

FREESEG: FREE MASK FROM INTERPRETABLE CONTRASTIVE LANGUAGE-IMAGE PRETRAINING FOR SEMANTIC SEGMENTATION

Anonymous authors

Paper under double-blind review

ABSTRACT

Fully supervised semantic segmentation learns from dense masks, which requires heavy annotation cost for closed set. In this paper, we use natural language as supervision without any pixel-level annotation for open world segmentation. We call the proposed framework as FreeSeg, where the mask is freely available from raw feature map of pretraining model. Compared with zero-shot or openset segmentation, FreeSeg doesn't require any annotated masks, and it widely predicts categories beyond class-agnostic unsupervised segmentation. Specifically, FreeSeg obtains free mask from Image-Text Similarity Map (ITSM) of Interpretable Contrastive Language-Image Pretraining (ICLIP). And our core improvements are the smoothed min pooling for dense ICLIP, with the partial label and pixel strategies for segmentation. Furthermore, FreeSeg is very straight forward without complex design like grouping, clustering or retrieval. Besides the simplicity, the performances of FreeSeg surpass previous state-of-the-art at large margins, e.g. 13.4% higher at mIoU on VOC dataset in the same settings.

1 INTRODUCTION

Semantic segmentation is an important and widely used task in computer vision, which requires dense and heavy pixel-level annotations for close set. To achieve data-efficient segmentation, prior works reduce the cost of annotation by applying weak supervisions like scribble Lin et al. (2016), box Dai et al. (2015), point Bearman et al. (2016) and image-level label Pinheiro & Collobert (2015); Ahn & Kwak (2018). In another way, semi-supervised segmentation Wei et al. (2018) reduces the cost via using partial mask annotations. Besides the cost, zero-shot segmentation Bucher et al. (2019), open-set segmentation da Silva et al. (2020) and unsupervised segmentation Van Gansbeke et al. (2021) broadens the categories to open world. Among them, unsupervised segmentation doesn't require mask annotations as the others, while it provides binary mask only.

To achieve a type of segmentation based on natural language for open world segmentation, we focus on open-vocabulary segmentation based on Contrastive Language-Image Pretraining (CLIP) Radford et al. (2021). Since the image-text pairs are readily available, and show powerful ability on zero-shot recognition. Some works Xu et al. (2021); Zhou et al. (2021); Li et al. (2022a) apply this advantage to zero-shot segmentation. However, they still rely on pixel-level annotations to predict mask for unseen classes. To get rid of costly pixel-level annotations, some works generate the segments by retrieval Shin et al. (2022), grouping Xu et al. (2022) or unsupervised segmentation Zabari & Hoshen (2021) with gradient-based interpretability method Chefer et al. (2021b). In this paper, we explore an easier, directer and better way to predict the segments from the raw feature map of pretraining model.

We wonder whether CLIP can provide freely available segments from its predicted tokens (last feature map), like Class Attention Map (CAM) Zhou et al. (2016) in weakly supervised semantic segmentation. In the Interpretable Contrastive Language-Image Pretraining (ICLIP) Li et al. (2022c), the answer is yes. This paper successfully locates discriminative regions, via the Image-Text Similarity Map (ITSM), after solving a problem called semantic shift. However, there is still a gap from heatmap to segmentation mask. Firstly, the segmentation task requires higher output resolution for dense predictions, and we find dense ICLIP focuses more on the local and most discriminative re-




Figure 1: Visualization of Image-Text Similarity Map (ITSM) Li et al. (2022c). Dense ITSM is prone to locate at partial regions than original ITSM, and our method solves this problem via the proposed smoothed min pooling. Note, the output sides of Dense ITSM and Ours are 56, and that of ITSM is 14. Besides, regions close to red are the target and background is colored in blue, where foreground pixels (score >0.5) are surrounded by dark lines.

regions, thus performing worse than original ICLIP as Fig. 1. Secondly, how to project 3D ITSM to 2D mask is a key problem, and inappropriate approach leads to unacceptable results.

To solve these problems, and fulfil the process from ITSM to segments, we propose the framework FreeSeg. Specifically, we scale up ICLIP for dense predictions, and make up the consequent partial activation problem via the proposed smoothed min pooling. After the generation of high quality dense ITSM, we project it to segments via a partial category scheme, and keep partial credible pixels of segments as pseudo mask for the refinement stage. Because of the partial labels and pixels, we call this part as partially supervised semantic segmentation. It’s also a straight forward and powerful framework for vocabulary based semantic segmentation. Experiments on PASCAL VOC 2014 Everingham et al. (2010) and MS COCO 2017 Lin et al. (2014) show the nontrivial performances compared with previous state-of-the-art, where the improvements are 13.4% and 4.0% at mIoU, respectively. Our main contributions are summarized below:

- We propose a straight forward and effective framework, FreeSeg, for open-vocabulary segmentation from the raw feature map of pretraining model.
- We extend ICLIP to dense predictions for semantic segmentation, then solve the consequent partial activation problem by the proposed smoothed min pooling, and fulfil the process from ITSM to segments with partially supervised semantic segmentation.
- FreeSeg achieves nontrivial improvements compared with previous state-of-the-art, and it’s very simple to use, without additional algorithms or proprietary encoder.

2 RELATED WORK

Data-efficient semantic segmentation via text supervision. Most data-efficient semantic segmentation methods are single modality task, via manual structured supervision, such as weak supervisions Lin et al. (2016); Dai et al. (2015); Bearman et al. (2016); Pinheiro & Collobert (2015) or semi mask supervision Wei et al. (2018). Different from previous methods, text supervision is a burgeoning approach for data-efficient segmentation, and provides new features. Firstly, images on the internet are often tagged with text descriptions, and it’s easy to obtain massive data without manual annotation. Secondly, text based pretraining model Radford et al. (2021) show powerful zero-shot ability, making zero-shot and open-set segmentation realizable. Some works Xu et al. (2021); Zhou et al. (2021); Li et al. (2022a) use CLIP Radford et al. (2021) to identify unseen classes with annotated mask from seen classes. However, the pixel-level mask annotation is still inefficient. Then,

GroupViT Xu et al. (2022) generates segments by merging the image tokens with a special backbone. And Shin et al. (2022) uses extra retrieval algorithm across multiple images to synthesis the mask. Besides, Zabari & Hoshen (2021) obtains segments via interpretability Chefer et al. (2021b) based on gradient, but owing to the limited localization quality, additional unsupervised segmentation algorithm is required. Compare with above methods, we obtain qualitative masks from the raw feature map of pretraining model, which is straight forward, effective and almost freely available.

