

ACSeg: Adaptive Conceptualization for Unsupervised Semantic Segmentation

Kehan Li^{1,3} Zhennan Wang² Zesen Cheng^{1,3} Runyi Yu^{1,3} Yian Zhao⁵ Guoli Song²
Chang Liu⁴ Li Yuan^{1,2,3*} Jie Chen^{1,2,3*}

¹ School of Electronic and Computer Engineering, Peking University, Shenzhen, China ² Peng Cheng Laboratory, Shenzhen, China

³ AI for Science (AI4S)-Preferred Program, Peking University Shenzhen Graduate School, Shenzhen, China

⁴ Department of Automation and BNRist, Tsinghua University, Beijing, China ⁵ Dalian University of Technology

Abstract

Recently, self-supervised large-scale visual pre-training models have shown great promise in representing pixel-level semantic relationships, significantly promoting the development of unsupervised dense prediction tasks, e.g., unsupervised semantic segmentation (USS). The extracted relationship among pixel-level representations typically contains rich class-aware information that semantically identical pixel embeddings in the representation space gather together to form sophisticated concepts. However, leveraging the learned models to ascertain semantically consistent pixel groups or regions in the image is non-trivial since over/under-clustering overwhelms the conceptualization procedure under various semantic distributions of different images. In this work, we investigate the pixel-level semantic aggregation in self-supervised ViT pre-trained models as image Segmentation and propose the Adaptive Conceptualization approach for USS, termed **ACSeg**. Concretely, we explicitly encode concepts into learnable prototypes and design the **Adaptive Concept Generator (ACG)**, which adaptively maps these prototypes to informative concepts for each image. Meanwhile, considering the scene complexity of different images, we propose the modularity loss to optimize ACG independent of the concept number based on estimating the intensity of pixel pairs belonging to the same concept. Finally, we turn the USS task into classifying the discovered concepts in an unsupervised manner. Extensive experiments with state-of-the-art results demonstrate the effectiveness of the proposed ACSeg.

1. Introduction

Semantic segmentation is one of the primary tasks in computer vision, which has been widely used in many domains, such as autonomous driving [7, 14] and medical

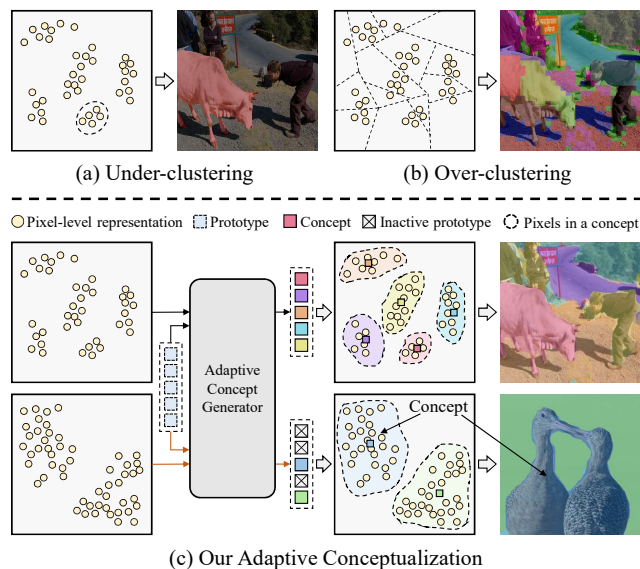


Figure 1. Comparison between existing methods and our adaptive conceptualization on finding underlying “concepts” in the pixel-level representations produced by a pre-trained model. While under-clustering just focuses on a single object and over-clustering splits objects, our adaptive conceptualization processes different images adaptively through updating the initialized prototypes with the representations for each image.

imaging [12, 24, 41]. With the development of deep learning and the increasing amount of data [7, 11, 28, 57], uplifting performance has been achieved on this task by optimizing deep neural networks with pixel-level annotations [29]. However, large-scale pixel-level annotations are expensive and laborious to obtain. Different kinds of weak supervision have been explored to achieve label efficiency [37], e.g., image-level [1, 49], scribble-level [27], and box-level supervision [35]. More than this, some methods also achieve semantic segmentation without relying on any labels [19, 20], namely unsupervised semantic segmentation (USS).

