

ing procedure. Most of the above methods necessitate the use of additional complex frameworks to facilitate training, resulting in a complicated and time-consuming training process.

In recent years, vision transformers [12] have steadily evolved in various fields [13, 14]. The self-attention mechanism in the vision transformer can capture effective relationships between pixels, which is beneficial for information propagation from labeled pixels to unlabeled pixels. Inspired by this, We investigate single-stage effective sparsely annotated semantic segmentation with vision transformer, aiming to achieve the state-of-the-art results in a straightforward manner, as shown in Fig 1.

In this work, we propose SASFormer, a simple yet effective framework based on segformer[12], as the first effort at sparsely annotated semantic segmentation with Vision Transformer. With the inherent patch-to-patch attention of cascaded transformer blocks, we first create patch attention maps at multiple levels. The patch-level attention maps are then multiplied by network predictions in order to correlate labeled and unlabeled image regions. Following this, an affinity loss function is applied to assure the consistency of the correlation statistics and network predictions. We finally consider embedding priors from unlabeled areas to boost the performance by combining affinity loss and conventional segmentation loss. Our contributions can be summarized in the following manner:

- We propose SASFormer as the first trustworthy baseline for SASS with a visual transformer to model relationships between distinct areas and offer category recommendations to unlabeled regions.
- We provide a novel affinity loss function to assure the similarity prior in SASS: areas of identical objects share similarity in both low-dimensional and high-dimensional feature space.
- The proposed SASFormer manifests remarkable performances in both point- and scribble-annotated SASS assignments on PASCAL VOC.

2. RELATED WORKS

2.1. Sparsely Annotated Semantic Segmentation

Sparsely annotated semantic segmentation aims to address the segmentation problem using minimally labeled visual areas. What's the point[2] brought point-level labels to the semantic segmentation problem for the first time, which boost the performance by introducing objectiveness priors. After that, ScribbleSup[4] extends point-level labels into the scribbling ones in order to bridge the performance gap between fully annotated methods and unlabeled ones. ScribeSup constructs a graphical model to transmit information from scribbles to unknown pixels, hence providing semantic guidance

for unknown pixels. Since then, approaches for sparse annotated semantic segmentation have emerged and developed. In an attempt to enhance performance at the lowest feasible cost, researchers have begun investigating the association between labeled and unlabeled regions. For example, Tang et al. [7] offered a variety of regularized losses for sparsely supervised segmentation based on dense CRF and kernel cut. BPG [9] tries to integrate semantic characteristics and textural information by constructing a prediction refinement network, whilst a boundary regression network is introduced to facilitate performance by generating clearly defined semantically separate parts. Recently, SPML [11] tackles this challenge by developing a semi-supervised metric learning approach with four unique forms of attraction and repulsion relationships. TEL [1] presents a tree energy loss, in which minimal spanning trees are constructed to represent low-level and high-level pair-wise affinities. The majority of these techniques either use multi-stage processes for progressive inference or include intricate frameworks for supplementary training. While in this work, we make the first attempt to model relationships between distinct areas and offer category recommendations to unlabeled regions based on vision transformer.

2.2. Transformer

The Transformer was initially proposed to model long-term dependencies in natural language processing tasks [15]. In 2020, Alexey Dosovitskiy introduced the pure Transformer architecture, which achieved remarkable results in image classification [16]. Subsequently, the Transformer has been widely employed in various computer vision tasks, such as object detection [17, 18, 19], semantic segmentation [12, 20] and video processing [21, 22]. Recently, researchers have started to utilize the long-term dependencies mechanism of the Transformer to capture the relationship between image categories and local image features, aiming to optimize weakly supervised problems [23]. However, these methods mainly focus on utilizing the correlation between class tokens and patch tokens, while few works consider the optimization and utilization of the correlation among different patch tokens. Our work takes an early attempt to delve deeply into the issue of imprecise correlation among different patch tokens in the Transformer block and proposes a simple yet effective solution.