Interpretable Contrastive Language-Image Pretraining (ICLIP). To obtain freely available segments, we focus on ICLIP Li et al. (2022c), a new work for the interpretability of CLIP. This work allows the pretraining model to locate the discriminative regions, which suits the segmentation task well. It interprets the predictions of image encoder by the Image-Text Similarity Map (ITSM $\mathbf{M} \in \mathbb{R}^{H,W,N_t}$) as Eq. 1, where N_t is the number of texts, and $\mathbf{X}_{1,:} \times \mathbf{Y}^\top$ indicates inner production between image features $\mathbf{X}_{1,:}$ (class token $\mathbf{X}_{:1}$ is excluded) and transposed text features \mathbf{Y}^\top , with operations of reshape, resize and min-max normalization.

$$\mathbf{M} = \text{Norm}(\text{Resize}(\text{Reshape}(\mathbf{X}_{1,:} \times \mathbf{Y}^\top))) \quad (1)$$

But ITSM cannot be directly applied to the original CLIP, because of one observed phenomenon: semantic shift between foregrounds and backgrounds. This problem leads to erroneous visualization results, and ICLIP solves it by the Masked Max Pooling (MMP) as Fig. 2, where \mathbf{F}_c is the pooled class feature, and $\hat{f}_s(x)_{1,:}$ indicate the outputs except the first class token before pooling. Besides, $\odot \mathbf{A}$ indicates that the outputs are element-wisely multiplied to the mean attention map \mathbf{A} from the last transformer layer. Note, *max* means max pooling over the token dimension, and $\mathbf{X}_{1,:}$ in Eq. 1 comes from $\hat{f}_s(x)_{1,:} \odot \mathbf{A}$ without pooling after a linear projection.

$$\mathbf{F}_c = \text{Mmp}(\hat{f}_s(x)_{1,:}) = \text{max}(\hat{f}_s(x)_{1,:} \odot \mathbf{A}) \quad (2)$$

Compared with original CLIP, the network architecture of ICLIP only adds a pair of linear projection layers with MMP. These extra projections are designed for independent representations of interpretability task, which is much simpler than GroupViT Xu et al. (2022) in architecture. In some degree, the ITSM is a by-product of the interpretable pretraining model, and the segments are almost freely available from ICLIP. However, there is still a gap from ITSM to segmentation results, and the output size of ICLIP is too small for semantic segmentation.

3 METHODOLOGY

3.1 OVERALL FRAMEWORK

From the raw feature map interpretable pretraining model, ICLIP Li et al. (2022c) to segments, there are two necessary steps: (1) Scaling up the output size to generate dense and detailed Image-Text Similarity Map (ITSM) for segmentation. (2) Project the ITSM to segments and refine the results. We later write the details of these steps in the next two subsections. Here, we generally introduce the overall framework of FreeSeg. For the first step, we find enlarged dense ICLIP performs bad, and focus on partial regions as Fig. 1. Then we solve this problem by the proposed Smoothed Min Pooling (SMP) to mitigate the most discriminative representations, thus emphasize the rest regions. For the second step, we keep confidential labels to generate the segments, and synthesis pseudo masks using partial credible pixels as the partially supervised semantic segmentation.

The overall framework of FreeSeg is demonstrated in Fig 2. The left part is the training phase, where there are two pairs of linear projections $\phi, \hat{\phi}$ for classification and interpretability task, respectively. Our goal of this phase is to learn a pair of linear projections $\hat{\phi}$ to generate dense ITSM out of partial activation. We achieve it by the proposed smoothed min pooling. Firstly, the image tokens $\mathbf{X}_{1,:}$ (without first class token $\mathbf{X}_{:1}$) are smoothed by an average pooling layer with sliding windows. Then we pick minimum signals as class features \mathbf{F}_c to update $\hat{\phi}$ with contrastive loss.

The right part is the test phase, with partially supervised semantic segmentator. Specifically, the zero-shot classification is evaluated by metric mAP from the original class token $\mathbf{X}_{:1}$ with its linear projections ϕ . And the dense ITSM \mathbf{M} is generated with $\hat{\phi}$ and image tokens $\mathbf{X}_{1,:}$ before smoothing