Early approaches for USS are based on pixel-level self-supervised representation learning by introducing cross-

*Corresponding author. Project page: <https://lklh.github.io/ACSeg>.

view consistency [6,20], edge detection [19,56], or saliency prior [43]. Recently, the self-supervised ViT [4] provides a new paradigm for USS due to its property of containing semantic information in pixel-level representations. We make it more intuitive through Figure 1, which shows that in the representation space of an image, the pixel-level representations produced by the self-supervised ViT contain underlying clusters. When projecting these clusters into the image, they become semantically consistent groups of pixels or regions representing “concepts”.

In this work, we aim to achieve USS by accurately extracting and classifying these “concepts” in the pixel representation space of each image. Unlike the previous attempts which only consider foreground-background partition [40,44,48] or divide each image into a fixed number of clusters [18,32], we argue that it is crucial to consider different images distinguishably due to the complexity of various scenarios (Figure 1). We thus propose the **Adaptive Conceptualization for unsupervised semantic Segmentation (ACSeg)**, a framework that finds these underlying concepts adaptively for each image and achieves USS by classifying the discovered concepts in an unsupervised manner.

To achieve conceptualization, we explicitly encode concepts to learnable prototypes and adaptively update them for different images by a network, as shown in Figure 2. This network, named as **Adaptive Concept Generator (ACG)**, is implemented by iteratively applying scaled dot-product attention [45] on the prototypes and pixel-level representations in the image to be processed. Through such a structure, the ACG learns to project the initial prototypes to the concept in the representation space depending on the input pixel-level representations. Then the concepts are explicitly presented in the image as different regions by assigning each pixel to the nearest concept in the representation space. The ACG is end-to-end optimized without any annotations by the proposed modularity loss. Specifically, we construct an affinity graph on the pixel-level representations and use the connection relationship of two pixels in the affinity graph to adjust the strength of assigning two pixels to the same concept, motivated by the modularity [34].

As the main part of ACSeg, the ACG achieves precise conceptualization for different images due to its adaptiveness, which is reflected in two aspect: Firstly, it can adaptively operate on pixel-level representations of different images thanks to the dynamic update structure. Secondly, the training objective does not enforce the number of concepts, resulting in adaptive number of concepts for different images. With these properties, we get accurate partition for images with different scene complexity via the concepts produced by the ACG, as shown in Figure 1(c). Therefore, in ACSeg, the semantic segmentation of an image can finally be achieved by matting the corresponding regions in the image and classifying them with the help of powerful




Figure 2. **Intuitive explanation for the basic idea of the ACG.** The concepts are explicitly encoded to learnable prototypes and dynamically updated according to the input pixel-level representations. After update, the pixels are assigned to the nearest concept in the representation space.

image-level pre-trained models.

For evaluation, we apply ACSeg on commonly used semantic segmentation datasets, including PASCAL VOC 2012 [11] and COCO-Stuff [20,28]. The experimental results show that the proposed ACSeg surpasses previous methods on different settings of unsupervised semantic segmentation tasks and achieves state-of-the-art performance on the PASCAL VOC 2012 unsupervised semantic segmentation benchmark without post-processing and re-training. Moreover, the visualization of the pixel-level representations and the concepts shows that the ACG is applicable for decomposing images with various scene complexity. Since the ACG is fast to converge without learning new representations and the concept classifier is employed in a zero-shot manner, we draw the proposed ACSeg as a generalizable method which is easy to modify and adapt to a wide range of unsupervised image understanding.

2. Related Works

Vison Transformer. Transformer, a model mainly based on self-attention mechanism, is widely used in natural language processing [3,8] and cross-modal understanding [21,25,26,36]. Vison Transformer (ViT) [9] is the first pure visual transformer model to process images. Recently, Caron *et al.* [4] propose self-distillation with **no** labels (DINO) to train the ViT, and found a property that its features contain explicit information about the segmentation of an image. Based on DINO, some previous studies [15,32,40,44,48,53] successfully demonstrate extending this property to unsupervised dense prediction tasks.

