Progressive Proxy Anchor Propagation for Unsupervised Semantic Segmentation

Hyun Seok Seong^o, WonJun Moon^o, SuBeen Lee^o, and Jae-Pil Heo^{*}

Sungkyunkwan University

Abstract. The labor-intensive labeling for semantic segmentation has spurred the emergence of Unsupervised Semantic Segmentation. Recent studies utilize patch-wise contrastive learning based on features from image-level self-supervised pretrained models. However, relying solely on similarity-based supervision from image-level pretrained models often leads to unreliable guidance due to insufficient patch-level semantic representations. To address this, we propose a Progressive Proxy Anchor Propagation (PPAP) strategy. This method gradually identifies more trustworthy positives for each anchor by relocating its proxy to regions densely populated with semantically similar samples. Specifically, we initially establish a tight boundary to gather a few reliable positive samples around each anchor. Then, considering the distribution of positive samples, we relocate the proxy anchor towards areas with a higher concentration of positives and adjust the positiveness boundary based on the propagation degree of the proxy anchor. Moreover, to account for ambiguous regions where positive and negative samples may coexist near the positiveness boundary, we introduce an instance-wise ambiguous zone. Samples within these zones are excluded from the negative set, further enhancing the reliability of the negative set. Our state-of-the-art performances on various datasets validate the effectiveness of the proposed method for Unsupervised Semantic Segmentation. Our code is available at https://github.com/hynnsk/PPAP.

Keywords: Unsupervised Semantic Segmentation, Contrastive Learning

1 Introduction

Semantic Segmentation plays a vital role in various fields, including robotics and autonomous driving [6,8,18,36,37,49,51,55]. With the abundant data available in media, developing high-quality semantic segmentation models has become feasible [32], though this has also increased the demand for extensive human annotations. Likewise, the increasing burden on human labor has spurred the emergence of Unsupervised Semantic Segmentation (USS) [13, 22, 27, 29–31, 45, 46, 48, 54].

The main challenge in USS stems from the lack of supervision to train the model. To overcome this, prior works [22, 46, 54] suggested first learning the

^{*} Corresponding author



Fig. 1: (a) Illustration of how positive and negative sets are determined in HP [46] and Ours. Different colors indicate different potential classes. In HP, k-th nearest neighbor on a per-sample basis becomes an instance-wise positiveness criterion, and all other samples farther than k-th neighbor are considered as the negatives. On the other hand, we progressively propagate the proxy anchor to be relocated in the region surrounded by semantically similar samples. Consequently, we enable the trustworthy positive collection with numerous samples in dense regions. Additionally, we define an ambiguous region around the positiveness boundary where semantic boundaries might be ambiguous. By excluding the samples in the ambiguous region in training, we avoid the undesired repulsion between the anchor and the possibly *false positives* in the negative set. (b) The number of positives and their precision with respect to the ground truth label for a randomly sampled subset of the dataset. The X-axis represents anchors within the range of [0%, q%] in the anchor list, sorted in ascending order by the number of identified positives. The bar plot displays the average number of gathered positives, and the line plot illustrates their precision, as determined by the ground truth labels of both the anchor and the positives.

image-level representation space and then leveraging this knowledge to develop the ability of pixel-level understanding. Likewise, utilizing self-supervised pretrained models [3, 4, 24] to provide supervision in USS became mainstream. By employing these foundation models, previous techniques have demonstrated promising results, particularly by learning the relationship among image patches in the dataset [22, 46, 54].

Yet, we notice that existing methods still encounter challenges in discovering trustworthy relationships between patches. For instance, HP [46], the latest USS technique based on contrastive learning, utilized k-th nearest neighbor of each anchor to determine an appropriate boundary for positive set selection. While discretizing samples with such a boundary provides an intuitive basis, it often leads to unreliable supervision. This is because they exclusively rely on the similarity metric on a per-patch basis within an imperfect embedding learned in an unsupervised manner at the image-level [4]. Consequently, as illustrated in Fig. 1a, this approach may cause anchors to gather *false positives* (FP)¹ in the positive set especially when they are located in data-sparse areas or near the

¹ The italic 'false positives (FP)' and 'true positives (TP)' [50] represent samples that incorrectly and correctly included in the set, respectively, throughout the paper.

semantic boundaries while encouraging anchors in dense regions to repel FP in the negative set. Specifically, in Fig. 1b, we present a quantitative comparison between HP and our method, focusing on the number and precision of collected positives. As observed, we note that only 77.5% of the positives identified by HP match the ground truth labels, even though the number of positives is insufficient. Precision further decreases for instances in data-sparse regions; samples with a small number of gathered positives exhibit even lower precision (bottom 10% samples retain only 33.81% precision on discovered positives). This indicates the inclusion of a substantial number of FP in the positive set, thereby attracting semantically dissimilar samples and leading to unstable learning.