3. METHODS

3.1. Overview

As depicted in Fig 2, given an image $I \in \mathbb{R}^{W \times H \times 3}$, we partition it into $p \times p$ resolution patches. These patches are fed into hierarchical transformer encoder to provide multilevel characteristics. We pass multi-level features into the decoder block in order to obtain segmentation probabilities \mathbf{P} with resolution $M \times C$, where M is the patch number and C is the num-

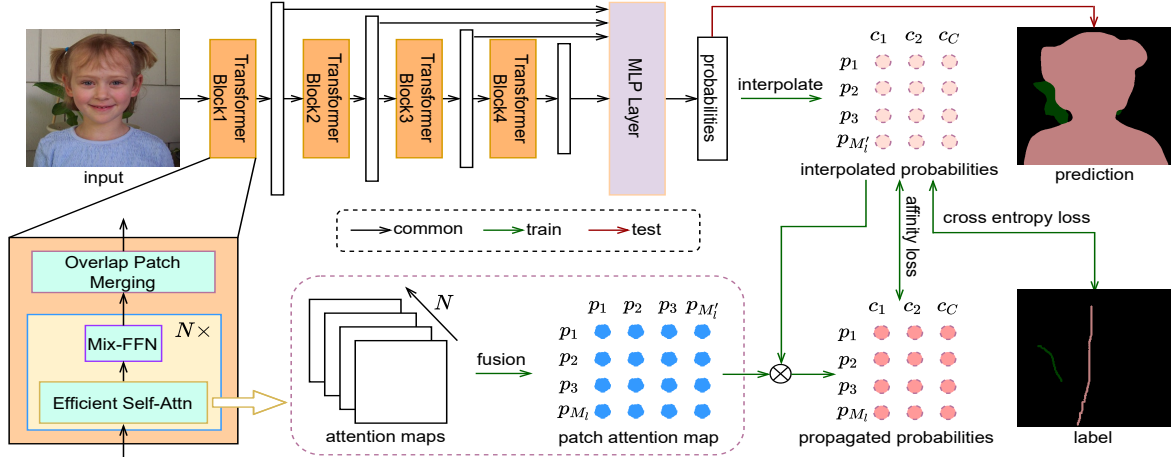


Fig. 2. An overview of the proposed Sparsely Annotated Segmentation Transformer(SasFormer)

ber of categories. On the one hand, we build multi-level patch attention maps with transformer layers in encoder blocks during the training phase. In addition, segmentation probabilities are interpolated into resolutions corresponding to the width of patch attention maps. Patch attention maps are then multiplied by interpolated segmentation probabilities, yielding propagated segmentation probabilities \mathbf{Y} . For labeled regions of interest, we perform standard cross entropy loss function L_{seg} to supervise segmentation probabilities with ground truth labels. For unlabeled regions, we introduce an affinity loss function L_{aff} to match segmentation probabilities to propagated segmentation probabilities. During testing, segmentation probabilities are immediately assigned to the class with the maximum probability at each pixel, resulting in the final segmentation map. The overall loss function is defined as follows:

$$L = L_{seg} + \alpha * L_{aff} \quad (1)$$

3.2. Patch Attention Map Generation

We introduce Segformer [12] as our backbone in order to describe our method more conveniently. More results of different backbones can be referred to Sec 4.3. Hierarchical transformer encoder contains $L = 4$ encoder blocks. Each encoder block contains N_l consecutive transformer layers. We get the patch attention map with the efficient self-attention of transformer layer, which is defined as follows:

$$\mathbf{A}_{l,n} = \text{softmax}\left(\frac{\mathbf{Q}_{l,n}\mathbf{K}_{l,n}^T}{\sqrt{D}}\right) \quad (2)$$

where $\mathbf{Q}_{l,n} \in \mathbb{R}^{M_l \times D}$ and $\mathbf{K}_{l,n} \in \mathbb{R}^{M'_l \times D}$ are the query and key representations of n -th transformer layer in l -th encoder block, respectively. $M_l \in \{\frac{W*H}{4}, \frac{W*H}{8}, \frac{W*H}{16}, \frac{W*H}{32}\}$ denotes resolution size in different encoder blocks. M'_l denotes resolution size from M_l after sequence reduction process[12]

to be more efficient. D indicates dimension of patch embeddings. \mathbf{T} is the transpose operator.

Efficient attention map $\mathbf{A}_{l,n} \in \mathbb{R}^{M_l \times M'_l}$ records dependency of each patch token of n -th transformer layer in l -th encoder block. We aggregate these attention maps over transformer layers in each decoder block as the following:

$$\mathbf{A}_l = \frac{1}{N} \sum_{n=1}^N \mathbf{A}_{l,n} \quad (3)$$

where $\mathbf{A}_l \in \mathbb{R}^{M_l \times M'_l}$ means patch attention map in l -th encoder block.