Unsupervised Semantic Segmentation. With the development of self-supervised and unsupervised learning, unsupervised methods for semantic segmentation task start to emerge. Among them, some methods focus on pixel-level self-supervised representation learning by introducing cross-view consistency [6,20,22,50,51,55,59], visual prior [19,43,56], and continuity of video frames [2]. In contrast, Zadaianchuk *et al.* [54] adopt pre-trained object-centric representations and cluster them to segment objects. Other methods exploit pixel-level knowledge of

pre-trained generative models [31] or self-supervised pre-trained convolutional neural network [18, 47]. Recently, self-supervised ViTs trained with DINO has recently been explored for unsupervised dense prediction tasks due to the ability of representing pixel-level semantic relationships. For semantic segmentation, Hamilton *et al.* [15] train a segmentation head by distilling the feature correspondences, which further encourages pixel features to form compact clusters and learn better pixel-level representations. TransFGU [53] obtains semantic segmentation in a top-down manner by extracting class activate maps from DINO models. Some approaches use the representations from DINO to segment images into regions. Melas *et al.* [32] adopt spectral decomposition on the affinity graph to discover meaningful parts in an image and implement semantic segmentation of an image. MaskDistill [44] uses some hand-made rules based on pixel-level representations to find the salient region in an image. In contrast, we aim at better extracting underlying concepts among the representations from DINO in an image by tackling the over/under-clustering problem.

Semantic Segmentation with Text. Vision-language pre-training models enable learning without annotations or zero-shot transfer on vision task [36]. For semantic segmentation, MaskCLIP [58] modifies the visual encoder of CLIP [36] and applies the text-based classifier on pixel level. Xu *et al.* [52] propose GroupViT, a hierarchical grouping vision transformer, and train it with image-to-text contrastive loss. Finally, the semantic segmentation results can be obtained by the grouping result and text embeddings. ReCo [39] leverages the retrieval abilities of CLIP and the robust correspondences offered by modern image representations to co-segment entities. Shin *et al.* [38] use CLIP to construct category-specific images and produce pseudo-label with a category-agnostic salient object detector bootstrapped from DINO. In this paper, we also show that our method can be combined with recent vision-language pre-trained model to perform semantic segmentation with only text-image supervision.

3. The Proposed ACSeg

In this section, we describe the proposed method for USS in detail, starting from the whole framework.

3.1. Overall Approach

Figure 3 illustrates the overall structure of ACSeg. Starting with an image, we first apply a self-supervised ViT to generate pixel-level representations. As mentioned above, these representations contain underlying concepts, which represents meaningful groups or regions of pixels. The Adaptive Concept Generator (ACG) is designed to output the concepts explicitly. Specifically, the ACG takes a series of learnable prototypes as input and iteratively updates

them by interacting with the pixel-level representations, resulting in adaptive concept representations for each image. Finally, the concepts are explicitly represented by pixel groups, which are obtained by assigning each pixel to the nearest concept in the representation space.

For optimization, we propose a novel loss function called modularity loss to train the ACG without any annotations. Intuitively, the modularity loss works on pixel pairs. We construct an affinity graph taking the pixel-level representations as vertices and their cosine similarity as edges. The modularity loss calculates the intensity of two pixels belonging to the same concept using the metric defined in modularity [34], thus adjusting the concept representations. At last, the concept classifier assigns each concept to a pre-defined category to obtain per-pixel class prediction, *i.e.*, semantic segmentation of an image. We introduce the details of each component in ACSeg in the following sections.

3.2. Adaptive Concept Generator

The role of ACG is to map the initial prototypes to the concept representations in each image. Since the concept representations are different in different images and depend on the pixel-level representations of the image, we introduce the scaled dot-product attention [45] to iteratively update the prototypes according to the pixel-level representations. Specifically, we first apply cross-attention taking the prototypes as the *query* and the pixel-level representations as the *key* and *value*. Let $\mathbf{C}^l \in \mathbb{R}^{k \times d}$ denote k prototypes after l -th update and $\mathbf{X} \in \mathbb{R}^{n \times d}$ denote n pixel-level representations from an image, the cross attention can be formulated as

$$\bar{\mathbf{C}}^l = \text{Softmax}\left(\frac{\mathbf{C}^{l-1}\mathbf{W}_q(\mathbf{X}\mathbf{W}_k)^T}{\sqrt{d}}\right)(\mathbf{X}\mathbf{W}_v), \quad (1)$$

$$\mathbf{C}^l = \mathbf{C}^{l-1} + \bar{\mathbf{C}}^l\mathbf{W}_o, \quad (2)$$

where $\mathbf{W}_q, \mathbf{W}_k, \mathbf{W}_v, \mathbf{W}_o \in \mathbb{R}^{d \times d}$ are learnable linear projections. The cross-attention updates prototypes adaptively with the pixel-level representations, which makes it possible to generate concepts adaptively for different images.

