neural information processing systems*, vol. 34, pp. 10919–10930.

- Chen, X., Ding, M., Wang, X., Xin, Y., Mo, S., Wang, Y., Han, S., Luo, P., Zeng, G., Wang, J. (2022). Context autoencoder for self-supervised representation learning, *arXiv preprint arXiv:2202.03026*.
- Chen, Z., Liu, B. (2016). Lifelong machine learning. Synthesis lectures on artificial intelligence and machine learning.
- Chen, L.-C., Papandreou, G., Schroff, F., Adam, H. (2017) Rethinking atrous convolution for semantic image segmentation, *arXiv preprint arXiv:1706.05587*
- Chen, L.-C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H. (2018). Encoder-decoder with atrous separable convolution for semantic image segmentation, in *Proceedings of the European conference on computer vision*, pp. 801–818.
- Chen, W., Zhu, X., Sun, R., He, J., Li, R., Shen, X., Yu, B. (2020). Tensor low-rank reconstruction for semantic segmentation, in *Proceedings European conference on computer vision* (pp. 52–69) Springer.
- Cheng, B., Misra, I., Schwing, A. G., Kirillov, A., Girdhar, R. (2022). Masked-attention mask transformer for universal image segmentation.
- Cheng, B., Schwing, A., Kirillov, A. (2021). Per-pixel classification is not all you need for semantic segmentation, *Proceedings of the advances in neural information processing systems*, vol. 34.
- Cheng, B., Schwing, A. G., Kirillov, A. (2021). Per-pixel classification is not all you need for semantic segmentation.
- Chu, X., Tian, Z., Wang, Y., Zhang, B., Ren, H., Wei, X., Xia, H., Shen, C. (2021). Twins: Revisiting the design of spatial attention in vision transformers, *Proceedings of the advances in neural information processing systems*, vol. 34.
- Ding, H., Jiang, X., Shuai, B., Liu, A. Q., Wang, G. (2018). Context contrasted feature and gated multi-scale aggregation for scene segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2393–2402.
- Dong, X., Bao, J., Chen, D., Zhang, W., Yu, N., Yuan, L., Chen, D., Guo, B. (2022). Cswin transformer: A general vision transformer backbone with cross-shaped windows, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 12124–12134.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2021). An image is worth 16x16 words: Transformers for image recognition at scale, *Proceedings International Conference on Learning Representations*
- Douillard, A., Chen, Y., Dapogny, A., Cord, M. (2021). Plop: Learning without forgetting for continual semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*
- Douillard, A., Cord, M., Ollion, C., Robert, T., Valle, E. (2020). Podnet: Pooled outputs distillation for small-tasks incremental learning, in *Proceedings European conference on computer vision* (pp. 86–102), Springer.
- Douillard, A., Ramé, A., Couairon, G., Cord, M. (2022). Dyttox: Transformers for continual learning with dynamic token expansion, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 9285–9295.
- French, R. M. (1999). Catastrophic forgetting in connectionist networks. *Trends in Cognitive Sciences*, 3(4), 128–135.
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H. (2019) Dual attention network for scene segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3146–3154.
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R. (2022). Masked autoencoders are scalable vision learners, in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009.
- Jain, J., Li, J., Chiu, M., Hassani, A., Orlov, N., Shi, H. (2022). One-former: One transformer to rule universal image segmentation, *arXiv preprint arXiv:2211.06220*
- Jin, Z., Liu, B., Chu, Q., Yu, N. (2021). Isnet: Integrate image-level and semantic-level context for semantic segmentation, in *Proceedings of the IEEE international conference on computer vision*, pp. 7189–7198.
- Kang, M., Park, J., Han, B. (2022). Class-incremental learning by knowledge distillation with adaptive feature consolidation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 16071–16080.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13), 3521–3526.
- Kong, Z., Dong, P., Ma, X., Meng, X., Niu, W., Sun, M., Shen, X., Yuan, G., Ren, B., Tang, H. et al. (2022). Spvit: Enabling faster vision transformers via latency-aware soft token pruning, in *Proceedings European conference on computer vision* (pp. 620–640), Springer.