To better comprehend what information patch attention maps capture at various levels, we randomly draw a point with a red pentagram in the inputs and analyze the dependency between the reference point and other pixels on patch attention maps of various encoder blocks, as depicted in the baseline of Fig 3. $\mathbf{A}_{l,n}^{i,j}$ means the amount of j -th input token embedding contributes to i -th output token of n -th transformer layer in l -th encoder block. Likewise, \mathbf{A}_l^i means the aggregated amount of each input token embedding contributes to i -th output token in l -th encoder block. We construct \mathbf{A}_1^r , \mathbf{A}_2^r , \mathbf{A}_3^r and \mathbf{A}_4^r to illustrate dependency between the reference point and each pixel in the first, second, third, forth encoder block, respectively.

It can be seen that \mathbf{A}_1 and \mathbf{A}_2 prefer to assign larger weights to patches that have the same color and texture as the reference point. For instance, \mathbf{A}_1^r and \mathbf{A}_2^r of the first input emphasize the white color of the bird, while the second input highlights the dark brown color of the material. \mathbf{A}_3 and \mathbf{A}_4 are more concerned with semantic consistency, although \mathbf{A}_4 is more class-specific. Patch attention maps establish correlations between pixels in color space and high-level characteristics, facilitating the propagation of segmentation information from labeled to unlabeled pixels. However, transformer tends to distract attention from targets to background

area[18], which can be clearly seen in those of baseline(A_2) and baseline(A_3). We hypothesize that the inability of attention maps to capture dependence, which impacts the transfer of information across pixels, is partly manifested in distractions from irrelevant background.

3.3. Affinity Loss Function

For each encoder block, we interpolate segmentation probabilities \mathbf{P} into the resolution of M'_l and M_l , yielding \mathbf{P}'_l and \mathbf{P}_l respectively. Then we transfer segmentation predictions from labeled pixels to unlabeled pixels by multiplying patch attention map \mathbf{A}_l and interpolated segmentation probabilities \mathbf{P}'_l , which is described as follows:

$$\mathbf{Y}_l = \mathbf{A}_l \otimes \mathbf{P}'_l \quad (4)$$

where \mathbf{Y}_l is propagated segmentation probabilities, \otimes denotes matrix multiplication.

\mathbf{Y}_l and \mathbf{P}_l are normalized along the category dimension, which is defined as:

$$\mathbf{Y}_l^{*(m_l, c)} = \frac{\exp(\mathbf{Y}_l^{(m_l, c)})}{\sum_{c=1}^C \exp(\mathbf{Y}_l^{(m_l, c)})} \quad 1 \leq m_l \leq M_l \quad (5)$$

$$\mathbf{P}_l^{*(m_l, c)} = \frac{\exp(\mathbf{P}_l^{(m_l, c)})}{\sum_{c=1}^C \exp(\mathbf{P}_l^{(m_l, c)})} \quad 1 \leq m_l \leq M_l \quad (6)$$

where m_l and c denote patch index and class index. $\mathbf{Y}_l^{(m_l, c)}$ and $\mathbf{P}_l^{(m_l, c)}$ mean the position of row m_l , column c of matrix \mathbf{Y}_l and \mathbf{P}_l .

The affinity loss L_{aff} attempts to maximize the similarity between propagated probabilities and the interpolated segmentation probabilities over all encoder blocks. Formally, the affinity loss is defined as:

$$L_{aff} = \frac{1}{L} \sum_{l=1}^L \|\mathbf{Y}_l^* - \mathbf{P}_l^*\|_1 \quad (7)$$

where $\|\cdot\|_1$ and L denote L1 regularization term and encoder block number, respectively.

In the process of segmentation training, we obtain additional bonus information, as illustrated in Fig 3. Under the constraint of consistency between propagated and interpolated segmentation probabilities, pair-wise dependencies between patch tokens with different labels on patch attention maps are simultaneously deteriorated, while pair-wise dependencies between patch tokens with the same label are enhanced. By comparing to the baseline in Fig 3, it is evident that the background noise in the patch attention map has been suppressed in our proposed SASFormer.

4. EXPERIMENTS

4.1. Datasets and Settings

The point-wise annotation [2] and scribble-wise annotation [4] of Pascal VOC 2012 dataset [24] are utilized for point-supervised and scribble-supervised settings, respectively. We adopt the segformer with pre-trained weights on

Table 1. Experimental results for **point-wise** annotations on Pascal VOC 2012 validation set.