After that, self-attention is used to model the connections for different concepts. Formulaically, it can be expressed as

$$\bar{\mathbf{C}}^l = \text{Softmax}\left(\frac{\mathbf{C}^{l-1}\mathbf{W}_q(\mathbf{C}^{l-1}\mathbf{W}_k)^T}{\sqrt{d}}\right)(\mathbf{C}^{l-1}\mathbf{W}_v), \quad (3)$$

$$\mathbf{C}^l = \mathbf{C}^{l-1} + \bar{\mathbf{C}}^l\mathbf{W}_o. \quad (4)$$

The self-attention updates each prototype with other prototypes and makes it aware of the presence of others, for better adjusting their relative positions in the embedding space.

The ACG consists of N update steps, and each update step is made up of cross-attention, self-attention, and

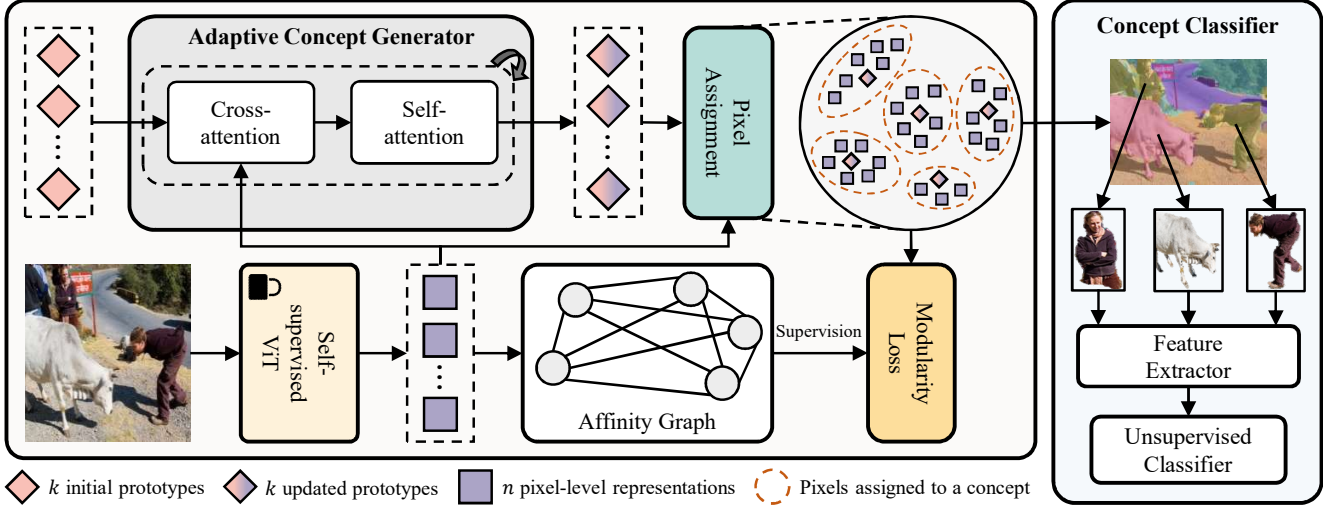


Figure 3. **Illustration of the proposed ACSeg.** For an image, we first use a self-supervised ViT to extract pixel-level representations, which imply the semantic relationship of pixels. The Adaptive Concept Generator (ACG) dynamically updates the initial prototypes to the underlying concepts in the representation space through scaled dot-product attention. Finally, the assignment of pixels is produced by the cosine similarity between pixel-level representations and the concepts, and the modularity loss is used to optimize the ACG. At last, the concept classifier is used to assign each concept to a pre-defined category thus obtain semantic segmentation of an image.

a Feed-forward Network (FFN) [45]. With the attention mechanism, the ACG can learn the map from initial prototypes to concepts adaptively for different images. For implementation, we adopt multi-head attention, layer normalization, and residual connection after the attention operation and the FFN, following the transformer [45].

3.3. Pixel Assignment

After ACG, each image has its own concepts. Abstractly, each concept is a vector in the representation space, approximately the average of a group of gathered pixels. Concretely, a concept consists of pixels with the same semantics. This abstract-to-concrete transformation is achieved by assigning each pixel to a concept in ACSeg.