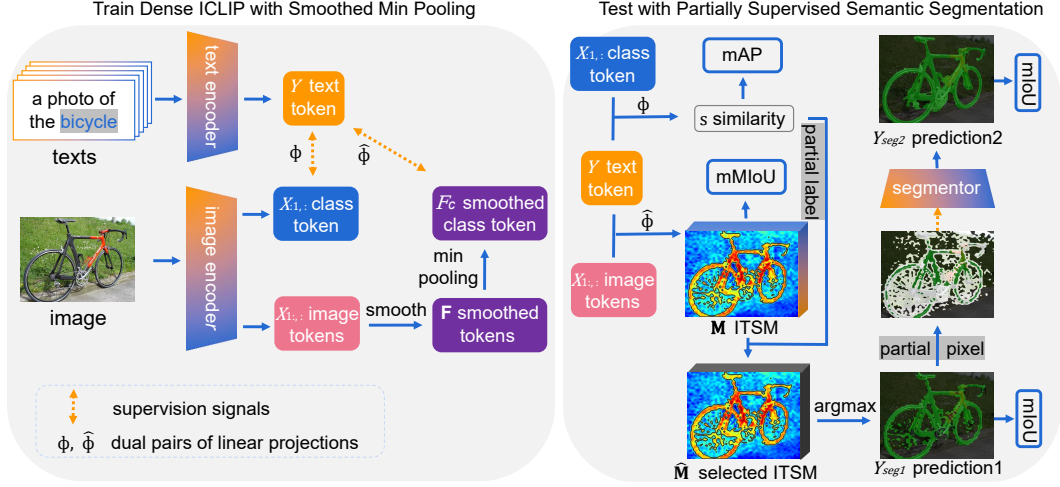


Figure 2: Overall framework of FreeSeg. We apply Smoothed Min Pooling (SMP) for Dense ICLIP, to solve the partial activation problem in Fig. 1. SMP learns a pair of linear projections $\hat{\phi}$ which is less sensitive to partial regions, then generates dense Image-Text Similarity Map (ITSM) with it in the test phase. We deploy the partial label and pixel strategies to refine the results with partial and credible supervision signals. Note that the blue boxes indicate the evaluation metrics.

(SMP is designed for training only). Then we keep confidential labels guided by similarity s , and apply argmax operation to get the segments Y_{seg1} of first stage. Finally, the segmentation model supervised by partial credible supervision signals refines the results, and predict the Y_{seg2} . Note, segmentation results are evaluated by mIoU, and mMIoU measures the interpretability.

3.2 SMOOTHED MIN POOLING FOR DENSE IMAGE-TEXT SIMILARITY MAP

As shown in Fig. 1, the ITSM of the original ICLIP is insufficient to generate fine and detailed segments. The first task is to obtain a dense ICLIP in high quality for the segmentation. We increase the input resolution and reduce the patch size of the image encoder, and then get the dense ICLIP, where the size of ITSM is 56×56 . However, the dense ICLIP is prone to highlight the most discriminative regions, causing the problem of partial activation. This problem leads to serious performance degradation. So, how to ensure the ICLIP works well for dense ITSM is our target. Our motivation is to mitigate those highly similar image-text pixel pairs, and thus reduce the contract between most discriminative region and other parts. To achieve this goal, we proposed the Smoothed Min Pooling (SMP), a novel pooling method designed for dense ICLIP, which is consisted of smoothing operation and a rarely used global minimum pooling.

For the smoothing operation, we realize the smoothing operation via the average pooling with sliding window. Before it, we firstly reshape the feature matrix of image tokens $F_i \in \mathbb{R}^{N_i, C}$ to feature map $\hat{F} \in \mathbb{R}^{H, W, C}$ before smoothing, as Eq. 3, where $\hat{f}_s(x)_{1..:}$ is the masked features as Eq. 2. Here, C is the channel size (embedding width), and H, W are the height and width.

$$F_i = \hat{f}_s(x)_{1..:} \odot A, \hat{F} = \text{reshape}(F_i) \quad (3)$$

Then we apply average pooling to smooth the masked and reshaped feature map \hat{F} to get the smoothed feature $F \in \mathbb{R}^{H, W, C}$ as Eq. 4, where k is the size of sliding window, $///$ indicates exact division, and $\hat{F}_{i,j,c} = 0$ is the padding value.

$$F_{h,w,c} = \frac{\sum_{i=h-k//2}^k \sum_{j=w-k//2}^k \hat{F}_{i,j,c}}{k^2}, \forall h, w, c \quad (4)$$

$$s.t. \hat{F}_{i,j,c} = 0, \text{ if } i \notin H \text{ or } j \notin W$$

To further suppress peak activations, we choose a seldom used pooling operation: minimum pooling. Thanks to the existent of extra projections of ICLIP and the proposed smoothing operation, min pooling is able to learn the non-partial representations, without interfere to the classification branch. Then, we write the smoothed min pooling Smp in Eq. 5, where F_c is the pooled feature, and min operation is carried out over H and W dimensions, besides avg indicates Eq. 4.

$$\mathbf{F}_c = Smp(\hat{\mathbf{F}}) = min(avg(\hat{\mathbf{F}})) = min(\mathbf{F}) \quad (5)$$

Note that SMP is *only applied in training phase*, to learn the added pair of linear projections. This pooling make the projections less sensitive to the most discriminative parts, thus avoiding the partial activation problem. During inference phase, SMP including smoothing operation is not applied, to make sure clear and fine ITSM. So, compare to image smoothing operations, our results are not blurry and performs well as Fig. 1. Besides, we carry out further analysis towards the hardly used min pooling in Appendix A.

3.3 PARTIALLY SUPERVISED SEMANTIC SEGMENTATION

After obtaining high quality dense ITSM, we need to convert it to segments. For fully supervised semantic segmentation, this step is realized by argmax operation. However, it leads to unacceptable results for ITSM. As shown in Tab. 1, argmax performs badly, even the quality of dense ITSM is high. We fulfil the process from dense ITSM to segments in this part, and achieve a result 50.27%, which is near to the upper bound 56.62% (mMIoU).

Table 1: Argmax performs badly when convert dense ITSM to segments on VOC dataset. Note, mMIoU measures the interpretability of ITSM and mIoU measures the segmentation results.

Operation	mMIoU (%) \uparrow	mIoU (%) \uparrow
argmax	56.62	5.70
Ours	56.62	50.27

Note, the mean Match Inter of Union (mMIoU) is proposed in ICLIP Li et al. (2022c), to evaluate ITSM \mathbf{M} without interference of classification and threshold engineering as Eq. 6. In the equation, c, n are numbers of class (without background class) and samples, respectively, and $\mathbf{G} \in \mathbb{R}^{H,W}$ is the ground truth, t indicates the best foreground threshold, $i \in \mathbf{G}$ means the usage of image-level labels.

$$mMIoU = \sum_{i=1}^c \sum_{j=1}^n \frac{(\mathbf{M}_{:,i}^j > t) \cap (\mathbf{G}^j = i)}{(\mathbf{M}_{:,i}^j > t) \cup (\mathbf{G}^j = i)}, s.t. i \in \mathbf{G} \quad (6)$$

