

Spectral Prompt Tuning: Unveiling Unseen Classes for Zero-Shot Semantic Segmentation

Wenhao Xu^{1,*}, Rongtao Xu^{2,3,*}, Changwei Wang^{2,3}, Shibiao Xu^{1,†},
Li Guo¹, Man Zhang¹, Xiaopeng Zhang²

¹School of Artificial Intelligence, Beijing University of Posts and Telecommunications

²State Key Laboratory of Multimodal Artificial Intelligence Systems, Institute of Automation, Chinese Academy of Sciences

³School of Artificial Intelligence, University of Chinese Academy of Sciences
shibiaoxu@bupt.edu.cn

Abstract

Recently, CLIP has found practical utility in the domain of pixel-level zero-shot segmentation tasks. The present landscape features two-stage methodologies beset by issues such as intricate pipelines and elevated computational costs. While current one-stage approaches alleviate these concerns and incorporate Visual Prompt Training (VPT) to uphold CLIP’s generalization capacity, they still fall short in fully harnessing CLIP’s potential for pixel-level unseen class demarcation and precise pixel predictions. To further stimulate CLIP’s zero-shot dense prediction capability, we propose SPT-SEG, a one-stage approach that improves CLIP’s adaptability from image to pixel. Specifically, we initially introduce Spectral Prompt Tuning (SPT), incorporating spectral prompts into the CLIP visual encoder’s shallow layers to capture structural intricacies of images, thereby enhancing comprehension of unseen classes. Subsequently, we introduce the Spectral Guided Decoder (SGD), utilizing both high and low-frequency information to steer the network’s spatial focus towards more prominent classification features, enabling precise pixel-level prediction outcomes. Through extensive experiments on two public datasets, we demonstrate the superiority of our method over state-of-the-art approaches, performing well across all classes and particularly excelling in handling unseen classes.

Introduction

Semantic segmentation is one of the fundamental tasks in computer vision, aiming to predict the class for each pixel in an image (Xu et al. 2023d, 2021b; Chen et al. 2021; Dong et al. 2021). Despite the existence of numerous related works (Lu et al. 2020; Dong et al. 2020; Xu et al. 2023b; Wang et al. 2023a), the success of deep semantic segmentation models heavily relies on a large amount of annotated training images, which requires significant efforts. In recent years, interest has been growing in unsupervised or weakly supervised semantic segmentation methods, including semi-supervised (Chen et al. 2021), weakly supervised (Xu et al. 2023a,c; Wang et al. 2023b), few-shot (Xie et al. 2021), and zero-shot semantic segmentation (Bucher et al. 2019; Pastore et al. 2021; Xian et al. 2019). Among

*These authors contributed equally.

†Shibiao Xu is the corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

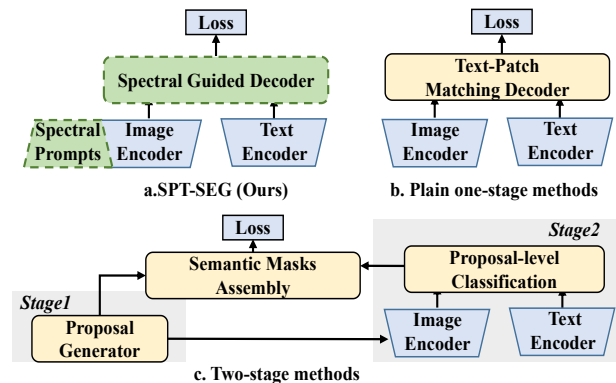


Figure 1: (a) Our SPT-SEG method demonstrates outstanding performance across all classes. (b) While yielding favorable results within the seen classes, it exhibits relatively poorer performance in the unseen classes. (c) Its performance is unsatisfactory across all classes.

them, zero-shot semantic segmentation tasks are particularly challenging and appealing, as they require generating accurate semantic segmentation results with only the semantic descriptions of the classes given.

To incorporate zero-shot capability into visual systems, researchers have proposed large-scale vision-and-language pretraining models, such as CLIP (Radford et al. 2021) and ALIGN (Jia et al. 2021a). Specifically, CLIP encodes semantic concepts into model parameters by contrastive training on a massive collection of image-text pairs, forming a zero-shot knowledge base for downstream tasks. However, contrastive pretraining mainly focuses on capturing image-level concepts. In CLIP, the training texts primarily describe the global context of images, and the encoded image and text embeddings are used together to compute contrastive losses. Consequently, CLIP is more suitable for image-level classification tasks (Zhou et al. 2022b,a; Lu et al. 2022; Zhang et al. 2022). The pretrained visual-language model CLIP (Radford et al. 2021) has recently found applications in various dense prediction tasks, including semantic segmentation (Pakhomov et al. 2021), referring segmentation (Wang et al. 2022), and object detection (Esmailpour

et al. 2022). In the zero shot semantic segmentation task, approaches like zsseg (Xu et al. 2021a) and Zegformer (Ding et al. 2022) adopt a similar strategy that requires two-stage processing: first generating region proposals and then feeding the cropped regions into CLIP for zero-shot classification. However, this strategy involves encoding images twice as FI 1(c), once for proposal generation and another for CLIP encoding of each proposal. This design introduces additional computational overhead and fails to fully leverage the knowledge in the CLIP encoder to guide the proposal generation stage. To streamline the process, ZegCLIP (Zhou et al. 2023) introduces a one-stage approach by incorporating visual prompt tuning into CLIP, then extending CLIP’s zero-shot capabilities from image-level to pixel-level.

The inclusion of Visual Prompt Tuning (VPT) in CLIP significantly enhances its downstream task generalization with few learnable parameters. However, since the original CLIP’s training primarily revolves around image-level contrastive learning, its features tend to emphasize only the most discriminative parts of objects. Even with the introduction of VPT, the observed phenomenon persists even during pre-training with image-level contrastive loss. Consequently, this phenomenon leads to incomplete and biased segmentation in dense prediction tasks.