To mitigate these issues, we propose a Progressive Proxy Anchor Propagation (PPAP) strategy to deal with the vulnerability of the per-patch-based similarity metric in an image-level pretrained embedding space. Our goal is to establish a reliable proxy anchor by considering the data distribution surrounding each anchor, thereby gathering patches with more trustworthy positive and negative relationships minimizing ambiguity. This approach can also obtain a larger number of training guidance as it enhances the precision of gathered relationships. Specifically, to discover the position of the proxy anchor, we begin by defining a tight boundary around each anchor to construct a small, reliable positive set. The rationale behind establishing a tight boundary is rooted in the observation that samples within closely adjacent regions are highly likely to share similar semantics even within the image-level pretrained embedding space. Subsequently, we iteratively undertake the following two steps to enlarge a trustworthy positive set per anchor: 1) Re-define the position of a proxy anchor based on the distribution of identified positive samples, 2) Lower the similarity threshold for the positiveness criterion, *i.e.*, expand the boundary, according to the reliability of the new proxy anchor position, and gather the updated positive set. Likewise, by discovering the samples with similar semantics and moving the proxy anchor towards the center point of such samples, we expect the assembly of trustworthy positives. This strategy enables collecting a large number of positive patches with high precision, as shown in Fig 1b. Still, a positiveness boundary might not be a perfect measure to detect all the positive samples. In other words, there exists a degree of ambiguity around the boundary, where both positive and negative instances might coexist. To address this, we expand the original binary relationship categorization of samples, *i.e.*, positive and negative, for contrastive learning into tri-partite groups, *i.e.*, positive, negative, and ambiguous. The size of the ambiguous set is determined based on the reliability of the relocated proxy anchor. Consequently, while utilizing the positive and negative sets in contrastive learning, we disregard the ambiguous set, as including FP in the negative set often disrupts the stable training [11].

Overall, our contributions are summarized as follows:

 We propose Progressive Proxy Anchor Propagation (PPAP), which systematically gathers trustworthy positive samples for each anchor by progressively analyzing the distribution of the positive samples.

- 4 H.S. Seong et al.
- We establish an ambiguity-excluded negative set based on the propagated proxy anchor, defining a semantically ambiguous zone for each anchor. This approach effectively eliminates potential *FP* in the negative set.
- The efficacy of our trustworthy contrastive learning is validated by achieving new state-of-the-art performances across diverse datasets.

2 Related Work

2.1 Unsupervised Semantic Segmentation

Semantic segmentation aims to classify the semantics of individual pixels within an image [5, 7, 8, 17, 37, 42, 44, 55, 58]. In recent years, the integration of transformers into semantic segmentation has emerged as a promising research direction [12,35,51]. However, achieving pixel-wise supervision requires extensive human labor. The necessity of learning semantic segmentation without supervision has become apparent in recent literature [13, 27, 29, 30, 45, 46, 48, 54]. Earlier trials [13,27] learned to maintain consistent semantics across the paired features. In contrast, recent techniques [22, 46, 54] have employed Vision Transformer (ViT) models trained in a self-supervised manner as backbone networks to transfer knowledge to the segmentation head. For instance, transFGU [54] grouped target datasets based on prior knowledge and generated pseudo-labels to train the segmentation model. STEGO [22] tried to maintain the patch relationships in the segmentation head by distilling feature correspondences to segmentation correspondences. HP [46], on the other hand, identified hidden positives using the k-th nearest neighbor criterion to guide the contrastive objective. Our goal aligns with previous works in seeking pseudo-supervision by considering patch relationships. However, the key difference lies in our approach of considering data distribution to find trustworthy pseudo-supervision within the imperfect embedding space.

2.2 Self-supervised Representation Learning

Self-supervised representation learning has long been spotlighted for its effectiveness in providing a decent initialization point for various downstream tasks [4,22]. There are several prevalent approaches in this domain, including pretext tasks which learn the representation by reconstructing the original input from augmented images [9, 16, 20, 34, 40, 41, 57], relation-based approaches [3, 10, 21, 24, 39] and masked-modeling approaches [23, 47, 52, 53]. While masked-modeling approaches excel at preserving local context, they are often less efficient in learning discriminative representations [23, 25, 53]. Therefore, relation-based approaches [4, 56], particularly DINO [4], are popularly employed in the realm of USS [22, 46, 54]. Although the features from DINO are powerful in describing semantics for the whole image, their direct use for semantic segmentation proves effective due to the model being trained at the image-level, as shown in Fig 1b. In this regard, we have developed algorithms to complement the representation of DINO for the promotion of such features to reflect pixel-level semantics.



Fig. 2: Overall procedure of Progressive Proxy Anchor Propagation (PPAP). Our backbone consists of two branches: one for acquiring the training guidance, and the other for task adaptive finetuning. Specifically, the former feature extractor produces feature \mathbf{f} used to compute training guidance via trustworthy positive and ambiguity-excluded negative sets by PPAP, and its parameters are frozen for stable guidance. On the other hand, the latter branch is being finetuned with the training guidance to learn taskadaptive feature \mathbf{z} .

3 Method

3.1 Background and Overview

Recently, it has become mainstream to utilize the positive relationships among patches in training for Unsupervised Semantic Segmentation (USS) [22, 46, 54]. They exploited the patch-wise embeddings from a pretrained foundation model. However, we claim that they heavily relied on the similarity measured in imperfect embedding space for inferring patch-level training guidance. Instead, in this paper, we suggest the importance of considering the data distribution; since not all anchors are highly likely to be densely surrounded by semantically similar patch features in the embedding space, we aim to search for better spots to gather a sufficient number of trustworthy positive and negative samples.