- Li, X., Yang, Y., Zhao, Q., Shen, T., Lin, Z., Liu, H. (2020). Spatial pyramid based graph reasoning for semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8950–8959.
- Li, F., Zhang, H., Liu, S., Zhang, L., Ni, L. M., Shum, H. -Y. et al. (2022). Mask dino: Towards a unified transformer-based framework for object detection and segmentation,” *arXiv preprint arXiv:2206.02777*
- Li, X., Zhao, H., Han, L., Tong, Y., Tan, S., Yang, K. (2020). Gated fully fusion for semantic segmentation, in *Proceedings of the AAAI conference on artificial intelligence*, vol. 34, no. 07, pp. 11418–11425.
- Li, X., Zhong, Z., Wu, J., Yang, Y., Lin, Z., Liu, H. (2019). Expectation-maximization attention networks for semantic segmentation, in *Proceedings of the IEEE international conference on computer vision*, pp. 9167–9176.
- Li, X., Zhong, Z., Wu, J., Yang, Y., Lin, Z., Liu, H. (2019). Expectation-maximization attention networks for semantic segmentation, in *Proceedings of the IEEE international conference on computer vision*, pp. 9167–9176.
- Li, Z., & Hoiem, D. (2018). Learning without forgetting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40, 2935–2947.
- Lin, T. -Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S. (2017). Feature pyramid networks for object detection, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2117–2125.
- Lin, T. -Y., Goyal, P., Girshick, R., He, K., Dollár, P. (2017). Focal loss for dense object detection, in *Proceedings of the IEEE international conference on computer vision*, pp. 2980–2988.
- Lin, F., Liang, Z., He, J., Zheng, M., Tian, S., Chen, K. (2022). Struct-token: Rethinking semantic segmentation with structural prior.
- Lin, G., Milan, A., Shen, C., Reid, I. (2017). RefineNet: Multi-path refinement networks for high-resolution semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1925–1934.
- Liu, J., He, J., Zhang, J., Ren, J., Li, H. (2020). EfficientFCN: Holistically-guided decoding for semantic segmentation, in *Proceedings European conference on computer vision*
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B. (2021). Swin transformer: Hierarchical vision transformer using shifted windows, in *Proceedings of the IEEE international conference on computer vision*, pp. 10012–10022.
- Long, J., Shelhamer, J., Darrell, T. (2015). Fully convolutional networks for semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pp. 3431–3440.

- Lu, H., Fei, N., Huo, Y., Gao, Y., Lu, Z., Wen, J.-R. (2022). Cots: Collaborative two-stream vision-language pre-training model for cross-modal retrieval, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 15692–15701.
- Maracani, A., Michieli, U., Toldo, M., Zanuttigh, P. (2021). Recall: Replay-based continual learning in semantic segmentation, in *Proceedings of the IEEE international conference on computer vision*.
- Michieli, U., Zanuttigh, P. (2019). Incremental learning techniques for semantic segmentation, in *Proceedings of the IEEE conference on computer vision workshops*, pp. 3205–3212.
- Michieli, U., Zanuttigh, P. (2021). Continual semantic segmentation via repulsion-attraction of sparse and disentangled latent representations, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1114–1124.
- Milletari, F., Navab, N., Ahmadi, S.-A. (2016) V-net: Fully convolutional neural networks for volumetric medical image segmentation, in *3DV. IEEE*, pp. 565–571.
- MMSegmentation. (2020). MMSegmentation: OpenMMLab semantic segmentation toolbox and benchmark, <https://github.com/open-mmlab/mms Segmentation>
- Mottaghi, R., Chen, X., Liu, X., Cho, N. -G., Lee, S. -W., Fidler, S., Urtasun, R., Yuille, A. (2014). The role of context for object detection and semantic segmentation in the wild, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 891–898.
- Ostapenko, O., Lesort, T., Rodríguez, P., Arefin, M.R. , Douillard, A., Rish, I., Charlin, L. (2022). Continual learning with foundation models: An empirical study of latent replay, in *Conference on lifelong learning agents*. PMLR, pp. 60–91.
- Peng, Z., Dong, L., Bao, H., Ye, Q., Wei, F. (2022). BEiT v2: Masked image modeling with vector-quantized visual tokenizers.