Based on the aforementioned observations, we believe that further enhancing the image-to-pixel adaptability of CLIP (Radford et al. 2021) would contribute to improved zero-shot segmentation performance. Therefore, we propose an innovative one-stage method called SPT-SEG, as shown in Fig. 1(b). SPT-SEG differs from plain one-stage methods, as depicted in Fig.1(a). In our approach, we integrate spectral cues into the shallow layers of the CLIP visual encoder, which provides additional structural information that enhances the model’s comprehension of various object components. We also utilize high-frequency and low-frequency information to guide the alignment of text and pixels, directing the network’s spatial focus towards more salient classification features. The synergy of these two designs enhances the model’s semantic understanding and reasoning capabilities, effectively addressing the issues of inadequate pixel generalization and incomplete segmentation present in the current CLIP-based zero-shot semantic segmentation methods.

In summary, our contributions are listed as follows:

- We introduce **Spectral Prompt Tuning (SPT)**, which builds upon VPT by incorporating a small set of learnable spectral parameters. These parameters are integrated into the shallow layers of the CLIP visual encoder to introduce spectral information.
- We propose the **Spectral Guided Decoder (SGD)** layer, which is a novel component that utilizes high-frequency and low-frequency information to guide the matching process between textual and pixel representations.
- We comprehensively assess our method on two public datasets, and the results clearly show that our approach significantly surpasses state-of-the-art methods.

Related Work

Vision-Language Model. Extensive research has been conducted on Visual-Language Models (VLM)(Hong et al. 2021; Huang et al. 2021; Kamath et al. 2021; Kim, Son, and Kim 2021), showcasing significant advancements in downstream vision tasks, especially in settings with unannotated or restricted data. These tasks encompass diverse areas such as image retrieval(Liu et al. 2021), dense prediction (Rao et al. 2022), visual referring expression (Wang et al. 2022), and visual question answering (Jiang, Liu, and Zheng 2022). CLIP (Radford et al. 2021) is widely recognized as one of the most popular vision-language models. It is pretrained using contrastive learning on a massive dataset of 400 million text-image pairs. ALIGN (Jia et al. 2021b) utilized an even larger dataset, comprising 1.8 billion pairs, for pre-training its model. However, this larger dataset also introduced a significant amount of noise. In more recent works, CoCa (Yu et al. 2022) and Beit-V3 (Wang et al. 2023c) have further emphasized the superior performance of VLM pre-trained features.

Prompt Tuning. The concept of prompts originated from natural language processing and is mainly used in VLM to enhance its understanding of downstream specific tasks. By providing prompts, we can avoid massive parameter learning for VLM and instead use it as a fixed knowledge base, focusing only on task-relevant information. These prompts can be manually created for downstream tasks or automatically learned during fine-tuning. Full fine-tuning and linear probe (Gao et al. 2021) are two typical methods for adapting the VLM (i.e. CLIP) to downstream tasks. Full fine-tuning leads to a reduced VL representation of previously learned, while linear probe limits the zero-shot capability of CLIP. Inspired by the prompt learning in NLP, many works propose to adapt VLM by adding learnable tokens during end-to-end training. CoOp (Zhou et al. 2022b) introduced continuous prompt learning, where a set of continuous vectors are optimized end-to-end with down-stream supervision. Additionally, learnable prompts are applied by CoOp on the text encoder of CLIP to replace sub-optimal hand-crafted templates. Co-CoOp (Zhou et al. 2022a) highlights the poor performance of CoOp on novel classes and addresses the generalization problem by explicitly conditioning the prompts on image instances. Recently, prompting (Jia et al. 2022; Sandler et al. 2022) has been adapted to vision tasks. (Sandler et al. 2022) proposes memory tokens which is a set of learnable embedding vectors for each transformer layer. VPT (Jia et al. 2022) proposes similar ideas and investigates the generality and feasibility of visual prompting via extensive experiments spanning multiple kinds of recognition tasks across multiple domains and backbone architectures. Our research further extends the paradigm of visual prompt learning by introducing spectral prompt, addressing the limitations of previous visual prompt learning methods in fully leveraging the structural information of images and their limited adaptability to pixel-level tasks.

Zero-shot Semantic Segmentation. It remains a challenging task to achieve zero-shot semantic segmentation due to the presence of an imbalance problem in seen classes. Previous studies such as SPNet (Xian et al. 2019), ZS3 (Bucher