We achieve the result "Ours" in Tab. 1 by choosing partial credible label set \mathbb{P} before argmax operation. The motivation is that keeping partial credible labels excludes the interference of irrelevant categories. Specifically, pixel-level classification of image tokens is worse than image-level results, and min-max normalization in ITSM emphasizes this error. We call this method as Partial Label Strategy (PLS). It firstly selects a partial credible label set \mathbb{P} from all text labels \mathbb{A} . Then values of ITSM on classes out of \mathbb{P} are set to 0, before argmax operation as Eq. 7. Note, background matrix \mathbf{B} is concatenated to selected ITSM $\hat{\mathbf{M}}$, and its value is set to background threshold t .

$$\mathbf{Y}_{seg1} = Pls(\mathbf{M}) = argmax(concat(\mathbf{B}, \hat{\mathbf{M}})), \quad (7)$$

where $\hat{\mathbf{M}}_{:,c} = 0$, if $c \notin \mathbb{P}$

In above equation, \mathbf{Y}_{seg1} is the segmentation results of the first stage from ICLIP. Then, we introduce the details of how the \mathbb{P} is selected. We need two thresholds for the zero-shot scores $\mathbf{s} = \mathbf{X}_{:,1:} \times \mathbf{Y}^\top$. As Eq. 8, we add any text label c to \mathbb{P} , whose score is larger than the strict threshold t_s . If there is no label selected of one image, we pick the top 1 label if its score beyond the loose threshold t_l .

$$c \in \mathbb{P}, \text{ if } s_c > t_s$$

$$argmax(s) \in \mathbb{P}, \text{ if } max(s) > t_l \ \& \ \mathbb{P} = \emptyset \quad (8)$$

Note that, the scores s is not uniform, we record the all the predicted scores of a dataset, and apply min-max normalization to the class dimension to calibrate the score to $[0, 1]$ for label selection.

Then we get the segmentation results Y_{seg1} via the Partial Label Strategy (PLS) as Eq. 7. Followed by pseudo mask based weakly supervised segmentation methods Li et al. (2021; 2022b), we refine the results via extra segmentation model. Besides the PLS, we also apply a partial pixel strategy for the second stage. This method is simple while effective. To be specific, we add a threshold delta Δt for background threshold t , to note the binary weight mask \mathbf{W} of credible backgrounds (right part in Eq. 9) and foregrounds based on the max value of partial labeled ITSM $\hat{\mathbf{M}}$.

$$\mathbf{W} = (\max(\hat{\mathbf{M}}) > (t + \Delta t)) \cup (\max(\hat{\mathbf{M}}) < (t - \Delta t)) \quad (9)$$

Then we keep the corresponding loss scatters in stage 2 as Eq. 10, Where \mathbf{L} is the loss matrix for each pixel, and the segmentation loss is \mathcal{L} with input x , supervision Y_{seg1} from zero-shot segmentation. Finally, we get the refined segmentation predictions Y_{seg2} from this model.

$$\mathcal{L}(x, Y_{seg1}) = \text{mean}(\mathbf{L} \cdot \mathbf{W}) \quad (10)$$

Since this process from ITSM to segments with refinement use partial label and pixels, we call the method as partially supervised semantic segmentation.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP.

Dataset. Followed by ICLIP Li et al. (2022c), we use Google Conceptual Captions 3 million (GCC3M) Sharma et al. (2018) as the training dataset, where there are only about 3 million image-text pairs. Compared with CLIP Radford et al. (2021), whose train size is 400 million, we need much fewer data. For test phase, we evaluate the segmentation performances on PASCAL VOC 2012 validation dataset Everingham et al. (2010) and MS COCO 2017 validation set Lin et al. (2014), where the category numbers are 21 and 81, respectively, with one background class. These datasets are widely used in data-efficient segmentation, also applied to interpretability Li et al. (2022c) task.

Framework. The image encoder of FreeSeg is ViT-B Dosovitskiy et al. (2020) (patch size 8), and we use the self-supervised weights of DINO Caron et al. (2021) to accelerate convergence with fewer training data. And the text encoder is the same as CLIP Radford et al. (2021). For the smoothed min pooling, we realize it via average pooling, where the kernel size is 19 with padding 9 and stride 1. In the test phase, there are two segmentation outputs, and the second result is predicted from the PSPNet Zhao et al. (2017) with backbone ResNet-101 He et al. (2016), supervised by the first predictions with partial label and pixel.

Training. During training, the image-encoder is locked without gradient for fast convergence as LiT Zhai et al. (2022). And rest weights are updated by AdamW Loshchilov & Hutter (2017), at learning rate $6.25e-5$, total batch size 512, weight decay 0.05 for 10 epochs without warm up. Other training settings including text encoder, augmentation, scheduler, etc. are followed by Xu et al. (2022); Touvron et al. (2021). To get dense ITSM, we set the training resolution of image encoder to 320×320 , and the test input is 448×448 . For the partially supervised semantic segmentation, we use PSPNet implemented by mmsegmentation MMLab (2020) to refine the results. Its input size is 512×512 , and train 20k iterations at batch size 16 for VOC and 40k interactions at batch size 32 for COCO, where the learn rate is 0.01 with SGD optimizer and poly learn rate scheduler.

Evaluation. As shown in Fig. 2, there are four outputs for evaluation. The zero-shot multi-label classification is evaluated by mean Precision-Recall (mAP) at test resolution 320 as the training size. The test size is also 320 for ISTSM, evaluated by the mean Match Inter of Union (mMIoU) as Eq. 6 for interpretability task. The segmentation results are measured by mean Inter of Union (mIoU). For the result Y_{seg1} from ITSM, its test resolution is 448×448 and output size is 56×56 , then interpolated to original size to count mIoU without post-processing. Besides, Y_{seg2} is the refined result from partially supervised semantic segmentation with dense CRF Krähenbühl & Koltun (2011) as the post-processing. For the thresholds of partial label strategy in Eq. 8, the strict label threshold t_s is

0.8 and the loose label threshold t_l is 0.6. For partial pixel strategy, the value of background B in Eq. 7 is set to 0.5, and the range Δt of ignored pixel in Eq.10 is 0.2.