- Peng, Y., Qi, J., Ye, Z., & Zhuo, Y. (2021). Hierarchical visual-textual knowledge distillation for life-long correlation learning. *International Journal of Computer Vision*, 129, 921–941.
- Phan, M. H., Phung, S. L., Tran-Thanh, L., Bouzerdoum, A. et al. (2022). Class similarity weighted knowledge distillation for continual semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pp. 16866–16875.
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J. et al. (2021). Learning transferable visual models from natural language supervision, in *International conference on machine learning*. PMLR, pp. 8748–8763.
- Ramasesh, V. V., Lewkowycz, A., Dyer, E. (2022). Effect of scale on catastrophic forgetting in neural networks, in *Proceedings of the international conference on learning representation*.
- Ranftl, R., Bochkovskiy, A., Koltun, V. (2021) Vision transformers for dense prediction, in *Proceedings of the IEEE international conference on computer vision*, pp. 12179–12188.
- Rao, Y., Zhao, W., Liu, B., Lu, J., Zhou, J., Hsieh, C.-J. (2021). Dynamicvit: Efficient vision transformers with dynamic token sparsification, in *Proceedings of the advances in neural information processing systems*, vol. 34, pp. 13937–13949.
- Ronneberger, O., Fischer, P., & Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. *Medical image computing and computer-assisted intervention* (pp. 234–241). Springer.
- Ryoo, M., Piergiovanni, A., Arnab, A., Dehghani, M., Angelova, A. (2021). Tokenlearner: Adaptive space-time tokenization for videos, *Proceedings of the advances in neural information processing systems*, vol. 34, pp. 12786–12797.
- Shao, C., Feng, Y. (2022) Overcoming catastrophic forgetting beyond continual learning: Balanced training for neural machine translation, *arXiv preprint arXiv:2203.03910*
- Steiner, A., Kolesnikov, A., Zhai, X., Wightman, R., Uszkoreit, J., Beyer, L. (2021). How to train your vit? Data, augmentation, and regularization in vision transformers.
- Strudel, R., Garcia, R., Laptev, I., Schmid, C. (2021) Segformer: Transformer for semantic segmentation, in *Proceedings of the IEEE international conference on computer vision*, pp. 7262–7272.
- Sun, K., Zhao, Y., Jiang, B., Cheng, T., Xiao, B., Liu, D., Mu, Y., Wang, X., Liu, W., Wang, J. (2019). High-resolution representations for labeling pixels and regions.
- Touvron, H., Cord, M., Jégou, H. (2022). Deit iii: Revenge of the vit, in *Computer Vision-ECCV. 17th European conference*, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV (pp. 516–533) Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., Polosukhin, I. (2017). Attention is all you need,” *Proceedings of the advances in neural information processing systems*, vol. 30.
- Wang, Z., Liu, L., Duan, Y., Kong, Y., Tao, D. (2022). Continual learning with lifelong vision transformer, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 171–181.
- Wang, Z., Liu, L., Kong, Y., Guo, J., Tao, D. (2022). Online continual learning with contrastive vision transformer, in *Proceedings European conference on computer vision* (pp. 631–650), Springer.
- Wang, W., Xie, E., Li, X., Fan, D.-P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L. (2021). Pyramid vision transformer: A versatile backbone for dense prediction without convolutions, in *Proceedings of the IEEE international conference on computer vision*, pp. 568–578.
- Wang, Z., Zhang, Z., Ebrahimi, S., Sun, R., Zhang, H., Lee, C.-Y., Ren, X., Su, G., Perot, V., Dy, J., et al. (2022). Dualprompt: Complementary prompting for rehearsal-free continual learning, in *Computer Vision-ECCV, 17th European Conference*, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI (pp. 631–648), Springer.
- Wang, Z., Zhang, Z., Lee, C.-Y., Zhang, H., Sun, R., Ren, X., Su, G., Perot, V., Dy, J., Pfister, T. (2022). Learning to prompt for continual learning, in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 139–149.
- Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., et al. (2020). Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(10), 3349–3364.
- Wei, L., Xie, L., Zhou, W., Li, H., Tian, Q. (2022). Mvp: Multimodality-guided visual pre-training, in *Proceedings European conference on computer vision* (pp. 337–353), Springer.