Methods	Backbone	multi-stage	mIoU(%)
DenseCRF Loss [7]	ResNet101	✓	57.00
A2GNN [25]	ResNet101	✓	66.80
Seminar [10]	ResNet101	✓	72.51
What's the point [2]	VGG16	-	43.40
TEL [1]	ResNet101	-	68.40
Baseline*	Segformer	-	64.21
TEL* [1]	Segformer	-	69.33
Sasformer (Ours)	Segformer	-	73.13

ImageNet as our backbone for experiments. Our model is trained for 80k training epochs with initial learning rate 0.001, batch size 16 and input resolution 512x512. The SGD optimizer is used with momentum 0.9 and weight decay 1e-4. Random horizontal flip, random brightness in [-10, 10], random resize in [0.5, 2.0] and random crop are all employed in data augmentation. In practice, we set α in Eq.1 1.2 for scribble setting and 0.2 for point setting. All experiments are conducted on PyTorch with four 32GB V100 GPUs. * indicates that we reproduce the results using the provided codes, unless otherwise specified.

4.2. Comparison with State-of-the-art Methods

Table 2. Experimental results for **scribble-wise** annotations on Pascal VOC 2012 validation set.

Methods	Backbone	multi-stage	mIoU (%)
ScribbleSup [4]	VGG16	✓	63.10
DenseCRF Loss [7]	ResNet101	✓	75.00
GridCRF Loss [8]	ResNet101	✓	72.80
URSS [26]	ResNet101	✓	76.10
Seminar [10]	ResNet101	✓	76.20
A2GNN [25]	ResNet101	✓	76.20
BPG [9]	ResNet101	-	76.00
SPML [11]	ResNet101	-	76.10
PSI [5]	ResNet101	-	74.90
TEL [1]	ResNet101	-	77.30
Baseline*	Segformer	-	73.37
TEL* [1]	Segformer	-	78.53
Sasformer (Ours)	Segformer	-	79.49

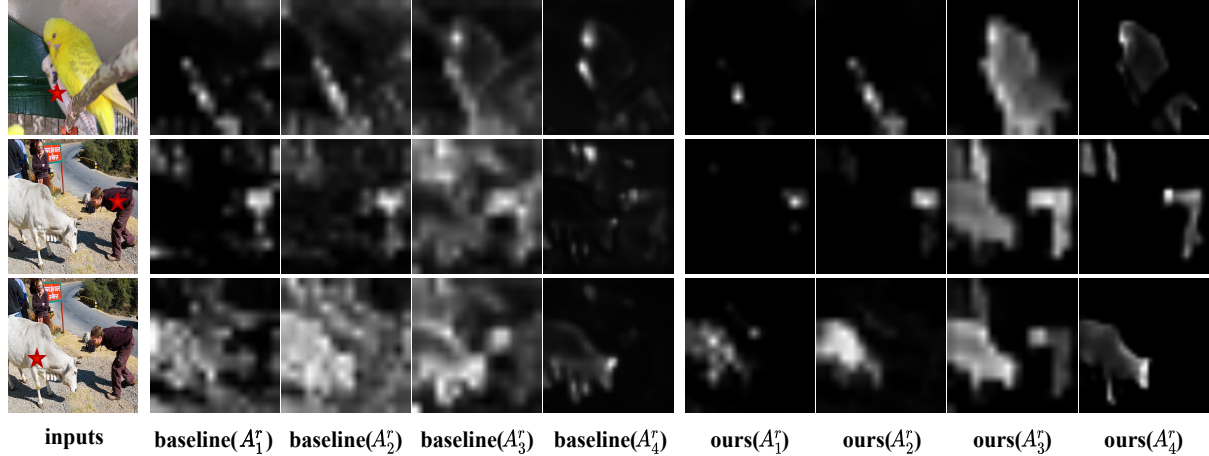


Fig. 3. Visualization of inputs with reference points (red pentagram) and patch attention maps at various levels. The baseline indicates the Segformer [12] approach. A_1^r , A_2^r , A_3^r and A_4^r illustrate the dependency between the reference point and each pixel of patch attention map in the first, second, third and forth encoder block, respectively.