4.2 ABLATION STUDY

Necessity of Smoothed Min Pooling (SMP). The proposed SMP is a core module to enable dense ITSM. To investigate the necessity of it, we compared it with different pooling methods. As shown in Tab. 2, dense ICLIP Li et al. (2022c) improves the classification performance slightly, but performances of interpretability (mMIoU) and segmentation (mIoU) drop obviously, where this problem is depicted in Fig. 1. To solve this problem, we add a smoothing operation to ICLIP as smoothed max pooling. This pooling mitigates the peak signals of image tokens, thus learn smoothed representations to highlight the under discriminative regions. To further mitigate most discriminative features, we replace the global max pooling to min pooling, which is seldom used before, then get the SMP. Notably, the proposed SMP achieves the best mMIoU and mIoU, beyond MMP by 12.67% and 13.03%, respectively, with comparable mAP. These results indicate SMP is vital important to generate qualitative dense predictions, and we visualize its improvements in Appendix C.

Table 2: Necessity of Smoothed Min Pooling (SMP). The output side of Dense ICLIP Li et al. (2022c) is 4 times larger than ICLIP, but it performs worse with original Masked Max Pooling. Smoothed Max Pooling indicates the min pooling of SMP is replaced by max pooling. Note, mMIoU is the metric for interpretability, and mIoU measures the segmentation result (stage 1) on VOC 2012 validation dataset. The best result is marked in bold and the second is noted with the underline.

Experiments	Pooling Method	mAP (%) \uparrow	mMIoU (%) \uparrow	mIoU (%) \uparrow
ICLIP	Masked Max Pooling	78.06	50.31	43.22
Dense ICLIP	Masked Max Pooling	79.71	43.55	37.24
Dense ICLIP	Smoothed Max Pooling	79.14	<u>53.72</u>	<u>48.02</u>
Dense ICLIP	Smoothed Min Pooling	<u>79.22</u>	56.22	50.27

Kernel size of smoothing. For the Smoothed Min Pooling (SMP), the most important hyperparameter is the kernel size, and we investigate it as Tab. 3, the mAP ranks first as kernel size 7, but its mMIoU is lower than 11, 13, 19. Since SMP is designed for dense ICLIP, where the major metric is mMIoU, we select kernel size 19 based this metric. Because its mMIoU is almost the same to the best, meanwhile it performs well on mAP. Besides, we can see that SMP performs better than masked max pooling in Tab. 2 at any kernel size on mMIoU.

Table 3: Selection of smoothing kernel size.

Kernel Size	mAP (%)	mMIoU (%)
3	78.99	50.41
5	78.19	53.91
7	79.76	55.62
11	78.95	56.06
13	78.67	56.68
19	<u>79.22</u>	<u>56.62</u>
39	<u>78.17</u>	53.46

Table 4: Effectiveness of partially supervised semantic segmentation. Argmax is the baseline from ITSM to Y_{seg1} (stage 1), and partial label strategy significantly improves the result on VOC 2012 validation set. Besides, partial pixel strategy keeps credible supervisions and helps to stage 2.

Experiments	Stage	mIoU (%) \uparrow
Directly Argmax	1	5.70
Partial Label Strategy	1	50.27
Partial Label Strategy	2	62.18
Partial Label Strategy + Partial Pixel Strategy	2	65.68

Effectiveness of partially supervised semantic segmentation. Here, we analyze the effectiveness of the partially supervised semantic segmentation part of FreeSeg. After dense ITSM, the simplest way to get a segmentation map is argmax. However, as seen in Table 4, directly applying the argmax operation produces undesirable results. In comparison, our partial label strategy can produce a significant improvement, increasing from 5.7%

to 50.27%. Furthermore, we use this prediction as pseudo mask in stage 2, and we can achieve better results. In stage 2, we employ the partial pixel strategy, which can preserve reliable supervision signals and boost performance from 62.18 to 65.68.

Accumulative gains of FreeSeg. Table 5 shows the accumulative gains from each component of FreeSeg on the VOC 2012 dataset. The baseline is directly using Argmax on the dense ITSM. Based on this, we proposed two new techniques, smoothed min pooling and partially supervised semantic segmentation, to improve the performance. Without these two modules the baselines are merely 37.24% and 5.70%, respectively. The mIoU reaches to 62.18 after stage2 with pseudo mask from partial label, and partial pixel pushes the result to 65.68%.

Table 5: Accumulative gains of FreeSeg on VOC 2012 validation set at metric mIoU (%). Note, SMP indicates Smoothed Max Pooling, and stage 2 means the usage of segmentation network.

Dense ITSM	SMP	Partial Label	Stage 2	Partial Pixel	mIoU (%)
✓					5.70
✓		✓			37.24
		✓			43.22
✓	✓	✓			50.27
✓	✓	✓	✓		62.18
✓	✓	✓	✓	✓	65.68

4.3 RESULTS

Compared with methods in the same settings. We compare our FreeSeg with existing methods in the same settings as Tab. 6 on VOC 2012 dataset. The listed methods are supervised by text supervisions only, and we don’t compare zero-shot segmentation methods Xu et al. (2021); Zhou et al. (2021); Li et al. (2022a) based on manual mask of seen classes, or methods requiring additional algorithms for mask. Note that, the CLIP based methods use the output feature map based on non-parameter grouping methods like spectral clustering Shi & Malik (2000), mean-shift Comaniciu & Meer (2002) and K-means, which are implemented by GroupViT Xu et al. (2022).

Table 6: Comparison with previous state-of-the-arts on VOC 2012 validation set. Note, GroupViT trained with GCC 3M Sharma et al. (2018) is reproduced via its official code for stage 1, and stage 2 uses the same segmentation code as ours.