The architecture of our method is illustrated in Fig 2. Following the recent works [22, 46, 54], our goal is to learn an appropriate projection function for the features extracted from a pretrained model suitable to the USS task. To achieve this, we define two streams using the pretrained ViT; for the first stream we keep all the blocks frozen to provide reliable supervision, and for the other stream we finetune the last block for adapting features to the semantic segmentation task. Given a mini-batch of images $\{\mathbf{x}_b\}_{b=1}^B$, the former stream computes pairs of $B \times H \times W$ patch features $\mathbf{f}_i \in \mathbb{R}^D$ where $H \times W$ is the number of patch features for an

image and D stands for the dimension of embedding space. On the other hand, the latter stream producing projected patch features $\mathbf{z}_i \in \mathbb{R}^D$ is being finetuned with the gathered positive and negative sets. Specifically, the process begins by determining the positive set \mathcal{P}_i with \mathbf{f}_i . Through the iterative process of positive gathering and proxy anchor relocation, we construct trustworthy positive set. Afterward, we determine the ambiguity-excluded negative set to train \mathbf{z}_i . In the following sections, we discuss the Progressive Proxy Anchor Propagation strategy to obtain trustworthy positive and ambiguity-excluded negative sets. in Sec. 3.2 and Sec. 3.3, respectively.

3.2 Progressive Proxy Anchor Propagation

Collecting a sufficient amount of trustworthy pseudo-supervision is a cumbersome task but crucial for the performance in USS [11].

To this end, we propose a Progressive Proxy Anchor Propagation algorithm to identify the reliable region for each anchor, where semantically similar samples to each anchor are densely located, as described in Fig. 2. The propagation process begins by forming an initial positive set comprising samples that are highly adjacent to each anchor. Subsequently, the algorithm employs an iterative process composed of two following steps: 1) relocate the proxy anchor towards more densely populated regions identified with the distribution of gathered positives, and 2) identify positive samples around the proxy anchor according to the expanded boundary. Note that the proxy anchor (*i.e.*, relocated anchor) provides a positive collection criterion on behalf of the anchor where its boundary for positive collection is proportional to the reliability of the proxy anchor's new position that is measured by the propagation degree of the proxy anchor. The proxy anchor position is considered more reliable if it does not move significantly, suggesting it is already surrounded by samples with similar semantics. This enables each anchor to gather numerous trustworthy positive samples.

Specifically, the initial positive set \mathcal{P}_i^0 of a given anchor \mathbf{f}_i is obtained by applying the initial positiveness criterion Φ^0 to gather as below:

$$\mathcal{P}_i^0 = \{ j \mid \mathbf{f}_i \cdot \mathbf{f}_j > \Phi_i^0, j \in \mathcal{B} \}, \ \forall_i \Phi_i^0 = \Phi^0, \tag{1}$$

where \mathcal{B} denotes the set containing all patch features within the mini-batch, *i.e.*, $|\mathcal{B}| = B \times H \times W$, and (·) refers to the similarity measure between two vectors (typically the cosine similarity). Here, the criterion Φ^0 is to decide whether all other patch features in the mini-batch are positive or not, based on the similarity threshold. This initial threshold is set to be big enough to make a tight criterion and is shared across all anchors. Such a tight boundary from a large initial criterion is for stable proxy anchor relocation since the samples with close proximity are more likely to be semantically similar.

Then, to propagate the proxy anchor toward the positives-dominant region, we derive the new proxy anchor position by averaging the collected positive set. This way, we take the distribution of the gathered positives into account. Formally, out of total T steps of relocation, we present the t-th relocated position of an anchor \mathbf{f}_i as \mathbf{v}_i^t , as follows:

$$\mathbf{v}_i^t = \frac{1}{|\mathcal{P}_i^{t-1}|} \sum_{j \in \mathcal{P}_i^{t-1}} \mathbf{f}_j.$$
⁽²⁾

To account for using the average points of the positive set, we posit that the close vicinity of the proxy anchor is highly likely to retain the same semantic. Thus, we claim that the center point of the gathered positives will move the proxy anchor closer to a dense region populated with semantically similar samples.

Furthermore, the positiveness criterion Φ_i^t should be reliability-adaptively adjusted. We determine the reliability of \mathbf{v}_i^t based on the similarity to the previous proxy anchor position \mathbf{v}_i^{t-1} since we assume that the proxy anchor point is converged to the center point of semantically similar patches if the scope of the propagation in a single step is limited (*i.e.*, high similarity between \mathbf{v}_i^{t-1} and \mathbf{v}_i^t). With such intuition, we revise the criterion Φ_i^t to be loosened when there is high reliability (\mathbf{v}_i^t and \mathbf{v}_i^{t-1} are in close proximity) on the position \mathbf{v}_i^t as follows:

$$\Phi_i^t = \Phi_i^{t-1} - \left(1 - (\mathbf{v}_i^{t-1} \cdot \mathbf{v}_i^t)) / \sigma_{\text{pos}},$$
(3)

where σ_{pos} is a coefficient used to prevent excessive reduction of the criterion. Note that $\mathbf{v}_i^0 = \mathbf{f}_i$. Consequently, the positive set at iteration t is expressed with the revised criterion Φ_i^t with the new proxy anchor point \mathbf{v}_i^t as follows:

$$\mathcal{P}_i^t = \{ j \mid \mathbf{v}_i^t \cdot \mathbf{f}_j > \Phi_i^t, j \in \mathcal{B} \}.$$
(4)

The process above is iteratively performed (Eq. 2 - Eq. 4) for T times to discover a reliable zone to sample the positives \mathcal{P}_i^T .

3.3 Ambiguity-excluded Negative Set

Along with the importance of gathering trustworthy positive sets for contrastive learning, preserving the reliability of the negative sets is another important factor [11]. Accordingly, we utilize the propagated proxy anchor \mathbf{v}_i^T as the base to compose the negative set to prevent the conflict to the positive set \mathcal{P}_i^T .