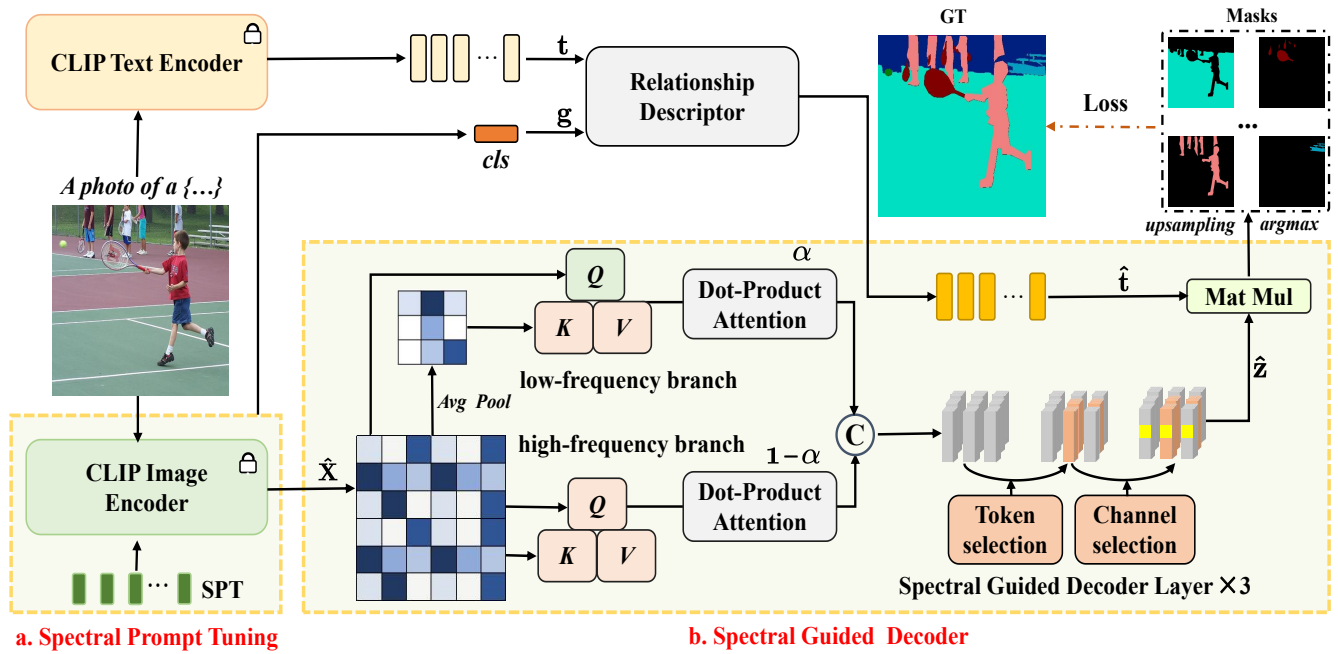


Figure 2: Overview of our proposed SPT-SEG. The main contribution of our work lies in two simple but effective designs (Red marks a,b in the figure): (a) Spectral prompt tuning which adds learnable spectral prompts to the first two layers of the CLIP’s visual encoder; (b) Spectral guided decoder which utilizes high- and low-frequency feature information to guide the text to match with pixels, and decodes the predicted results.

et al. 2019), CaGNet (Gu et al. 2020) and STRICT (Pastore et al. 2021) adopt strategies to improve the generalization ability of semantic mappings from visible to invisible classes. Since the popular pre-trained visual language model CLIP has shown powerful zero-shot classification capabilities, it has recently been applied to zero-shot semantic segmentation as well. Zegformer (Ding et al. 2022) and zsseg (Xu et al. 2021a) developed an extensive proposal generator and used CLIP to classify each region and then integrate the predictions. Previous studies, such as SPNet (Xian et al. 2019), ZS3 (Bucher et al. 2019), CaGNet (Gu et al. 2020), SIGN (Cheng et al. 2021), Joint (Baek, Oh, and Ham 2021), and STRICT (Pastore et al. 2021), adopt the approach of improving the generalization capability of semantic mapping from the classes that have been encountered to unseen ones. Recently, a two-stage paradigm (Ding et al. 2022; Xu et al. 2021a) has been proposed to explore the use of CLIP for zero-shot segmentation. They leveraged the CLIP model to classify individual regions following a comprehensive proposal generator and then integrate the resulting predictions. Although effective, this design requires two image encoding processes, resulting in expensive computational costs. In order to simplify the pipeline of the two stages, ZegCLIP (Zhou et al. 2023) proposed a one-stage method that transfers CLIP’s powerful generalization ability from images to pixel-level classification. In this work, we use a one-stage method and achieve outstanding zero-shot segmentation performance through two effective designs.

Method

Problem Definition

We adopt the generalized zero-shot semantic segmentation (GZLSS) method (Xian et al. 2019), which requires to segment both seen classes C^s and unseen classes C^u after only training on a dataset with pixel-annotations of seen part. During training, the model generates per-pixel classification results based on the semantic descriptions of all visible classes. During testing, the model is evaluated on both seen and unseen classes. It is important to note that $C^s \cap C^u = \emptyset$ and that the label of C^u is not available during training.

SPT-SEG

The architecture of SPT-SEG is illustrated in Fig. 2. The basic one-stage methodology comprises four key components: the CLIP encoder that incorporates the text and visual encoders, the relationship descriptor between the cls token and the text embedding, a decoder, and a loss function. Our enhancements focus on two pivotal components: (1) Introducing an innovative **Spectral Prompt Tuning** approach within the visual encoder, aimed at extracting structural insights to bolster CLIP’s adaptability to dense prediction tasks, (2) Integrating a **Spectral Guided Decode Layer** into the decoder, which adeptly captures high and low-frequency features specific to the task.

Spectral Prompt Tuning Prompt tuning is a recently proposed fine-tuning technique that offers a valuable approach to adapt pre-trained transformer models to target

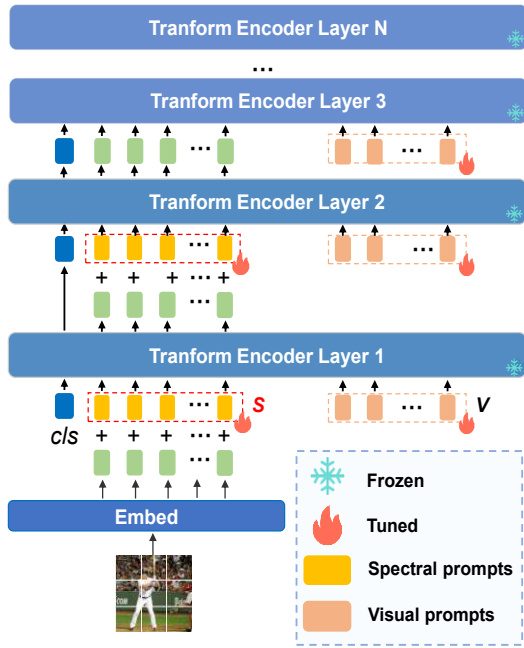


Figure 3: Overview of our proposed Spectral-Prompt Tuning. During training on downstream tasks, only the parameters of prompts and the linear head are updated while the whole Transformer encoder is frozen.

domains (Xing et al. 2022). However, fine-tuning zero-shot segmentation models solely on a limited set of visible classes often leads to overfitting. This occurs because the optimization process focuses solely on visible classes, disregarding knowledge relevant to visual concepts that cannot be obtained from the training set. To address this issue, Visual Prompt Tuning (VPT) (Jia et al. 2022) has emerged as a potential solution. VPT introduces a small number of task-specific learnable parameters in the input space while keeping the backbone frozen during downstream training. While VPT has shown promising results in certain cases, it does not fully leverage the intrinsic properties and structural characteristics of images, which may not be fully manifested in the spatial domain, thereby limiting its effectiveness in handling structure-aware tasks.