We first get a soft assignment for each pixel by calculating the cosine similarity with the concepts in the same image

$$S_{i,j} = \cos \langle \mathbf{x}_i, \mathbf{c}_j \rangle, \quad (5)$$

where $S \in \mathbb{R}^{n \times k}$ is the assignment matrix, $\mathbf{x}_i = \mathbf{X}_{i,:}$ is the i -th pixel embedding and $\mathbf{c}_j = \mathbf{C}_{j,:}$ is the j -th concept. The soft assignment is differentiable and is used to optimize the network when training, which is described in Section 3.4. We assign each pixel to a definite concept during inference by the maximum similarity

$$a_i = \operatorname{argmax}_j \cos \langle \mathbf{x}_i, \mathbf{c}_j \rangle. \quad (6)$$

By doing that, an image is segmented into m regions. Each region is identified by a concept and consists of pixels assigned to this concept. It is worth noting that m can be

different for different images because the assignment is obtained by argmax operation which do not guarantee that every concept is assigned at least once. Due to the adaptive nature of this assignment and the adaptive generating of concepts for each image, we name the network Adaptive Concept Generator.

3.4. Modularity Loss

For training the ACG, we design a loss function based on the idea of estimating the intensity of assigning two pixels to the same concept. To achieve this goal, we introduce modularity [34], which is commonly used in community detection. The modularity is built upon a graph, thus we first construct a fully connected undirected affinity graph for pixels from an image by taking them as vertices. The weight of edge between two pixels which represents their affinity is calculated by the cosine similarity of them

$$A_{i,j} = \max(0, \cos \langle \mathbf{x}_i, \mathbf{x}_j \rangle). \quad (7)$$

Here we truncate the value to a minimum of zero to avoid negative values in calculation. Given two vertices i, j , following the modularity, we estimate the intensity w_{ij} of assigning them to the same concept by

$$w_{ij} = A_{i,j} - \frac{k_i \cdot k_j}{2m}, \quad (8)$$

where $k_i = \sum_j A_{i,j}$ is the sum of edges that are connected to vertex i and $2m = \sum_{i,j} A_{i,j}$ is the sum of all edges in the graph. For intuition, w_{ij} reflects the intensity of dividing pixels i and j into the same cluster via comparing the actual

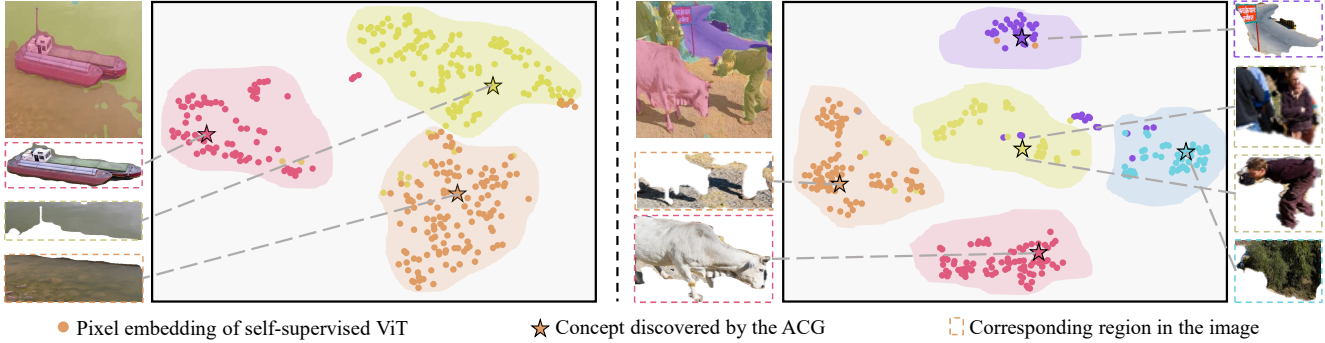


Figure 4. **t-SNE visualization of the pixel-level representations produced by self-supervised ViT and the corresponding concepts discovered by the ACG.** We mark the concepts found by the ACG in different colors. Based on the pixel-level representations, the ACG precisely finds the underlying concepts adaptively for different images with different scene complexity.

situation and the random situation. When preserving the degrees of vertices in the graph but connecting vertices randomly, the probability of an edge existing between vertices i and j is $k_i \cdot \frac{k_j}{2m}$. Finally, w_{ij} reveals the possible existence of clusters by the comparison between the actual density of edges (A_{ij}) and the expected density when vertices are attached randomly ($\frac{k_i \cdot k_j}{2m}$). We refer readers to [13,34] for the detailed derivation.