However, we point out the presence of an ambiguous zone for each anchor where positives and negatives are intermixed, making it unclear to categorize them exactly on one side. When the samples in such a zone are considered negatives, the model may face unwanted repulsion. Derived from such motivation, we additionally define an ambiguous set for each anchor that is neither included in the positive set nor the negative set, thereby excluding them from the learning process.

The method to establish the ambiguous set is similar to the process for positive set sampling: we update the criterion over the T steps of the proxy anchor propagation in an anchor-dependent manner and define the set according to this criterion. However, the key difference lies in how we determine the initial ambiguity criterion Ψ^0 . Unlike the boundary Φ for positive selection, we set 8 H.S. Seong et al.

 Ψ to a small value at the initial step to serve as a loose boundary since the vicinal areas of an initial anchor \mathbf{f}_i might not be reliable. This criterion is then progressively raised in the subsequent steps.

Given the initial anchor point \mathbf{v}_i^0 as \mathbf{f}_i and the *t*-th propagated proxy anchor point \mathbf{v}_i^t through proxy anchor propagation (Eq. 2), we progressively adjust the ambiguity criterion by:

$$\Psi_i^t = \Psi_i^{t-1} + \left(1 - (\mathbf{v}_i^{t-1} \cdot \mathbf{v}_i^t)\right) / \sigma_{\text{amb}},\tag{5}$$

where σ_{amb} is a coefficient to prevent excessive increase of the criterion and $\forall_i \Psi_i^0 = \Psi^0$. Through t steps, Ψ_i^t tightens the boundary if the position of the proxy anchor becomes densely surrounded by the positives. In other words, if the relocated proxy anchor is positioned in a densely populated area, there is less probability of having semantically alike samples outside the positive sampling region formed with Eq. 3.

After determining the ambiguity criterion through T steps, we then proceed to define the ambiguous set. Using the T-th relocated proxy anchor \mathbf{v}_i^T , the ambiguous set \mathcal{A}_i for the anchor \mathbf{f}_i is defined as follows:

$$\mathcal{A}_i = \{ j \mid (\mathbf{v}_i \cdot \mathbf{f}_j > \Psi_i^T) \land (\mathbf{v}_i \cdot \mathbf{f}_j < \Phi_i^T), j \in \mathcal{B} \}.$$
(6)

Finally, the negative set \mathcal{N} is organized as follows:

$$\mathcal{N}_i = \{ j \mid (j \notin \mathcal{P}_i) \land (j \notin \mathcal{A}_i), j \in \mathcal{B} \}.$$
(7)

3.4 Training Objective

Following existing works [13, 22, 27, 46] in USS, we utilize contrastive learning objective [10,28]. With the aim of distinguishing the semantically similar positive set \mathcal{P}_i^T and dissimilar negative set \mathcal{N}_i , the objective is expressed as:

$$L_{i}^{\text{con}} = \frac{-1}{|\mathcal{P}_{i}^{T}|} \sum_{p \in \mathcal{P}_{i}^{T}} \log \frac{\exp(\mathbf{z}_{i} \cdot \mathbf{z}_{p}/\tau)}{\sum_{n \in (\mathcal{N}_{i} \cup \mathcal{P}_{i}^{T})} \exp(\mathbf{z}_{i} \cdot \mathbf{z}_{n}/\tau)},$$
(8)

where τ is a temperature parameter, and \mathcal{P}_i^T and \mathcal{N}_i denote positive and negative sets for *i*-th anchor, respectively.

4 Experiments

4.1 Experimental Settings

Datasets. Following previous protocols [13,19,22,27,46], we evaluate our method on COCO-stuff [1] and Cityscapes [14], Potsdam-3, and ImageNet-S [19] datasets. Further details can be found in the Appendix. COCO-stuff is a dataset for scene understanding tasks, *e.g.*, semantic segmentation, detection, and image captioning, that consists of 172 classes. Among them, the COCO-stuff benchmark for

Method	Backbone	Unsupervised Acc. mIoU		Linear Acc. mIoU	
DC [2] MDC [2] IIC [27] PiCIE [13] PiCIE+H [13]	$\begin{array}{c} \mathrm{R18+FPN}\\ \mathrm{R18+FPN}\\ \mathrm{R18+FPN}\\ \mathrm{R18+FPN}\\ \mathrm{R18+FPN}\\ \mathrm{R18+FPN} \end{array}$	$ \begin{array}{r} 19.9 \\ 32.2 \\ 21.8 \\ 48.1 \\ 50.0 \\ \end{array} $	9.86.713.814.4		-13.3 8.4 13.9 14.8
DINO [4] TransFGU [54] STEGO [22] HP [46] PPAP (Ours)	ViT-S/8 ViT-S/8 ViT-S/8 ViT-S/8 ViT-S/8	28.7 52.7 48.3 57.2 59.0	11.3 17.5 24.5 24.6 27.2	68.6 - 74.4 75.6 76.9	33.9 - 38.3 42.7 46.3
DINO [4] STEGO [22] HP [46] PPAP (Ours)	ViT-S/16 ViT-S/16 ViT-S/16 ViT-S/16	22.0 52.5 54.5 62.9	8.0 23.7 24.3 26.5	50.3 70.6 74.1 76.0	18.1 34.5 39.1 43.3

 Table 1: Experimental results on COCO-stuff dataset.

USS utilizes 27 classes. Cityscapes is another large-scale dataset for scene understanding that consists of 30 classes captured across 50 different cities. Similarly to COCO-stuff, 27 subclasses are used for the benchmark. In addition, Potsdam-3 contains satellite images that are divided into 3 classes. Lastly, ImageNet-S is a large-scale dataset which has 1.2 million training images with 919 semantic classes.