To address this limitation, we propose the Spectral Prompt Tuning (SPT) method, as shown in Fig. 3. SPT extends the concept of VPT by incorporating prompt parameters learned from a spectral perspective.

In contrast to VPT’s exclusive reliance on visual prompts for fine-tuning, SPT capitalizes on frequency domain features to offer supplementary understanding of intricate attributes and structural characteristics. The features learned by SPT in the spectrum allow it to better capture and distinguish subtle visual features of different classes, even for those classes that do not have direct examples in the training data. In this way, when the model encounters images of completely new classes, it can extract common information about these classes from the spectrum features, enabling more accurate segmentation. This ability can alleviate the

”partial” or ”ambiguous” segmentation issues that occur in zero-shot scenarios, thus ensuring a more precise capture of unknown classes.

The input embeddings from the l -th layer of the image encoder in the CLIP model are denoted as $\{\mathbf{g}^l, \mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_N^l\}$. Here, \mathbf{g}^l represents the embedding for the [cls] token, and $\mathbf{H}^l = \{\mathbf{h}_1^l, \mathbf{h}_2^l, \dots, \mathbf{h}_N^l\}$ corresponds to the embeddings of image patches. In the context of SPT, the CLIP image encoder’s token sequences are extended with learnable tokens $\mathbf{V}^l = \{\mathbf{v}_1^l, \mathbf{v}_2^l, \dots, \mathbf{v}_M^l\}$ in each layer. Furthermore, learnable spectral prompts $\mathbf{S}^l = \{\mathbf{s}_1^l, \mathbf{s}_2^l, \dots, \mathbf{s}_N^l\}$ are added in the first two layers. These additions enhance the model’s ability to process image features at multiple levels of abstraction.

\mathbf{S}^l is calculated from \mathbf{H}^l and \mathbf{g}^l , and a set of learnable filter parameters \mathbf{w}_f , the process can be expressed as:

$$\mathbf{S}^l = \mathcal{F}^{-1}(\mathcal{F}(\mathbf{H}^l \odot \mathbf{g}^l) \odot \mathbf{w}_f), \quad (1)$$

where \mathcal{F} is the 2D fast fourier transform (FFT) and \mathcal{F}^{-1} is the inverse FFT (IFFT). Then, when $l \leq 2$ the layer processes the input token as:

$$[\mathbf{g}^l, -, \mathbf{H}^l] = \text{Layer}^l([\mathbf{g}^{l-1}, \mathbf{V}^{l-1}, \mathbf{H}^{l-1} + \mathbf{S}^{l-1}]), \quad (2)$$

when $l > 2$ the transform layer processes the input token as:

$$[\mathbf{g}^l, -, \mathbf{H}^l] = \text{Layer}^l([\mathbf{g}^{l-1}, \mathbf{V}^{l-1}, \mathbf{H}^{l-1}]). \quad (3)$$

Spectral Guided Decode Layer In practical semantic segmentation applications, high-quality segmentation results are crucial for the success of the task. Recent work (Patro, Namboodiri, and Agneeswaran 2023) combined spectral layers with multi-head attention in a transformer architecture to capture relevant features in initial layers. LiTv2 (Pan, Cai, and Zhuang 2022) introduced a novel attention mechanism that separately processes high and low-frequency components in attention layers, capturing local and global relationships effectively in classification and segmentation tasks. Drawing inspiration from these insights, we propose an innovative decoding method as shown Fig. 2(b) by introducing frequency domain features during the decoding stage, which significantly enhances the performance of image segmentation. Firstly, the frequency domain-guided decoder can balance the attention on small details and global structure, enabling the model to focus on both local and overall features simultaneously. Secondly, guided by frequency domain features, the decoder can capture object boundaries and textures more accurately, thereby improving the precision of the segmentation results. Most importantly, this decoder exhibits stronger generalization ability on unseen classes, which is crucial for unknown situations in real-world applications. The design comprises the following steps:

(1) The high-frequency branch captures fine-grained local dependencies through local window self-attention., while the low-frequency branch applies average pooling to each window, obtaining low-frequency signals that capture the global dependencies of the input. This high and low-frequency capturing is built on the multi-head self-attention

(MSA) mechanism, which allows for capturing distant relations labeled at different locations in the input sequence $\mathbf{X} \in \mathbb{R}^{N \times D}$. Here, N is the length of the input sequence, and D represents the hidden dimension. To achieve this, we divide the N_h heads in MSA into two groups with a split ratio α . Specifically, αN_h heads are used for the high-frequency branch, and the remaining $(1 - \alpha)N_h$ heads are utilized for the low-frequency branch. The high-frequency branch computes the output by linearly projecting the outputs of the α self-attention heads and then concatenating them as follows: The high-frequency branch performs a simple non-overlapping-window (3×3) partitioning of the inputs X , and then computes the outputs by α -sizing and concatenating them as follows:

$$\text{MSA}_\alpha(\hat{\mathbf{X}}) = \text{Concat}_{h \in [\alpha N_h]} [\text{SA}_h(\hat{\mathbf{X}})], \quad (4)$$

where $\text{SA}_h(\hat{\mathbf{X}})$ denotes the output of the h -th self-attention head, and note that $\hat{\mathbf{X}}$ denotes the input with the non-overlapping window already divided. Meanwhile, the low-frequency branch utilizes average pooling to extract low-frequency signals within each window, and its computation process can be expressed as:

$$\text{MSA}_{1-\alpha}(\hat{\mathbf{X}}) = \text{Concat}_{h \in [(1-\alpha)N_h]} [\text{SA}_h(\text{AvgPool}(\hat{\mathbf{X}}))], \quad (5)$$

Finally, the overall output is obtained by concatenating the outputs from each branch as follows:

$$\mathbf{z} = [\text{MSA}_\alpha(\hat{\mathbf{X}}); \text{MSA}_{1-\alpha}(\hat{\mathbf{X}})], \quad (6)$$

where $[\cdot]$ denotes the concatenation operation.