For constructing a differentiable function, we now go back to the representation space and define the degree of two pixels belong to the same concept considering the updated prototypes

$$\delta(i, j) = \max_c \bar{S}_{i,c} \cdot \bar{S}_{j,c}, \quad \bar{S} = \max(0, S), \quad (9)$$

where S is the soft assignment in equation 5, which makes the loss function differentiable. We ignore values less than zero in S to avoid the case where both $S_{i,c}$ and $S_{j,c}$ are negative. Since the pixel-level representations are fixed and only the prototypes are updated, the update of a prototype will be ambiguous if we take unrelated pixel pairs into account. Therefore, we only choose one related prototype c for each pair by $\max_c \bar{S}_{i,c} \cdot \bar{S}_{j,c}$ and calculating $\delta(i, j)$ in the same way. With w_{ij} and $\delta(i, j)$, the modularity loss for pixels in an image can be formulated as

$$\begin{aligned} \mathcal{L} &= -\frac{1}{2m} \sum_{i,j} w_{ij} \delta(i, j) \\ &= -\frac{1}{2m} \sum_{i,j} \left(A_{i,j} - \frac{k_i \cdot k_j}{2m} \right) \delta(i, j). \end{aligned} \quad (10)$$

Because we optimize the ACG with mini-batches, we use the sum of all edges $2m$ as the normalization factor for different images following the modularity.

Overall, the modularity loss adjusts the similarity of the prototypes to different pixel pairs according to the estimated intensity of i, j belonging to the same concept and then determines the position of updated prototypes in the embedding space. Meanwhile, the max function in equation 9

adaptively chooses different cluster centers, which achieve the dynamic number of cluster centers for different images. We summarize the reasons why the modularity loss is suitable for optimizing ACG as follows: Firstly, no hyperparameters are required, making it robust to process images with scenarios of different complexity. Secondly, the minimum of the modularity loss does not depend on the number of concepts, allowing the network to detect different counts of concepts in different images. Through optimizing the modularity loss, the ACG finally learns to adaptively predict concepts in different images.

3.5. Concept Classifier

In ACSeg, we finally obtain semantic segmentation of images by classifying the concepts produced by the ACG. We first discuss the classifier for the special “background” class, which is hard to describe when the labels are unavailable. For instance, the background class actually contains a lot of sub-classes such as water, sky, land, wall, etc. However, under the unsupervised setting, it is not known what or how many these sub-classes are.

We identify the background classes based on the attention of a self-supervised ViT, which is used when producing pixel-level representations so that no additional models are required. Specifically, we first get a foreground score for each region by summing the attention values of pixels within it. The attention values of pixels are from the last self-attention layer of the self-supervised ViT, taking the minimum of each attention head based on the idea that a pixel is likely to be foreground when it appears in at least one attention head. With the foreground score of regions in an image, we cluster them into two categories and classify the cluster with a smaller score to background.

For the concepts belonging to foreground, the classification can be achieved by multiple ways. One of them is to get a region-level representation for each concept by matting it from the original image and exploiting the ability to extract a discriminative image-level representation of