Tab 1 and Tab 2 present the experimental results on the Pascal VOC 2012 for point-wise and scribble-wise annotations, respectively. In Tab 1 for point-wise annotations, the baseline based on segformer-B4 employs the cross-entropy loss function only can reach 64.2% mIoU. TEL [1] improves mIoU performance by 5.13%. Our proposed SASFormer outperforms TEL by 4.43% with the same backbone, yielding a mIoU of 72.78%. As for scribble-wise annotations in Tab 2, the baseline framework achieves a mIoU of 73.3%. SASFormer produces the best performance with the mIoU of 79.49%, surpassing the baseline by 6.19%. It demonstrates that SASFormer achieves remarkable performance compared to all other state-of-the-art approaches, which adopt a single-stage training strategy without any additional data. The qualitative results for point-wise annotations and scribble-wise annotations showcase that not only does our technique capture object outlines more accurately, but it also minimizes category misidentification, as shown in Fig 4.

4.3. Ablations

Table 3. Effect of patch attention map in different encoder blocks.

	L_{aff}^1	L_{aff}^2	L_{aff}^3	L_{aff}^4	mIoU (%)
exp1	-	-	-	-	73.37
exp2	✓	-	-	-	77.38
exp3	-	✓	-	-	76.97
exp4	-	-	✓	-	75.98
exp5	-	-	-	✓	74.10
exp6	✓	✓	-	-	79.06
exp7	✓	✓	✓	-	79.24
exp8	✓	✓	✓	✓	79.49

Single scale vs. Multi scale. We examine the influence of

patch attention maps in various encoder blocks on segmentation performance in Tab 3, where L_{aff}^l denotes the affinity loss generated by the patch attention map of the l -th encoder block. The optimal segmentation performance is achieved when patch attention maps are incorporated into the generation of affinity loss in different encoder blocks.

Table 4. The effect of different loss configurations.

Configuration	KL	CE	L2	L1
mIoU (%)	78.62	76.80	77.25	79.49

The impact of distance metrics. For the similarity between propagated and interpolated segmentation probabilities, in addition to L1 distance in Eq.7, cross entropy, L2 distance, and KL divergence may also be considered. To evaluate the performance of our SASFormer, we evaluate different loss configurations. As seen in Tab 4, different forms of affinity loss can also improve performance. We choose L1 distance as the final implementation of our SASFormer since it produces the best results (79.49% mIoU).

Table 5. The flexibility of our proposed approach.

Backbone	Ours	mIoU (%)
Segmentor [27]	w/o	72.75
	w/	76.11 (+3.36)
SETR [28]	w/o	70.47
	w/	76.04 (+5.57)
segformer [12]	w/o	73.37
	w/	79.49 (+6.12)

The flexibility among backbones. To evaluate the flexibility of our technique across a variety of transformer-based segmentation networks, we also present the popular segmentor [27] and SETR [28] in order to compare segmentation performance with and without affinity loss. As shown in Tab 5,