Method	Data	Stage 1 (mIoU%)	Stage 2 (mIoU%)
CLIP + Spectral Clustering Shi & Malik (2000)	400M	19.7	-
CLIP + Mean-shift Comaniciu & Meer (2002)	400M	20.7	-
CLIP + K-means	400M	25.0	-
GroupViT Xu et al. (2022)	28M	<u>37.2</u>	<u>52.3</u>
GroupViT Xu et al. (2022)	3M	18.4	21.4
FreeSeg	3M	50.3 _{+13.1}	65.7 _{+13.4}

It is clear that our FreeSeg outperforms the prior state-of-the-art GroupViT Xu et al. (2022) by 13.1% and 13.4%, respectively, on stage 1 and stage 2 at mIoU. Note that the prior SoTA results were reported by GroupViT Xu et al. (2022) trained on about 28 millions samples (GCC12M Changpinyo et al. (2021) + YFCC14M Thomee et al. (2016)). Because of limited computation resources, we train with 3 million image-text pairs of GCC3M Sharma et al. (2018), and reproduce GroupViT with the same dataset. Compared with this setting, FreeSeg is 44.3% higher than GroupViT at second stage.

Besides VOC, our FreeSeg performs better than previous SoTA, GroupViT, on the challenging dataset COCO Lin et al. (2014) as Tab. 7. We can see that, in stage 2, our technique outperforms the GroupViT by 4% despite using fewer data. Furthermore, we discuss existing limitations about complex situations in Appendix E with additional experiments and visualization.

Table 7: Comparison with previous state-of-the-arts on COCO 2017 validation set. FreeSeg achieves higher results with fewer data on the challenged transfer dataset.

Method	Data	Stage 1 (%)	Stage 2(%)
GroupViT Xu et al. (2022)	GCC12M+YFCC14M	-	24.3
FreeSeg	GCC3M	20.4	28.3

Visualization. We compare the visual results between stage 1 and stage 2 in Fig. 3. The results suggest that there are some voids and noise in stage 1, while stage 2 solves these issue regions well. Because, background threshold cause these noisy pixels in the first stage, and stage 2 refines the results learning from other samples, without such threshing operation.

Then we draw qualitative visual results from VOC dataset in Figure 4, and compare our FreeSeg with GroupViT. From these examples, we can see that FreeSeg does a fantastic job for some cases, which are quite close to ground truths, and beyond previous SoTA a lot in visualization. Furthermore, we depict additional visualization results in Appendix B.



Figure 3: Visual comparison between stage 1 and stage 2 of FreeSeg. The partially supervised semantic segmentation refines the voids and noise via learning from other images.



Figure 4: Qualitative visualization results compared with GroupViT Xu et al. (2022)

4.4 CONCLUSION.

In summary, we propose a semantic segmentation framework FreeSeg, which uses natural language supervision, for open world segmentation, via freely available mask from the raw feature map. Different from most zero-shot segmentation methods, FreeSeg doesn't require manual mask and seen classes. Instead, it generates segments via the Image-Text Similarity Map (ITSM) from Interpretable Contrastive Language-Image Pretraining (ICLIP), without extra algorithm or proprietary encoder. We scale up ICLIP for dense ITSM to support the segmentation task, and solve the consequent partial activation problem, by the proposed Smoothed Min Pooling (SMP). Then we fulfil the process from dense ITSM to segments, with the proposed partially supervised semantic segmentation based on partial label and pixel. Compared with previous SoTA, our FreeSeg is straight forward, data-efficient, and greatly improvements the performance.

REFERENCES

- Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, pp. 4981–4990, 2018.
- Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What the point: Semantic segmentation with point supervision. In *ECCV*, pp. 549–565. Springer, 2016.
- Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32, 2019.
- Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 9650–9660, 2021.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3558–3568, 2021.
- Hila Chefer, Shir Gur, and Lior Wolf. Generic attention-model explainability for interpreting bimodal and encoder-decoder transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 397–406, 2021a.
- Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 782–791, 2021b.
- Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002.
- Caio CV da Silva, Keiller Nogueira, Hugo N Oliveira, and Jefersson A dos Santos. Towards open-set semantic segmentation of aerial images. In *2020 IEEE Latin American GRSS & ISPRS Remote Sensing Conference (LAGIRS)*, pp. 16–21. IEEE, 2020.
- Jifeng Dai, Kaiming He, and Jian Sun. Boxesup: Exploiting bounding boxes to supervise convolutional networks for semantic segmentation. In *ICCV*, pp. 1635–1643, 2015.
- Karan Desai, Gaurav Kaul, Zubin Aysola, and Justin Johnson. Redcaps: Web-curated image-text data created by the people, for the people. *arXiv preprint arXiv:2111.11431*, 2021.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88(2): 303–338, 2010.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Philipp Krähenbühl and Vladlen Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Q. Weinberger (eds.), *Advances in Neural Information Processing Systems*, volume 24, pp. 109–117. Curran Associates, Inc., 2011. URL <https://proceedings.neurips.cc/paper/2011/file/beda24c1e1b46055dff2c39c98fd6fc1-Paper.pdf>.
- Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022a.
- Yi Li, Zhanghui Kuang, Liyang Liu, Yimin Chen, and Wayne Zhang. Pseudo-mask matters in weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 6964–6973, 2021.

- Yi Li, Yiqun Duan, Zhanghui Kuang, Yimin Chen, Wayne Zhang, and Xiaomeng Li. Uncertainty estimation via response scaling for pseudo-mask noise mitigation in weakly-supervised semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 1447–1455, 2022b.
- Yi Li, Hualiang Wang, Yiqun Duan, Hang Xu, and Xiaomeng Li. Exploring visual interpretability for contrastive language-image pre-training. *arXiv preprint arXiv:2209.07046*, 2022c.
- Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for semantic segmentation. In *CVPR*, pp. 3159–3167, 2016.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pp. 740–755. Springer, 2014.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- MMLab. Mmsegmentation, an open source semantic segmentation toolbox. <https://github.com/open-mmlab/msegmentation>, 2020. Accessed: 2021-05-04.
- Pedro O Pinheiro and Ronan Collobert. From image-level to pixel-level labeling with convolutional networks. In *CVPR*, pp. 1713–1721, 2015.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pp. 8748–8763. PMLR, 2021.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2556–2565, 2018.
- Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Transactions on pattern analysis and machine intelligence*, 22(8):888–905, 2000.
- Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. *arXiv preprint arXiv:2206.07045*, 2022.
- Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pp. 10347–10357. PMLR, 2021.
- Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10052–10062, 2021.
- Yunchao Wei, Huaxin Xiao, Honghui Shi, Zequn Jie, Jiashi Feng, and Thomas S Huang. Revisiting dilated convolution: A simple approach for weakly-and semi-supervised semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7268–7277, 2018.
- Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18134–18144, 2022.
- Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*, 2021.