Evaluation Protocols. For COCO-stuff, Cityscapes, and Potsdam-3 datasets, we adopt two evaluation methods: clustering (unsupervised) and linear probe [22, 46]. Clustering evaluates the alignment between the prediction and the ground truth with the Hungarian matching algorithm. On the other hand, the linear probe utilizes an additional fully connected layer for classification. For both evaluations, we apply the post-processing step using a Conditional Random Field (CRF) [33] to refine the predictions. Accuracy (Acc.) and mean Intersection over Union (mIoU) are used to measure the performances. For the evaluation on the ImageNet-S dataset, we adopt mIoU with distance matching *i.e.*, the k-nearest neighbors classifier with k=10, following the evaluation protocol from PASS [19].

4.2 Experimental Results

Quantitative Result. We compare the performances of our method with various baselines [2,13,22,27,46,54]. In Tab. 1, we observe that the recent works with the ViT backbone outperform the other ones, and among them, our approach demonstrates state-of-the-art performances across all metrics. In particular, our PPAP, equipped only with the sampling strategies for both the positive and negative samples, exceeds HP [46] that utilizes contrastive learning also with locality learning in a task-specific perspective.

10 H.S. Seong et al.

Table 2: Experimental results on Cityscapesdataset.

Method	Backbone	Unsu Acc.	pervised mIoU	Lir Acc.	near mIoU
MDC [2] UC [27]	R18+FPN B18+FPN	40.7 47.9	$7.1 \\ 6.4$	-	-
PiCIE [13]	R18+FPN	65.5	12.3	-	-
DINO [4]	ViT-S/8	34.5	10.9	84.6	22.8
TransFGU [54]	ViT-S/8	77.9	16.8	-	- 20.6
PPAP (Ours)	VII-5/8	80.1 82.0	18.4 19.6	90.8	31.5
DINO [4]	ViT-B/8	43.6	11.8	84.2	23.0
STEGO [22]	ViT-B/8	73.2	21.0	90.3	26.8
HP [46]	ViT-B/8	79.5	18.4	90.9	33.0
PPAP (Ours)	ViT-B/8	83.3	21.2	91.4	36.5

Table 3: Experimental results onPotsdam-3 dataset.

Method	Backbone	Unsuj Acc.	pervised mIoU
Rand. CNN [27]	VGG11	38.2	-
K-Means [43]	VGG11	45.7	-
SIFT [38]	VGG11	38.2	-
CP [15]	VGG11	49.6	-
CC [26]	VGG11	63.9	-
DeepCluster [2]	VGG11	41.7	-
IIC [27]	VGG11	65.1	-
DINO [4]	ViT-B/8	53.0	-
STEGO [22]	ViT-B/8	77.0	-
HP [46]	ViT-B/8	82.4	69.7
PPAP (Ours)	ViT-B/8	83.2	71.0

 Table 4: Experimental results on ImageNet-S

 validation set. †: reproduced performance.

Method	Backbone	IN-S	IN-S_{300}	IN-S_{50}
SwAV [3] PASS [19]	ResNet-50 ResNet-50	$15.1 \\ 15.6$	$22.4 \\ 25.1$	-
$DINO^{\dagger}$ [4] HP^{\dagger} [46] PPAP (Ours)	ViT-S/16 ViT-S/16 ViT-S/16	7.3 8.6 25.7	12.0 14.4 37.3	22.8 29.5 59.3

Table 5: Ablation study varying eachcomponent.

PF	PPAP		Unsupervised		
TPS	ANS	Acc.	mIoU		
-	-	49.1	22.9		
\checkmark	-	53.7	25.0		
-	\checkmark	55.5	25.1		
✓	✓	62.9	26.5		

In terms of the backbone, we achieve greater improvements with the ViT-S/16, which uses larger-sized patch features. We attribute these results to the robust property of our PPAP. Unlike PPAP, other methods are shown to yield better results with the small patch size (ViT-S/8 backbone) since the patches with the larger size are more likely to include a mixture of semantics. Yet, our proposed PPAP is guided to search for semantically similar patches to learn its prototypical proxy point and even disregard the patches that retain an ambiguous relationship with the given anchor. As a result, PPAP achieves promising results by a margin up to 15.41% and 9.05% compared to HP [46] in unsupervised Acc. and mIoU, respectively.

Performance comparison on the Cityscapes dataset is displayed in Tab. 2. Similar to the results on the COCO-stuff dataset, our proposed PPAP achieves new state-of-the-art results except in one case. These results further verify the applicability of our components to different datasets.

PPAP also shows improvement on the Potsdam-3 as shown in Tab. 3, but the difference is modest due to the dataset's limited three distinct semantic classes, well clustered by pretrained ViT features. However, our method demonstrates greater enhancements on datasets with more and less distinct semantic classes, not as effectively distinguished by the pretrained backbones.

On the ImageNet-S dataset, PPAP significantly outperforms existing methods, PASS [19] and HP [46], as shown in Tab. 4. We also conduct experiments



Fig. 3: Qualitative comparison results of PPAP (Ours) with STEGO and HP on the COCO-stuff dataset with DINO pretrained ViT-S/8 backbone.

on the subsets ImageNet- S_{300} and ImageNet- S_{50} , which contain 300 and 50 classes, respectively. The superior performance of PPAP on these datasets further demonstrates its scalability compared to existing methods.