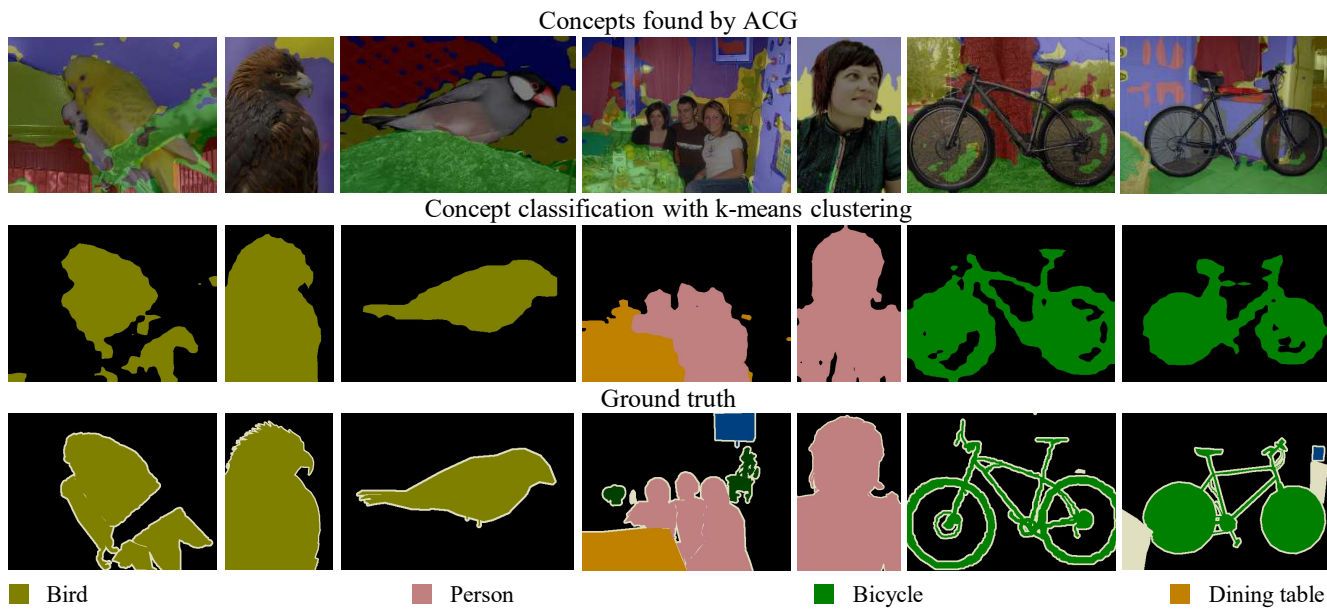


Figure 5. **Qualitative results on PASCAL VOC 2012 dataset.** We show the conceptualization results of ACG, semantic segmentation results based on k-means clustering and the ground truth in sequence. The proposed ACSeg can find different semantics within an image and realize precise semantic segmentation with image-level embedding from the self-supervised ViT.

the self-supervised ViT to get a region-level representation for each concept. With these discriminative region-level representations, we apply the k-means clustering algorithm and k -NN classifier on these representations to determine the class of each concept. Besides, the recent progress of vision-language pre-trained models [36] makes it possible to achieve unsupervised classification with the guidance of texts. Therefore, we design a text-based classifier for the obtained concepts. Specifically, we first obtain pixel-level representations using the visual encoder of the pre-trained model following MaskCLIP [58]. Then the visual representation for each concept is produced by averaging the pixel-level representations within it. Since the visual representations and language representations are aligned in CLIP [36], we use the text of the pre-defined categories as the classifier for each concept.

4. Experiment

4.1. Implementation Details

We use ViT-Small [9] trained with DINO [4] as the model to extract pixel-level and region-level representations. From the ViT, we take tokens except the cls token from the last layer as corresponding pixel-level representations. For training and inference of the ACG, we resize the width and height of images to 224 and do not use additional data augmentation. The ACG is optimized by AdamW [30] with a learning rate of 0.0001 and weight decay of 0.01. We train the ACG for 2500 iterations using a batch size of 32. We set the number of update steps in ACG to 6 and the num-

ber of prototypes to 5. During inference, we use bilinear upsampling to restore the soft assignment of pixels to the original resolution before getting the hard assignment. In the concept classifier, we first crop the image by the bounding box of the region and resize it to 224×224 . After that, we mask out the pixels that are not part of the region after image normalization. For qualitative evaluation, we adopt mean intersection over union (mIoU) and pixel accuracy as metrics, following most researches on semantic segmentation. We conduct experiments on two commonly used semantic segmentation datasets PASCAL VOC 2012 [11] and COCO-Stuff [28], and we follow the previous works [20] to adopt the 27-class subset of COCO-Stuff. Other detailed experiment settings are provided in Appendix.