- Wu, T., Caccia, M., Li, Z., Li, Y.-F., Qi, G., Haffari, G. (2022). Pre-trained language model in continual learning: A comparative study, in *Proceedings of the international conference on learning representation*.
- Wu, Y. -H., Liu, Y., Zhan, X., Cheng, M. -M. (2022). P2t: Pyramid pooling transformer for scene understanding, *IEEE transactions on pattern analysis and machine intelligence*.
- Wu, T., Lu, Y., Zhu, Y., Zhang, C., Wu, M., Ma, Z., Guo, G. (2020). Ginet: Graph interaction network for scene parsing, in *Proceedings European conference on computer vision* (pp. 34–51), Springer.
- Xiao, T., Liu, Y., Zhou, B., Jiang, Y., Sun, J. (2018). Unified perceptual parsing for scene understanding, in *Proceedings European conference on computer vision*, pp. 418–434.
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P. (2021) “Segformer: Simple and efficient design for semantic segmentation with transformers,” *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34.
- Xu, Y., Zhang, J., Zhang, Q., Tao, D. (2022). “Rethinking hierarchies in pre-trained plain vision transformer, arXiv preprint [arXiv:2211.01785](https://arxiv.org/abs/2211.01785)
- Yan, S., Xie, J., He, X. (2021). Der: Dynamically expandable representation for class incremental learning, in *Proceedings of the IEEE*

- conference on computer vision and pattern recognition*, pp. 3014–3023.
- Yuan, Y., Chen, X., Chen, X., Wang, J. (2019). Segmentation transformer: Object-contextual representations for semantic segmentation, *arXiv preprint arXiv:1909.11065*
- Yuan, Y., Chen, X., Wang, J. (2020). Object-contextual representations for semantic segmentation, in *Proceedings of the European Conference on Computer Vision* Springer, pp. 173–190.
- Zhang, H., Dana, K., Shi, J., Zhang, Z., Wang, X., Tyagi, A., Agrawal, A. (2018). Context encoding for semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7151–7160.
- Zhang, W., Pang, J., Chen, K., Loy, C. C. (2021). K-net: Towards unified image segmentation, *Proceedings of the advances in neural information processing systems*, vol. 34.
- Zhang, B., Tian, Z., Shen, C. et al. (2021). Dynamic neural representational decoders for high-resolution semantic segmentation, vol. 34.
- Zhang, B., Tian, Z., Tang, Q., Chu, X., Wei, X., Shen, C., Liu, Y. (2022). Segvit: Semantic segmentation with plain vision transformers, in *Proceedings of the advances in neural information processing systems*.
- Zhang, C.-B., Xiao, J.-W., Liu, X., Chen, Y.-C., Cheng, M.-M. (2022). “Representation compensation networks for continual semantic segmentation, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2022, pp. 7053–7064.
- Zhang, W., Pang, J., Chen, K., & Loy, C. C. (2021). K-net: Towards unified image segmentation. *Advances in Neural Information Processing Systems*, 34, 10326–10338.
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J. (2017). Pyramid scene parsing network, in *Proceedings of the IEEE conference on computer vision and pattern recognition*.
- Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P. H. (2021). et al., Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers, in *Proceedings of the IEEE conference on computer vision and pattern Recognition*, pp. 6881–6890.
- Zhou, Z., Siddiquee, M.M.R., Tajbakhsh, N., Liang, J. (2018). Unet++: A nested U-net architecture for medical image segmentation, in *Proceedings of the deep learning in medical image analysis workshop*, pp. 3–11.
- Zhou, J., Wei, C., Wang, H., Shen, W., Xie, C., Yuille, A., Kong, T. (2022). ibot: Image bert pre-training with online tokenizer, *Proceedings of the international conference on learning representation*.
- Zhou, Z., Zhang, B., Lei, Y., Liu, L., Liu, Y. (2022). Zegclip: Towards adapting clip for zero-shot semantic segmentation, *arXiv preprint arXiv:2212.03588*
- Zhou, B., Zhao, H., Puig, X., Fidler, S., Barriuso, A., Torralla, A. (2017) Scene parsing through ade20k dataset, in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 633–641.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.