Our PPAP algorithm shares similarities with k-means clustering in that both methods iteratively find and relocate points through averaging. However, there are three critical differences: 1) PPAP aims to discover patch-wise proxy anchor rather than having instances within a cluster share a single proxy. 2) PPAP enables a stable initial relocation process by using only the very close nearest neighbors of an anchor. 3) PPAP accurately determines the positiveness and negativeness of each anchor based on the degree of its relocation. In contrast, simply applying k-means clustering and assigning positive relationships to all features within each cluster leads to high recall but low precision in the positive set. For example, we observed that the precision of positives with k-means is only 22%, even when the number of classes is given as a prior. In comparison, PPAP achieves a precision of 72% on the COCO-stuff dataset using the DINO pretrained ViT-S/16 model without class prior.

Qualitative Result. We display our quantitative results in comparison to STEGO [22] and HP [46] in Fig. 3. Considering the complexity of the scenes, we plot simple to complicated scenes in order from left to right. To be brief, our PPAP shows consistent results to have fewer mispredicted pixels compared to the baselines. Particularly, our proposed method is robust to pixel-wise noises because each anchor is progressively propagated to search for reliable points that address the vulnerability of per-sample basis inference in the imperfect embedding space learned in an unsupervised manner at the image-level.

4.3 Ablation Study

We provide ablation studies to evaluate the individual components and key hyperparameters. The primary components under study are: 1) Trustworthy Pos-



Fig. 4: Ablation studies of various coefficients on three different datasets. Whereas the X-axis denotes the value of each hyperparameter, the Y-axis shows the performance.

itive Set (TPS) obtained based on PPAP and 2) Ambiguity-excluded Negative Set (ANS). Additionally, we examined the impact of hyperparameters, specifically: 1) coefficients $\sigma_{\rm pos}$ and $\sigma_{\rm amg}$ to regulate the criteria and 2) initial criteria Φ^0 and Ψ^0 .

Varying Components of PPAP. We perform an ablation study to assess the individual contributions of each component, which are presented in Tab. 5. These experiments are performed on the COCO-stuff dataset with ViT-S/16 backbone. For the baseline, we train the model using contrastive loss based on positives determined with the initial positiveness criterion \mathcal{P}^0 and negatives comprising all remaining samples. Incorporating the TPS yields improvements of 9.36% in unsupervised accuracy and 9.17% in mIoU. For the experiment on the third row, the ambiguous set for contrastive loss is defined after *T*-step propagation, while the positive set is defined with the initial positiveness criterion \mathcal{P}^0 . This brings 13.03% and 9.6% enhancements over the baseline. These results verify that both components significantly contribute to performance enhancement. Consequently, with both components combined, we observe a notable overall improvement of 28.11% in accuracy and 15.72% in mIoU.

Varying σ_{pos} and σ_{amb} . In Fig 4a and 4b, we carry out an ablation study on the coefficient σ , which is crucial for controlling the degree of reduction and increase in positiveness and ambiguity criteria, respectively. A smaller σ_{pos} leads to a more intensive reduction in the positiveness criterion, while a bigger σ_{pos} results in a more gradual reduction. Such property is reflected in the outcomes presented in Fig 4a. A smaller σ_{pos} tends to lower the performance due to FP erroneously included in the positive set. On the other hand, moderately larger σ_{pos} helps to mitigate the aforementioned problem. σ_{amb} follows the same principle: a smaller value results in a more substantial increase of the criterion, escalating the possibility of erroneously repelling FP in the negative set. And a bigger value leads to a more moderate increase. A proper σ can prevent both the excessive changes in the criteria. We note that it is advisable to choose values around 3 to ensure stability.

Table 6: Comparing both TP in positive set (\mathcal{P}) and FP in negative set (\mathcal{N}) between HP and Ours on 3 datasets. Numbers and percentages are the results averaged from subsamples of each dataset. S/8 and B/8 indicate ViT-S/8 and ViT-B/8 backbone, respectively.

	Cityscap HP [46]	es $(S/8)$	HP [46]	uff27 $(S/8)$ Ours	Potsdam- HP [46]	-3 (B/8) Ours
Number of positives in \mathcal{P}	78	1838	142	288	$ \begin{array}{c ccccccccccccccccccccccccccccccccccc$	194
% of TP in \mathcal{P}	88.83	90.18	81.98	87.15		88.62
Number of negatives in \mathcal{N}	50K	43K	50K	46K	50K	30K
% of FP in \mathcal{N}	21.33	14.22	7.34	5.92	34.08	29.01

Varying Φ^0 and Ψ^0 . Ablation studies for varying Φ^0 and Ψ^0 are shown in Fig. 4c and 4d. These parameters act as the initial criteria for selecting positive and ambiguous samples and serve as key hyperparameters in our method. For Φ^0 , as samples in close vicinity to the anchor are highly likely to share the same semantics, setting its initial value low incurs the existence of FP in the positive set. Still, setting the Φ^0 too high results in selecting only a few positives that the anchor relocation is implemented only within its very close proximity. Regarding Ψ^0 , setting Ψ^0 too small significantly reduces the size of the negative set and leads to a shortage of hard negative samples in the negative set, while too large Ψ^0 increases FP in the negative set. We found that values of 0.55 for Φ^0 and 0.15 for Ψ^0 consistently perform well across all datasets and models. Note that we also found that higher Φ^0 requires more propagation steps (*i.e.*, bigger T).