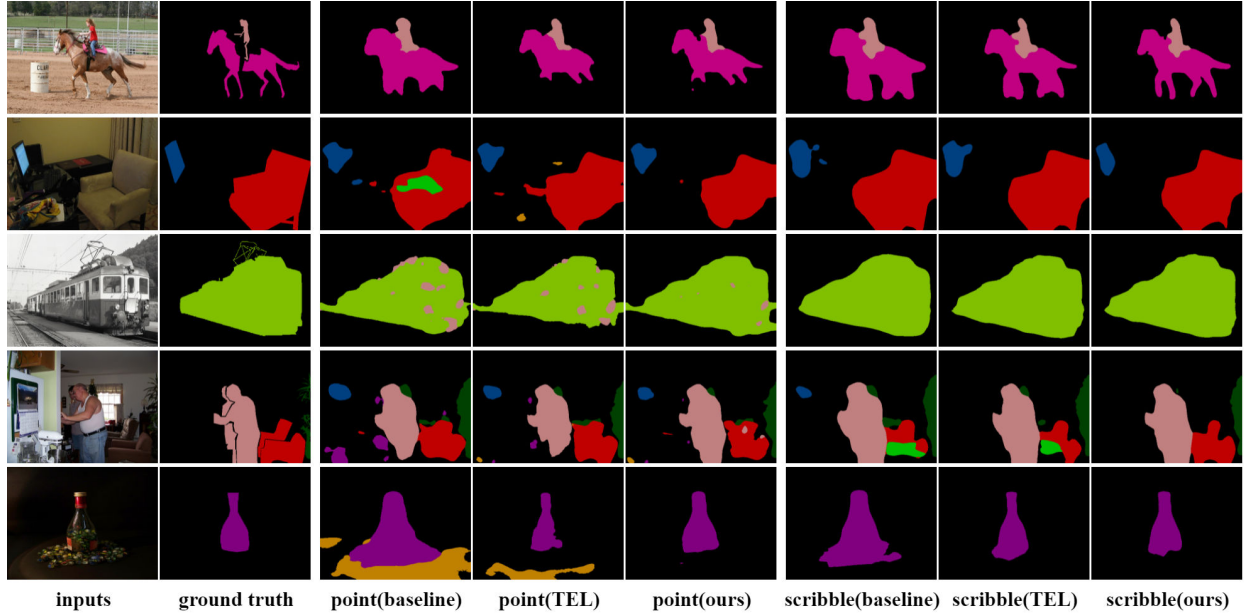


Fig. 4. Segmentation results on Pascal VOC 2012 dataset. The point and scribble indicate the point and scribble supervision. The baseline and TEL indicate the Segformer [12] and TEL [1] approach, respectively.

our affinity loss is robust across a variety of segmentation networks.

5. CONCLUSIONS

In this work, we propose an efficient framework, dubbed SASFormer, to deal with the issue of semantic segmentation under sparsely annotated supervision. Most of the existing approaches are always time-consuming and often necessitate a multi-stage training strategy. While the proposed SASFormer adopts an end-to-end scheme that leverages hierarchical patch attention maps. In addition, we also introduce the affinity loss to capture the consistency between the features of correlation results and network predictions. Exhaustive experiments validate the effectiveness of the proposed approach and showcase remarkable performance compared to state-of-the-art approaches.

Acknowledgments

This work is partially supported by the National Natural Science Foundation of China (Grant No. 62106235), by the Exploratory Research Project of Zhejiang Lab(2022PG0AN01), by the Zhejiang Provincial Natural Science Foundation of China (LQ21F020003).

6. REFERENCES

- [1] Zhiyuan Liang, Tiancai Wang, Xiangyu Zhang, Jian Sun, and Jianbing Shen, “Tree energy loss: Towards sparsely annotated semantic segmentation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 16907–16916.
- [2] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei, “What’s the point: Semantic segmentation with point supervision,” in *European conference on computer vision*. Springer, 2016, pp. 549–565.
- [3] Rui Qian, Yunchao Wei, Honghui Shi, Jiachen Li, Jiaying Liu, and Thomas Huang, “Weakly supervised scene parsing with point-based distance metric learning,” in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2019, vol. 33, pp. 8843–8850.
- [4] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun, “Scribblesup: Scribble-supervised convolutional networks for semantic segmentation,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 3159–3167.
- [5] Jingshan Xu, Chuanwei Zhou, Zhen Cui, Chunyan Xu, Yuge Huang, Pengcheng Shen, Shaoxin Li, and Jian Yang, “Scribble-supervised semantic segmentation inference,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 15354–15363.
- [6] Dingwen Zhang, Wenyan Zeng, Guangyu Guo, Chaowei Fang, Lechao Cheng, and Junwei Han, “Weakly supervised semantic segmentation via alternative self-dual teaching,” *arXiv preprint arXiv:2112.09459*, 2021.
- [7] Meng Tang, Federico Perazzi, Abdelaziz Djelouah, Ismail Ben Ayed, Christopher Schroers, and Yuri Boykov, “On regularized losses for weakly-supervised cnn seg-