(2) we emphasize task-relevant tokens and channels through frequency domain feature extraction to select specific characteristics. We perform frequency domain feature extraction on $\mathbf{z} \in \mathbb{R}^{N \times D}$ to identify task-related markers and channels. The output is obtained using the following operation:

$$\hat{\mathbf{z}} = P \cdot \text{sim}(\mathbf{z}, \xi) \cdot \mathbf{z}, \quad (7)$$

where $\xi \in \mathbb{R}^d$ and $P \in \mathbb{R}^{d \times d}$ are task-specific parameters, and $\text{sim}(\cdot, \cdot)$ represents the cosine similarity ranging between $[0, 1]$. The resulting $\hat{\mathbf{z}}$ can be represented as $[\hat{\mathbf{z}}_1, \hat{\mathbf{z}}_2, \dots, \hat{\mathbf{z}}_N] \in \mathbb{R}^{N \times D}$, where $\hat{\mathbf{z}}_j$ denotes the embedding for the j^{th} patch class. The matrix $\mathbf{t} = [\mathbf{t}^1, \mathbf{t}^2, \dots, \mathbf{t}^C] \in \mathbb{R}^{C \times D}$ represents C classes, with d as the feature dimension of the CLIP model. Here, \mathbf{t}^i denotes the representation of the i -th class, and $[\text{cls}]$ corresponds to the global feature represented as $\mathbf{g} \in \mathbb{R}^{N \times D}$. The relationship descriptor can be represented as:

$$\hat{\mathbf{t}} = \phi([\mathbf{t} \cdot \mathbf{g}; \mathbf{t}]), \quad (8)$$

where $\phi(\cdot)$ projects $[\mathbf{t} \cdot \mathbf{g}; \mathbf{t}]$ to the same dimension as $\hat{\mathbf{z}}$.

Semantic masks are calculated using matrix product:

$$\text{Masks} = \hat{\mathbf{t}} \cdot \hat{\mathbf{z}}^T \in \mathbb{R}^{C \times N}, \quad (9)$$

The final segmentation results are obtained by applying the *Argmax* operation along the class dimension of *Masks*.

Loss Function We employ a combination of the focal loss (Lin et al. 2017), and the structural similarity (SSIM) loss (Wang, Simoncelli, and Bovik 2003). The total loss \mathcal{L} is a linear combination of the focal loss and SSIM loss, with coefficients α and β to balance their contributions:

$$\mathcal{L} = \gamma \cdot \mathcal{L}_{\text{focal}} + \sigma \cdot \mathcal{L}_{\text{ssim}}, \quad (10)$$

The coefficients γ, σ are used to control the relative importance of the focal loss and SSIM loss in the overall loss function.

Experiments

Datasets

We conducted extensive experiments on two benchmark datasets to evaluate the effectiveness of our proposed method: PASCAL VOC 2012 (20), COCO-Stuff 164K. Here are the details of each dataset:

1. **PASCAL VOC 2012**: This dataset consists of 10,582 augmented images for training and 1,449 for validation. We focus on 15 seen classes, ignoring the "background" class, and 5 unseen classes.
2. **COCO-Stuff 164K**: It is a large-scale dataset with 118,287 training images and 5,000 testing images, covering 171 classes. Among them, 156 classes are seen, and 15 classes are unseen.

Evaluation Metrics

As in previous studies, we assess the performance using pixel-wise classification accuracy (*pAcc*) and the mean intersection over union (*mIoU*) for both seen and unseen classes, referred to as *mIoU(S)* and *mIoU(U)*, respectively. Additionally, we calculate the harmonic mean IoU (*hIoU*) between the seen and unseen classes as in Zeg-CLIP (Zhou et al. 2023), which is formulated as:

$$hIoU = \frac{2 * mIoU(S) * mIoU(U)}{mIoU(S) + mIoU(U)}. \quad (11)$$

Implementation Details

Our proposed method is implemented using the MMSegmentation open-source toolbox (Contributors 2020) with PyTorch 1.10.1. All experiments were conducted on two H800 GPUs using the pre-trained CLIP ViT-B/16 model. The batch size was set to 16, and the images were resized to a resolution of 512×512 . We performed a total of 20,000 training iterations on the PASCAL VOC 2012 dataset, and 96,000 iterations on the COCO-Stuff 164K dataset. Based on previous research works (Gu et al. 2020; Xu et al. 2021a; Ding et al. 2022; Zhou, Loy, and Dai 2022), we have set up the unseen classes. The optimizer used was AdamW, and we followed the default training schedule provided by the MM-Seg toolbox. In SPT-SEG, it should be noted that the model learns multiple prompts exclusively from seen classes during training. The optimizer used was AdamW, and we followed the default training schedule provided by the MMSeg toolbox.

Methods	PASCAL VOC 2012				COCO-Stuff 164K			
	pAcc	mIoU(S)	mIoU(U)	hIoU	pAcc	mIoU(S)	mIoU(U)	hIoU
SPNet _C VPR'19	/	78.0	15.6	26.1	/	35.2	8.7	14.0
ZS3 _{NeurIPS} '19	/	77.3	17.7	28.7	/	34.7	9.5	15.0
CaGNet _A CMMM'20	80.7	78.4	26.6	39.7	56.6	33.5	12.2	18.2
SIGN _I CCV'21	/	75.4	28.9	41.7	/	32.3	15.5	20.9
Joint _I CCV'21	/	77.7	32.5	45.9	/	/	/	/
ZegFormer _C VPR'22	/	86.4	63.6	73.3	/	36.6	33.2	34.8
zsseg _a Xiv'21	90.0	83.5	72.5	77.5	60.3	39.3	36.3	37.8
ZegCLIP _C VPR'23	94.6	91.9	77.8	84.3	62.0	40.2	41.4	40.8
SPT-SEG (Ours)	96.7	92.9	87.4	90.1	62.9	40.6	43.8	42.1
	(+2.1)	(+1.0)	(+9.6)	(+5.8)	(+0.9)	(+0.4)	(+2.4)	(+1.3)
ZegCLIP* _C VPR'23	96.3	92.4	90.9	91.6	69.9	40.7	63.2	49.6
SPT-SEG* (Ours)	97.6	93.6	92.9	93.2	72.5	41.6	66.0	51.0

Table 1: Comparison with state-of-the-art methods on the PASCAL VOC 2012 and COCO-Stuff 164K datasets. The asterisk (*) denotes training involving all classes. The best results are highlighted in bold.