4.4 Trustworthiness of Positive and Negative Sets.

If excessive FP and FN are present in contrastive loss, the model may face undesired attraction and repulsion, respectively. We contend that mitigating these issues can enhance the trustworthiness of contrastive learning. To verify the robust trustworthiness of our method, previously, we conducted a precision comparison of the positive sets between ours and HP [46] on Fig. 1b. Notably, ours not only collects more positive samples than HP but also achieves a higher precision (*i.e.*, ratio of TP). For further demonstration, we present the ratios of TP in the positive set (\mathcal{P}) and FP in the negative set (\mathcal{N}) for three datasets in Tab. 6. As shown, ours got a higher ratio of TP in \mathcal{P} and lower ratio of FPin \mathcal{N} than HP, even with a much larger size of the positive set. For example, we can observe that Ours got 23× more positives with a higher ratio of TP on Cityscapes dataset with ViT-S/8 backbone. Likewise, we ensure the trustworthiness of contrastive learning under more reliable positive and negative sets.

4.5 Visualization of Patches in Positive Set

We illustrate how HP [46] and our proposed PPAP organize the positive samples for each anchor through visualizations in Fig. 5. As depicted throughout the visualizations, there is a tendency for our method to collect more positive



Fig. 5: Comparison of gathered positives between HP and PPAP (Ours) with the visualizations. In all examples, blue boxes indicate the selected anchor patch and the red boxes denote the patches that are considered as positive to an anchor. In (d), yellow dotted circles exist to highlight the region where the FP are detected.

patches compared to the baseline. On top of that, we can also observe that positive samples that are semantically identical are progressively obtained through the propagation steps in (a), (b), and (c). Lastly, along with the numerically measured difference in the precision of the positive sets, we find that falsely detected positives by the baseline in (d) are not considered positive in ours. To account for such a phenomenon, we claim that the presence of ambiguous zones enables the filtering of the hard negative samples.

5 Conclusion

To tackle the challenge of USS, previous approaches have primarily relied on exploiting patch relationships to guide the training process. In this work, we extend this mainstream to ensure the reliability of the gathered guidance. Specifically, we consider the data distribution around the anchor to identify densely crowded regions containing samples with similar semantics. By relocating the proxy anchor to these regions, we expect it to be surrounded by trustworthy positives, creating a large positive set with high precision. In addition, we address instancewise ambiguous zones where samples with similar and dissimilar semantics coexist. By excluding samples from these regions during training, we aim to eliminate FP in the negative set, preventing unstable training. Our state-of-the-art results verify the importance of ensuring the reliability of the supervision in USS.

Acknowledgements. This work was supported in part by MSIT&KNPA/KIPoT (Police Lab 2.0, No. 210121M06), MSIT/IITP (No. 2022-0-00680, 2019-0-00421, 2020-0-01821, RS-2024-00437102), and SEMES-SKKU collaboration funded by SEMES.

References

- Caesar, H., Uijlings, J., Ferrari, V.: Coco-stuff: Thing and stuff classes in context. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1209–1218 (2018)
- Caron, M., Bojanowski, P., Joulin, A., Douze, M.: Deep clustering for unsupervised learning of visual features. In: Proceedings of the European conference on computer vision (ECCV). pp. 132–149 (2018)
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. Advances in Neural Information Processing Systems 33, 9912–9924 (2020)
- Caron, M., Touvron, H., Misra, I., Jégou, H., Mairal, J., Bojanowski, P., Joulin, A.: Emerging properties in self-supervised vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9650–9660 (2021)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. arXiv preprint arXiv:1412.7062 (2014)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Semantic image segmentation with deep convolutional nets and fully connected crfs. ICLR (2015)
- Chen, L.C., Papandreou, G., Kokkinos, I., Murphy, K., Yuille, A.L.: Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. IEEE transactions on pattern analysis and machine intelligence 40(4), 834–848 (2017)
- Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: Proceedings of the European conference on computer vision (ECCV). pp. 801–818 (2018)
- Chen, P., Liu, S., Jia, J.: Jigsaw clustering for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (2021)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, T.S., Hung, W.C., Tseng, H.Y., Chien, S.Y., Yang, M.H.: Incremental false negative detection for contrastive learning. arXiv preprint arXiv:2106.03719 (2021)
- Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1290–1299 (2022)
- Cho, J.H., Mall, U., Bala, K., Hariharan, B.: Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16794– 16804 (2021)

- 16 H.S. Seong et al.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
- Doersch, C., Gupta, A., Efros, A.A.: Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE international conference on computer vision. pp. 1422–1430 (2015)
- Dosovitskiy, A., Springenberg, J.T., Riedmiller, M., Brox, T.: Discriminative unsupervised feature learning with convolutional neural networks. In: Advances in Neural Information Processing Systems. Citeseer (2014)
- Farabet, C., Couprie, C., Najman, L., LeCun, Y.: Learning hierarchical features for scene labeling. IEEE transactions on pattern analysis and machine intelligence 35(8), 1915–1929 (2012)
- Fu, J., Liu, J., Tian, H., Li, Y., Bao, Y., Fang, Z., Lu, H.: Dual attention network for scene segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3146–3154 (2019)
- Gao, S., Li, Z.Y., Yang, M.H., Cheng, M.M., Han, J., Torr, P.: Large-scale unsupervised semantic segmentation. IEEE transactions on pattern analysis and machine intelligence 45(6), 7457–7476 (2022)
- Gidaris, S., Singh, P., Komodakis, N.: Unsupervised representation learning by predicting image rotations. In: International Conference on Learning Representations (2018), https://openreview.net/forum?id=S1v4N210-
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P., Buchatskaya, E., Doersch, C., Avila Pires, B., Guo, Z., Gheshlaghi Azar, M., et al.: Bootstrap your own latent-a new approach to self-supervised learning. Advances in neural information processing systems 33, 21271–21284 (2020)
- Hamilton, M., Zhang, Z., Hariharan, B., Snavely, N., Freeman, W.T.: Unsupervised semantic segmentation by distilling feature correspondences. In: International Conference on Learning Representations (2022), https://openreview.net/forum?id= SaK06z6H10c
- He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
- Huang, Z., Jin, X., Lu, C., Hou, Q., Cheng, M.M., Fu, D., Shen, X., Feng, J.: Contrastive masked autoencoders are stronger vision learners. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023)
- 26. Isola, P., Zoran, D., Krishnan, D., Adelson, E.H.: Learning visual groups from co-occurrences in space and time. arXiv preprint arXiv:1511.06811 (2015)
- Ji, X., Henriques, J.F., Vedaldi, A.: Invariant information clustering for unsupervised image classification and segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9865–9874 (2019)
- Khosla, P., Teterwak, P., Wang, C., Sarna, A., Tian, Y., Isola, P., Maschinot, A., Liu, C., Krishnan, D.: Supervised contrastive learning. Advances in Neural Information Processing Systems 33, 18661–18673 (2020)
- Kim, C., Han, W., Ju, D., Hwang, S.J.: Eagle: Eigen aggregation learning for object-centric unsupervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3523– 3533 (2024)