- mentation,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 507–522.
- [8] Dmitrii Marin, Meng Tang, Ismail Ben Ayed, and Yuri Boykov, “Beyond gradient descent for regularized segmentation losses,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 10187–10196.
 - [9] Bin Wang, Guojun Qi, Sheng Tang, Tianzhu Zhang, Yunchao Wei, Linghui Li, and Yongdong Zhang, “Boundary perception guidance: A scribble-supervised semantic segmentation approach,” in *IJCAI International joint conference on artificial intelligence*, 2019.
 - [10] Hongjun Chen, Jinbao Wang, Hong Cai Chen, Xiantong Zhen, Feng Zheng, Rongrong Ji, and Ling Shao, “Seminar learning for click-level weakly supervised semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 6920–6929.
 - [11] Tsung-Wei Ke, Jyh-Jing Hwang, and Stella X Yu, “Universal weakly supervised segmentation by pixel-to-segment contrastive learning,” *arXiv preprint arXiv:2105.00957*, 2021.
 - [12] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo, “Segformer: Simple and efficient design for semantic segmentation with transformers,” in *Neural Information Processing Systems (NeurIPS)*, 2021.
 - [13] Chaowei Fang, Dingwen Zhang, Liang Wang, Yulun Zhang, Lechao Cheng, and Junwei Han, “Cross-modality high-frequency transformer for mr image super-resolution,” *arXiv preprint arXiv:2203.15314*, 2022.
 - [14] Mengqi Xue, Qihan Huang, Haoifei Zhang, Lechao Cheng, Jie Song, Minghui Wu, and Mingli Song, “Protopformer: Concentrating on prototypical parts in vision transformers for interpretable image recognition,” *arXiv preprint arXiv:2208.10431*, 2022.
 - [15] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” *Advances in neural information processing systems*, vol. 30, 2017.
 - [16] 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.
 - [17] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko, “End-to-end object detection with transformers,” in *European conference on computer vision*. Springer, 2020, pp. 213–229.
 - [18] Hui Su, Yue Ye, Zhiwei Chen, Mingli Song, and Lechao Cheng, “Re-attention transformer for weakly supervised object localization,” *arXiv preprint arXiv:2208.01838*, 2022.
 - [19] Tian Qiu, Linyun Zhou, Wenxiang Xu, Lechao Cheng, Zunlei Feng, and Mingli Song, “Team-detr: Guide queries as a professional team in detection transformers,” *arXiv preprint arXiv:2302.07116*, 2023.
 - [20] Hao Li, Dingwen Zhang, Nian Liu, Lechao Cheng, Yalun Dai, Chao Zhang, Xinggang Wang, and Junwei Han, “Boosting low-data instance segmentation by unsupervised pre-training with saliency prompt,” *arXiv preprint arXiv:2302.01171*, 2023.
 - [21] Hao Zhang, Yanbin Hao, and Chong-Wah Ngo, “Token shift transformer for video classification,” in *Proceedings of the 29th ACM International Conference on Multimedia*, 2021, pp. 917–925.
 - [22] Hao Zhang, Lechao Cheng, Yanbin Hao, and Chongwah Ngo, “Long-term leap attention, short-term periodic shift for video classification,” in *Proceedings of the 30th ACM International Conference on Multimedia*, 2022, pp. 5773–5782.
 - [23] Wei Gao, Fang Wan, Xingjia Pan, Zhiliang Peng, Qi Tian, Zhenjun Han, Bolei Zhou, and Qixiang Ye, “Ts-cam: Token semantic coupled attention map for weakly supervised object localization,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 2886–2895.
 - [24] 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*, vol. 88, no. 2, pp. 303–338, 2010.
 - [25] Bingfeng Zhang, Jimin Xiao, Jianbo Jiao, Yunchao Wei, and Yao Zhao, “Affinity attention graph neural network for weakly supervised semantic segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021.
 - [26] Zhiyi Pan, Peng Jiang, Yunhai Wang, Changhe Tu, and Anthony G Cohn, “Scribble-supervised semantic segmentation by uncertainty reduction on neural representation and self-supervision on neural eigenspace,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7416–7425.
 - [27] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid, “Segmenter: Transformer for semantic segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 7262–7272.
 - [28] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip H.S. Torr, and Li Zhang, “Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers,” in *CVPR*, 2021.

A. SUPPLEMENTARY MATERIAL

To further demonstrate the effective of our approach, we conduct additional experiments on the Cityscapes and ADE20k datasets, which have segmentation annotations with varying levels of sparsity, including 10%, 20%, and 50% of full labels. As shown in Table 6, our SASFormer outperforms state-of-the-art methods at all sparsity levels. Moreover, we include quantitative results on the Cityscapes validation set in Fig. 5.

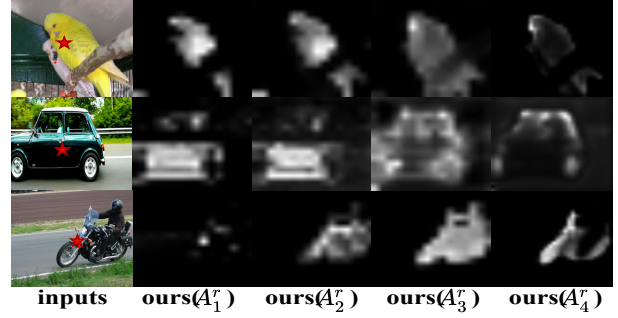


Fig. 6. Visualization of inputs with reference points (red pentagram) and patch attention maps at various levels.