Comparison with State-of-the-Art Methods

To showcase the effectiveness of our method, we present the evaluation results in comparison with previous state-of-the-art approaches, as shown in Tab. 1. Additionally, we include the results of fully supervised learning as an upper bound to demonstrate the performance gap between fully supervised segmentation and zero-shot segmentation on unseen classes. We provide qualitative results on the COCO-Stuff 164K dataset, depicted in Fig. 4. Our proposed method exhibits significant performance improvements, particularly for unseen classes, surpassing previous approaches, as depicted in Tab. 1. This highlights the superior generalization capability of our method compared to existing methods. Particularly noteworthy is the significant increase in mIoU for unseen classes in the VOC dataset **9.6%** and for unseen classes in the COCO dataset **2.4%**

Fig. 4 showcases the segmentation outcomes of the ZegCLIP (Zhou et al. 2023) and our proposed SPT-SEG, both on seen and unseen classes. With the integration of our proposed designs, SPT-SEG demonstrates impressive segmentation capabilities on both seen and unseen classes, effectively distinguishing similar unseen classes. For example, our approach effectively segments small target 'sport ball' objects and achieves full recognition of the unseen class 'playing field' (Fig. 4(1)). Furthermore, our method successfully discriminates "plastic" classes from skateboard regions (Fig. 4(2)), and accurately segments "dog" instances bearing resemblance to "horses" (Fig. 4(3)). Overall, SPT-SEG completely segments the unseen classes ("playing field", "plastic") and significantly outperforms other methods in terms of segmentation details. These results confirm the effectiveness of our proposed method in achieving superior segmentation performance, especially for unseen classes.

Ablation Study

Detailed results of applying designs on baseline To demonstrate the effectiveness of our proposed designs, we further report the improvements of applying designs on baseline (ZegCLIP) in Tab. 2. The addition of the SPT sig-

Bas.	SPT	SGD	PASCAL VOC 2012		
			mIoU(S)	mIoU(U)	hIoU
✓			91.9	77.8	84.3
✓	✓		92.6	86.7	89.6
✓		✓	92.0	79.9	85.5
✓	✓	✓	92.9	87.4	90.1

Table 2: Quantitative results on VOC dataset to demonstrate the effectiveness of our proposed two designs. Here ✓ means that this component is applied. Note that our baseline (Bas.) method is ZegCLIP (Zhou et al. 2023). The best results are highlighted in bold.

Depth	PASCAL VOC 2012		
	mIoU(S)	mIoU(U)	hIoU
1-6	92.5	86.4	89.3
6-12	92.1	80.9	86.1
1-12	92.6	86.5	89.4
11-12	92.0	78.3	84.6
1-2	92.9	87.4	90.1

Table 3: Ablation on Spectral Prompt Tuning depth. The 1-st layer refers to the one closest to input. ViT-B has 12 layers in total. The best results are highlighted in bold.

nificantly enhances the model's performance on unseen data. When both SPT and SGD are utilized, the SPT-SEG model exhibits excellent results on the VOC test dataset.

Effect of the depth of SPT Tab. 3 demonstrates the impact of SPT insertion positions and layers on SPT-SEG performance. The performance of SPT is notably superior when inserted in the earlier layers compared to the later ones. However, its overall performance is comparable when applied across all layers as well as with its application limited to the first two layers. This finding indicates the greater significance of early transformer layer spectral prompts over later layers' prompts.