- Kim, J., Shim, K., Lee, I., Shim, B.: Expand-and-quantize: Unsupervised semantic segmentation using high-dimensional space and product quantization. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 2768–2776 (2024)
- Kim, J., Lee, B.K., Ro, Y.M.: Causal unsupervised semantic segmentation. arXiv preprint arXiv:2310.07379 (2023)
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (2023)
- Krähenbühl, P., Koltun, V.: Efficient inference in fully connected crfs with gaussian edge potentials. Advances in neural information processing systems 24 (2011)
- 34. Larsson, G., Maire, M., Shakhnarovich, G.: Learning representations for automatic colorization. In: European conference on computer vision. Springer (2016)
- Li, Z., Wang, W., Xie, E., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P., Lu, T.: Panoptic segformer: Delving deeper into panoptic segmentation with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1280–1289 (2022)
- Lin, G., Milan, A., Shen, C., Reid, I.: Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1925–1934 (2017)
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3431–3440 (2015)
- Lowe, D.G.: Object recognition from local scale-invariant features. In: Proceedings of the seventh IEEE international conference on computer vision. vol. 2, pp. 1150– 1157. Ieee (1999)
- Misra, I., Maaten, L.v.d.: Self-supervised learning of pretext-invariant representations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6707–6717 (2020)
- 40. Noroozi, M., Favaro, P.: Unsupervised learning of visual representations by solving jigsaw puzzles. In: European conference on computer vision. Springer (2016)
- 41. Noroozi, M., Vinjimoor, A., Favaro, P., Pirsiavash, H.: Boosting self-supervised learning via knowledge transfer. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)
- 42. Park, S., Lee, S., Hyun, S., Seong, H.S., Heo, J.P.: Task-disruptive background suppression for few-shot segmentation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 4442–4449 (2024)
- 43. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. the Journal of machine Learning research 12, 2825–2830 (2011)
- 44. Pinheiro, P., Collobert, R.: Recurrent convolutional neural networks for scene labeling. In: International conference on machine learning. pp. 82–90. PMLR (2014)
- Seitzer, M., Horn, M., Zadaianchuk, A., Zietlow, D., Xiao, T., Simon-Gabriel, C.J., He, T., Zhang, Z., Schölkopf, B., Brox, T., et al.: Bridging the gap to real-world object-centric learning (2023)
- Seong, H.S., Moon, W., Lee, S., Heo, J.P.: Leveraging hidden positives for unsupervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 19540–19549 (2023)

- 18 H.S. Seong et al.
- Shi, Y., Siddharth, N., Torr, P., Kosiorek, A.R.: Adversarial masking for selfsupervised learning. In: International Conference on Machine Learning. pp. 20026– 20040. PMLR (2022)
- Sick, L., Engel, D., Hermosilla, P., Ropinski, T.: Unsupervised semantic segmentation through depth-guided feature correlation and sampling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3637– 3646 (2024)
- Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7262–7272 (2021)
- Wikipedia: False positives and false negatives Wikipedia, the free encyclopedia. http://en.wikipedia.org/w/index.php?title=False%20positives%20and% 20false%20negatives&oldid=1205293928 (2024), [Online; accessed 06-March-2024]
- Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. Advances in Neural Information Processing Systems 34, 12077–12090 (2021)
- 52. Xie, J., Li, W., Zhan, X., Liu, Z., Ong, Y.S., Loy, C.C.: Masked frequency modeling for self-supervised visual pre-training. In: The Eleventh International Conference on Learning Representations (2023), https://openreview.net/forum?id=9umxtNPx5E
- 53. Xie, Z., Zhang, Z., Cao, Y., Lin, Y., Bao, J., Yao, Z., Dai, Q., Hu, H.: Simmim: A simple framework for masked image modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9653–9663 (2022)
- Yin, Z., Wang, P., Wang, F., Xu, X., Zhang, H., Li, H., Jin, R.: Transfgu: a topdown approach to fine-grained unsupervised semantic segmentation. In: European Conference on Computer Vision. pp. 73–89. Springer (2022)
- Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2015)
- Yun, S., Lee, H., Kim, J., Shin, J.: Patch-level representation learning for selfsupervised vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8354–8363 (2022)
- 57. Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European conference on computer vision. Springer (2016)
- Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017)