Table 6. Experimental results on Cityscapes and ADE20k datasets with different sparsity levels.

Method	Params	Cityscapes			ADE20k		
		10%	20%	50%	10%	20%	50%
DenseCRF Loss [7]	70.0M	57.4	61.8	70.9	31.9	33.8	38.4
TEL [1]	70.0M	61.9	66.9	72.2	33.8	35.5	40.0
Segformer	61.4M	58.4	60.5	70.4	39.8	41.4	45.3
Ours	61.4M	64.6	70.0	75.6	41.8	44.6	46.8

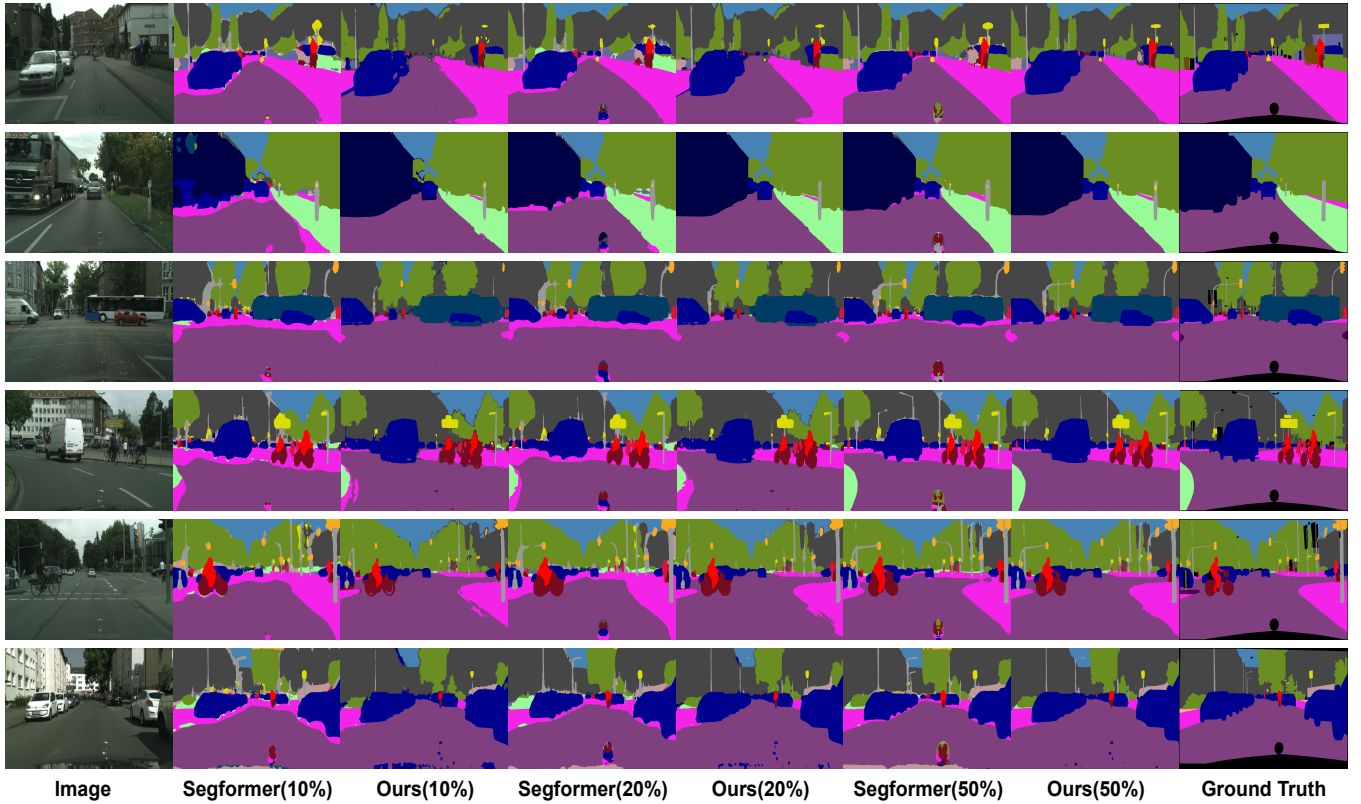


Fig. 5. Cityscapes val set segmentation results at 10%, 20%, 50% label proportions.