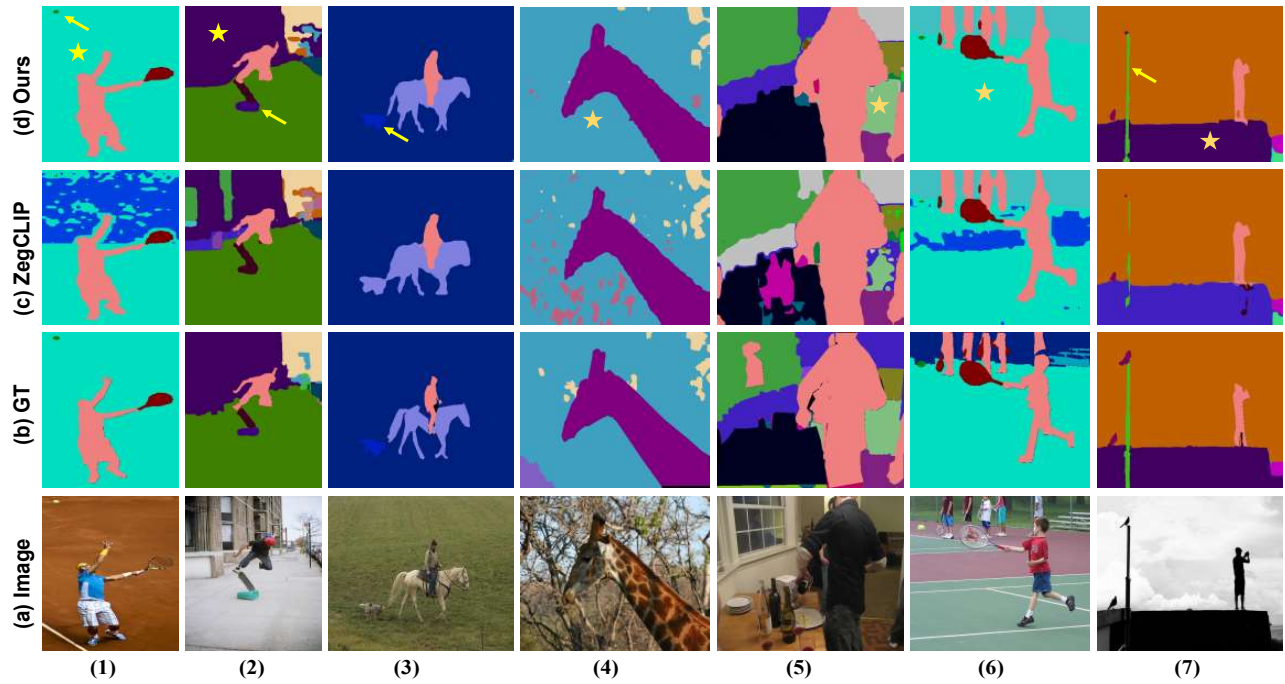


Figure 4: Qualitative results on COCO-Stuff 164K. (a) are the original testing images; (b) are the ground truths of each image. (c) represent the performance of ZegCLIP; (d) are the visualization results of our proposed SPT-SEG. Note that we have highlighted prominent regions using yellow arrows and marked other significant areas with yellow stars for emphasis.

Layers	PASCAL VOC 2012		
	mIoU(S)	mIoU(U)	hIoU
1	91.9	82.8	87.1
3	92.9	87.4	90.1
5	92.2	83.7	87.7

Table 4: Ablation on layers of Spectral Guided Decode Layer. The best results are highlighted in bold.

Effect of Spectral Guided Decode layers To investigate the impact of decoder layers on the performance of SPT-SEG, we conducted an ablation study on decoder layer depth. Tab. 4 demonstrates that within our research settings, the model achieved its optimal performance with 3 decoder layers. At this layer depth, the model exhibited excellent performance both at the pixel-level and class-level. However, when the decoder layers were increased to 5, we observed signs of overfitting, resulting in a decline in performance on the test set. Conversely, employing only 1 decoder layer significantly reduced the model’s performance.

Limitations

Limited by the recognition capability and resolution of CLIP, pixel classification may be prone to errors in complex scenes such as object occlusion and glass reflection (e.g. (Fig. 4(5))). Additionally, the ability to recognize details, such as object edges, also needs improvement. Resolving these limitations and enhancing the robustness of the SPT-

SEG method are important directions for future research.

Conclusion

In this work, we present an efficient one-stage direct zero-shot semantic segmentation method based on the pre-trained vision-language model CLIP. We introduce two innovative designs to transfer image classification capabilities to dense prediction tasks while maintaining a leading edge in zero-shot knowledge. These designs enable us to achieve competitive results on known classes and significantly improve performance on novel classes. To demonstrate the effectiveness of our approach, we comprehensively test its performance on two widely-used benchmark datasets, outperforming the previous state-of-the-art methods. Our research aims to explore the use of pre-trained visual language models for semantic segmentation. By integrating spectral information and enhancing the capability of CLIP, we successfully apply its zero-shot knowledge to downstream tasks, providing a flexible and accurate solution for zero-shot semantic segmentation.

Acknowledgements

This work was supported by Beijing Natural Science Foundation No. JQ23014, and in part by the National Natural Science Foundation of China (Nos. U21A20515, 62271074 and 62276031).

References

- Baek, D.; Oh, Y.; and Ham, B. 2021. Exploiting a joint embedding space for generalized zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 9536–9545.
- Bucher, M.; Vu, T.-H.; Cord, M.; and Pérez, P. 2019. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32.
- Chen, X.; Yuan, Y.; Zeng, G.; and Wang, J. 2021. Semi-supervised semantic segmentation with cross pseudo supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2613–2622.
- Cheng, J.; Nandi, S.; Natarajan, P.; and Abd-Almageed, W. 2021. Sign: Spatial-information incorporated generative network for generalized zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9556–9566.
- Contributors, M. 2020. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>.
- Ding, J.; Xue, N.; Xia, G.-S.; and Dai, D. 2022. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11583–11592.
- Dong, J.; Cong, Y.; Sun, G.; Fang, Z.; and Ding, Z. 2021. Where and How to Transfer: Knowledge Aggregation-Induced Transferability Perception for Unsupervised Domain Adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 1–1.
- Dong, J.; Cong, Y.; Sun, G.; Zhong, B.; and Xu, X. 2020. What Can Be Transferred: Unsupervised Domain Adaptation for Endoscopic Lesions Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4022–4031.
- Esmailpour, S.; Liu, B.; Robertson, E.; and Shu, L. 2022. Zero-shot out-of-distribution detection based on the pre-trained model clip. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, 6568–6576.
- Gao, P.; Geng, S.; Zhang, R.; Ma, T.; Fang, R.; Zhang, Y.; Li, H.; and Qiao, Y. 2021. Clip-adapter: Better vision-language models with feature adapters. *arXiv preprint arXiv:2110.04544*.
- Gu, Z.; Zhou, S.; Niu, L.; Zhao, Z.; and Zhang, L. 2020. Context-aware feature generation for zero-shot semantic segmentation. In *Proceedings of the 28th ACM International Conference on Multimedia*, 1921–1929.
- Hong, Y.; Wu, Q.; Qi, Y.; Rodriguez-Opazo, C.; and Gould, S. 2021. Vln bert: A recurrent vision-and-language bert for navigation. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, 1643–1653.
- Huang, Z.; Zeng, Z.; Huang, Y.; Liu, B.; Fu, D.; and Fu, J. 2021. Seeing out of the box: End-to-end pre-training for vision-language representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12976–12985.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021a. Scaling up visual and vision-language representation learning with noisy text supervision. In *International conference on machine learning*, 4904–4916. PMLR.
- Jia, C.; Yang, Y.; Xia, Y.; Chen, Y.-T.; Parekh, Z.; Pham, H.; Le, Q.; Sung, Y.-H.; Li, Z.; and Duerig, T. 2021b. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, 4904–4916. PMLR.
- Jia, M.; Tang, L.; Chen, B.-C.; Cardie, C.; Belongie, S.; Hariharan, B.; and Lim, S.-N. 2022. Visual prompt tuning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXIII*, 709–727. Springer.
- Jiang, J.; Liu, Z.; and Zheng, N. 2022. Finetuning Pre-trained Vision-Language Models with Correlation Information Bottleneck for Robust Visual Question Answering. *arXiv preprint arXiv:2209.06954*.
- Kamath, A.; Singh, M.; LeCun, Y.; Synnaeve, G.; Misra, I.; and Carion, N. 2021. Mdetr-modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1780–1790.
- Kim, W.; Son, B.; and Kim, I. 2021. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, 5583–5594. PMLR.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, 2980–2988.
- Liu, Z.; Rodriguez-Opazo, C.; Teney, D.; and Gould, S. 2021. Image retrieval on real-life images with pre-trained vision-and-language models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2125–2134.
- Lu, X.; Wang, W.; Danelljan, M.; Zhou, T.; Shen, J.; and Van Gool, L. 2020. Video object segmentation with episodic graph memory networks. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, 661–679. Springer.
- Lu, Y.; Liu, J.; Zhang, Y.; Liu, Y.; and Tian, X. 2022. Prompt distribution learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5206–5215.
- Pakhomov, D.; Hira, S.; Wagle, N.; Green, K. E.; and Navab, N. 2021. Segmentation in style: Unsupervised semantic image segmentation with stylegan and clip. *arXiv preprint arXiv:2107.12518*.
- Pan, Z.; Cai, J.; and Zhuang, B. 2022. Fast Vision Transformers with HiLo Attention. In *NeurIPS*.
- Pastore, G.; Cermelli, F.; Xian, Y.; Mancini, M.; Akata, Z.; and Caputo, B. 2021. A closer look at self-training for zero-label semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2693–2702.