Nir Zabari and Yedid Hoshen. Semantic segmentation in-the-wild without seeing any segmentation examples. *arXiv preprint arXiv:2112.03185*, 2021.

Xiaohua Zhai, Xiao Wang, Basil Mustafa, Andreas Steiner, Daniel Keysers, Alexander Kolesnikov, and Lucas Beyer. Lit: Zero-shot transfer with locked-image text tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18123–18133, 2022.

Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2881–2890, 2017.

Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 2921–2929, 2016.

Chong Zhou, Chen Change Loy, and Bo Dai. Denseclip: Extract free dense labels from clip. *arXiv preprint arXiv:2112.01071*, 2021.

A ANALYSIS FOR SMOOTHED MIN POOLING

Minimum pooling is hardly used in previous visual task, while it’s a core module of FreeSeg to achieve high quality dense Image-Text Similarity Map (ITSM). The reason of its success has been explained: min pooling mitigates the most discriminative features which causes the problem of partial activation.




Figure 5: Min pooling still focuses on discriminative regions without feature shift as average-like pooling methods. The points indicate the specific pixels of the feature map before pooling, which is most close to the pooled value for each channel. Larger points mean larger pooled values. Note AttnPool indicates the original class token (attention pooling).

In this part, we have two clarification about Smoothed Min Pooling (SMP). Firstly, SMP still focuses on discriminative regions without feature shift. For this concept, feature shift, it’s proposed in ICLIP Li et al. (2022c), which leads to dispersed feature localization from overlapped and aggregated

condition. This shift is related to global-average-like pooling methods, then cause the semantic shift and opposite visualization. We draw the matched points in Fig. 5, these points indicate the feature scatters before pooling, which is most close to the pooled value for each channel, and larger points mean larger pooled values. From this figure, we can see that min pooling still focuses on discriminative regions as max pooling, presenting overlapped and aggregated status.

The second clarification is that min pooling mitigates the most discriminative pixels via *cross representations*. To be specific, the concept, cross representations, indicates a group of features contributed to foreground regions are replaced by another group of features, and vice versa for background. Because, the max pooling selects the peak signals of a channel, while min pooling picks the foot signals at reversed meaning. In general, the foot signals are much smoother than peak signals, thus avoiding partial activation. To verify this cross representations, we compare the weights of last image linear projection layers. We find that there are 32227 weights occurring sign reversal between min pooling models (k19 and k13 in Fig. 3), and the number raises to 95381 between max pooling and min pooling models in Fig. 2. As the image encoder is fixed, and frequent sign reversal indicates different groups of features are selected. And this observation is corresponding to the behavior of cross representation via min pooling.

B ADDITIONAL QUALITATIVE VISUALIZATION RESULTS

We draw additional visualization results in this part. Firstly, we show the qualitative results of COCO dataset Lin et al. (2014) in Fig. 6. Although this data is challenging for zero-shot setting, our FreeSeg performs good in some examples and beyond the previous SoTA GroupViT Xu et al. (2022).




Figure 6: Qualitative visualization results on COCO dataset

We also show more visualization results on VOC datasets as Fig 7. We firstly compare our dense ITSM with smoothed min pooling with gradient based visualization method Bi-Module Chefer et al. (2021a), and our method performs much better than it. As shown in the appendix of ICLIP Li et al. (2022c), Bi-Modle performs much worse on larger output, so we use its best version ViT-B/32 to compare. Our dense ITSM is pretty good at details, which allows segmentation directly from raw feature maps. Secondly, we list our results of stage 1 & 2, compared with stage 2 of GroupViT. From these images, we can see that our methods are much better than previous work at details, and stage 1 beyond it even without refinement.

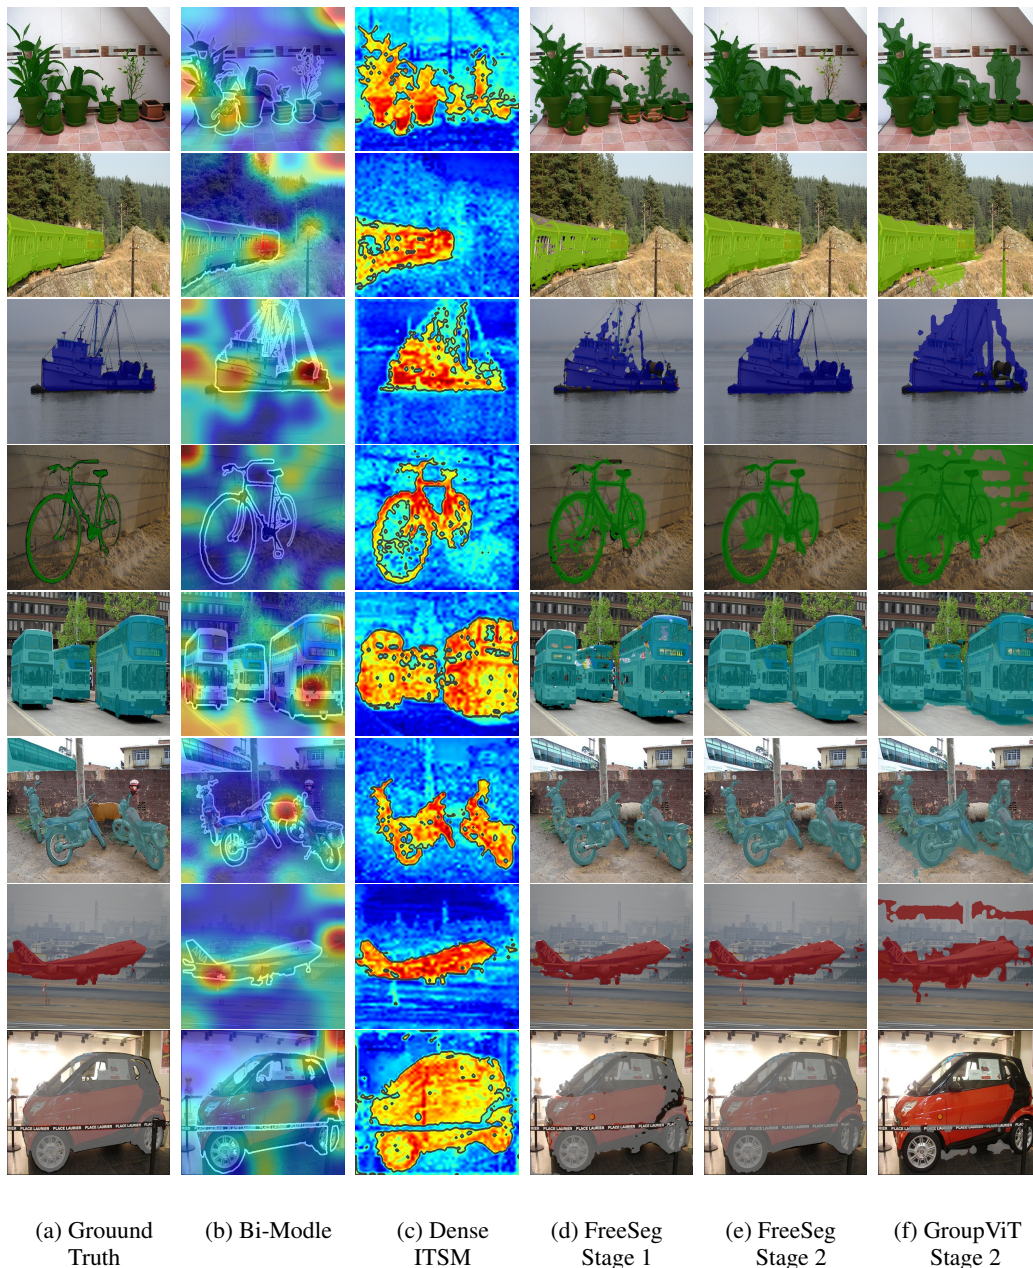


Figure 7: Visual comparison of interpretability (b vs. c) and segmentation (d, e, f) on VOC dataset.

C VISUAL COMPARISON AMONG POOLING METHODS

As shown in Tab. 2, the smoothing operation and min pooling greatly increase the performance of dense ITSM. Here, we visualize their improvements on VOC dataset Everingham et al. (2010) to depict how it works at dense ITSM. We can see that original max pooling focus on the most discriminative partial regions, and leaves many voids. After applying the smoothing operation, most voids are fixed. Since this smoothing is deployed in the training phase, the dense ITSM is not blurred. Then, the min pooling further improves the quality and achieves better details.




Figure 8: Visualization of dense ITSM from different pooling methods. Smoothed operation alleviates the problem of partial activation, which is most serious in the original max pooling. And the min pooling further solve this problem in above cases.

D FREELY AVAILABLE MASK

We emphasize that our approach generates segmentation masks with the least operation cost, and it's almost freely available from the pretraining model. Because, we directly obtain the mask from

its original predictions (image tokens), without extra design of image encoder or extra algorithm. And we compare the operations of mask generation in Tab. 8. In some extent, most text driven zero-shot segmentation methods Xu et al. (2021); Zhou et al. (2021); Li et al. (2022a) are similar to semi-supervised segmentation, with requires partial mask annotations. While some works Xu et al. (2022); Shin et al. (2022); Zabari & Hoshen (2021) are more close to unsupervised methods about mask generation. And our FreeSeg is similar to localization or weakly supervised segmentation, which generate masks from the last feature map, too. Besides, our settings are the same to GroupViT Xu et al. (2022), and we compare the quantitative results in Tab. 6. While other methods in Tab. 8 introduce manual mask or use additional algorithms in different settings, thus we only compare the operations of mask acquisition below.

Table 8: FreeSeg generates the segmentation mask from the raw feature map of the interpretable pretraining model, which is straight forward and almost freely available.

Method	Mask Acquisition
Xu et al. (2021); Zhou et al. (2021); Li et al. (2022a)	Learn from mask annotations
Shin et al. (2022)	Extra retrieval algorithm
Xu et al. (2022)	Proprietary grouping image encoder
Zabari & Hoshen (2021)	Unsupervised segmentation + Interpretability
FreeSeg	Raw feature map of pretraining model

E FAILURE CASES AND LIMITATIONS

We address the main limitations in this part, which requiring further improvements. Specifically, the CLIP is not good at zero-shot multi-label classification, because one image often contains multiple objects, while the text descriptions usually miss some of them. In another word, the supervision of natural language is noisy. As a result, CLIP only return few pass labels in Eq. 8, thus some objects are neglected. GroupViT Xu et al. (2022) handles this problem by text augmentation for multi-label contrastive loss, but it doesn't help when the object is not referred by the text. So, this phenomenon is universal in open vocabulary segmentation, as shown in Eq. 9. Besides, the limitations of ICLIP Li et al. (2022c) are inherited, such as worse performances on complex and certain categories.




Figure 9: Failure cases of missed categories.

Besides, our method requires retraining the model as GroupViT, and varied datasets influence the final results. For example, when GroupViT replaces FYCC Thomee et al. (2016) to RedCaps Desai et al. (2021), the gap between FreeSeg and GroupViT raises to 14.9% on VOC but meets a reduction by 3.2% on COCO. Besides, if GroupViT uses the same training dataset as ours, its result on COCO is merely 6.8% of stage 1.