- Patro, B. N.; Namboodiri, V. P.; and Agneeswaran, V. S. 2023. SpectFormer: Frequency and Attention is what you need in a Vision Transformer. *arXiv preprint arXiv:2304.06446*.
- 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*, 8748–8763. PMLR.
- Rao, Y.; Zhao, W.; Chen, G.; Tang, Y.; Zhu, Z.; Huang, G.; Zhou, J.; and Lu, J. 2022. Densclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18082–18091.
- Sandler, M.; Zhmoginov, A.; Vladymyrov, M.; and Jackson, A. 2022. Fine-tuning image transformers using learnable memory. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12155–12164.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; and Zhang, X. 2023a. Automatic polyp segmentation via image-level and surrounding-level context fusion deep neural network. *Engineering Applications of Artificial Intelligence*, 123: 106168.
- Wang, C.; Xu, R.; Xu, S.; Meng, W.; and Zhang, X. 2023b. Treating Pseudo-labels Generation as Image Matting for Weakly Supervised Semantic Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 755–765.
- Wang, W.; Bao, H.; Dong, L.; Bjorck, J.; Peng, Z.; Liu, Q.; Aggarwal, K.; Mohammed, O. K.; Singhal, S.; Som, S.; et al. 2023c. Image as a Foreign Language: BEiT Pretraining for Vision and Vision-Language Tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19175–19186.
- Wang, Z.; Lu, Y.; Li, Q.; Tao, X.; Guo, Y.; Gong, M.; and Liu, T. 2022. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11686–11695.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thirty-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Xian, Y.; Choudhury, S.; He, Y.; Schiele, B.; and Akata, Z. 2019. Semantic projection network for zero-and few-label semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8256–8265.
- Xie, G.-S.; Liu, J.; Xiong, H.; and Shao, L. 2021. Scale-aware graph neural network for few-shot semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5475–5484.
- Xing, Y.; Wu, Q.; Cheng, D.; Zhang, S.; Liang, G.; and Zhang, Y. 2022. Class-aware visual prompt tuning for vision-language pre-trained model. *arXiv preprint arXiv:2208.08340*.
- Xu, M.; Zhang, Z.; Wei, F.; Lin, Y.; Cao, Y.; Hu, H.; and Bai, X. 2021a. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*.
- Xu, R.; Wang, C.; Sun, J.; Xu, S.; Meng, W.; and Zhang, X. 2023a. Self Correspondence Distillation For End-to-End Weakly-Supervised Semantic Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.
- Xu, R.; Wang, C.; Xu, S.; Meng, W.; and Zhang, X. 2021b. DC-net: Dual context network for 2D medical image segmentation. In *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part I 24*, 503–513. Springer.
- Xu, R.; Wang, C.; Xu, S.; Meng, W.; and Zhang, X. 2023b. Dual-stream Representation Fusion Learning for accurate medical image segmentation. *Engineering Applications of Artificial Intelligence*, 123: 106402.
- Xu, R.; Wang, C.; Xu, S.; Meng, W.; and Zhang, X. 2023c. Wave-Like Class Activation Map With Representation Fusion for Weakly-Supervised Semantic Segmentation. *IEEE Transactions on Multimedia*.
- Xu, R.; Wang, C.; Zhang, J.; Xu, S.; Meng, W.; and Zhang, X. 2023d. Rssformer: Foreground saliency enhancement for remote sensing land-cover segmentation. *IEEE Transactions on Image Processing*, 32: 1052–1064.
- Yu, J.; Wang, Z.; Vasudevan, V.; Yeung, L.; Seyedhosseini, M.; and Wu, Y. 2022. Coca: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*.
- Zhang, R.; Zhang, W.; Fang, R.; Gao, P.; Li, K.; Dai, J.; Qiao, Y.; and Li, H. 2022. Tip-adapter: Training-free adaptation of clip for few-shot classification. In *European Conference on Computer Vision*, 493–510. Springer.
- Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract free dense labels from clip. In *European Conference on Computer Vision*, 696–712. Springer.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022a. Conditional prompt learning for vision-language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16816–16825.
- Zhou, K.; Yang, J.; Loy, C. C.; and Liu, Z. 2022b. Learning to prompt for vision-language models. *International Journal of Computer Vision*, 130(9): 2337–2348.
- Zhou, Z.; Lei, Y.; Zhang, B.; Liu, L.; and Liu, Y. 2023. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11175–11185.