4.2. Qualitative Results

To demonstrate the adaptiveness of the ACG, we first visualize the pixel-level representations and the concepts discovered by the ACG using t-SNE [42], as shown in Figure 4. Benefiting from the powerful representation capabilities of a ViT trained with DINO, the pixel-level representations reflect the semantic relationship between pixels, and pixels with the same semantics are aggregated into concepts. It is obvious that the representations and counts of concepts are various for different images. The ACG handles this situation correctly. It accurately maps the initial prototypes to concepts in the pixel-level representation space for different images, which is shown by meaningful regions when projecting the assignment of each pixel to the original image. Moreover, after optimizing by the proposed modularity




Figure 6. **Visualization of k -NN retrieval results.** We show five concepts with the highest similarity following each query concept (with red frame). The concepts is shown by the highlighted area in the image.

Method	mIoU
IIC [20]	9.8
MaskContrast [43]	35.0
DSM [†] [32]	37.2 ± 3.8
Leopart [59]	41.7
TransFGU [53]	37.2
MaskDistill [44]	42.0
MaskDistill [†] [44]	45.8
ACSeg (Ours)	47.1 ± 2.4

Table 1. **Unsupervised semantic segmentation results on PASCAL VOC.** [†] denotes results with re-training.

Method	mIoU
MoCo v2 [5]	4.4
IIC [20]	6.7
ImageNet [17]	8.9
DINO [4]	9.6
Modified DC [6]	9.8
PiCIE [6]	13.8
PiCIE+H [6]	14.4
ACSeg (Ours)	16.4 ± 0.9

Table 2. **Unsupervised semantic segmentation results on COCO-Stuff-27 dataset.**

Dataset	Method	K=1	K=5
VOC	MaskContrast [43]	43.3	-
	DSM [32]	32.1	31.9
	K-means	45.1	49.1
	Spectral	43.0	47.3
COCO	ACSeg (Ours)	57.8	61.0
	K-means	29.9	33.1
	Spectral	28.5	31.3
	ACSeg (Ours)	30.4	34.0

Table 3. **k -NN retrieval results on PASCAL VOC and COCO-Stuff-27.** Our ACSeg surpasses the clustering baselines and previous methods.

loss, the ACG can produce an adaptive number of concepts for different images, which is achieved by dropping the concepts assigned to no pixels.

In addition, we show the segmentation results after the ACG and the concept classifier in the image in Figure 5. It can be found that the ACG achieves high-quality pixel-level localization (e.g., birds, fence, and tree branch in the first image) and can process scenes of varying complexity (e.g., there are more detected concepts in a complex scene (the first image) than the simple one (the second image), thanks to the flexibility of the proposed modularity loss). The second row of Figure 5 shows the unsupervised semantic segmentation results by clustering the region-level representations. Through clustering, concepts in different images are linked to further form region-level semantic groups.

4.3. Quantitative Results

K-means clustering. We first evaluate the performance of ACSeg for fully unsupervised semantic segmentation and show the results in Table 1 and Table 2. Following recent researches, we adopt k-means clustering on region-level

representations to classify foreground regions and evaluate the quality of clusters with ground truth via Hungarian matching [23]. Since the k-means algorithm is greatly affected by the initial value of cluster centers, we run it ten times and report the results by *mean ± std*. Comparing with the methods which also adopt the region clustering pipeline [32, 44], we reach higher segmentation performance without re-training and achieve state-of-the-art performance on PASCAL VOC 2012 dataset, demonstrating the superiority of the ACSeg brought by its adaptive conceptualization process. In addition, the proposed ACSeg can be viewed as a transfer from self-supervised image-level models to dense prediction tasks, which only exploit the extracted representations of the pre-trained models rather than learning new representations. Therefore, it only requires an extremely small cost and only takes tens of minutes to train and get segmentation results. To this end, it has the advantage of being flexible and easy to use while achieving leading performance compared to the methods that train a segmentation models from scratch.

k -NN retrieval. We evaluate the effect of ACSeg method

Method	mIoU	
	VOC	COCO
MaskCLIP [58]	-	19.6
GroupViT [52]	51.2	20.3
ReCo [39]	-	26.3
ACSeg (<i>Ours</i>)	53.9	28.1

Table 4. Comparison of unsupervised semantic segmentation with text.

with k -NN classifier, to show the quality of region-level representation and the localization quality of the concepts generated by the ACG. In this setting, the class prediction of each concept in the validation set is obtained by the label of its nearest concepts in the training set, and the label of each concept in the training set is set to the same as the most overlapping ground-truth region. For comparison, we report the results of previous works [32, 43] which produce region-level representations and some clustering baselines (k-means clustering and spectral clustering) which replace the ACG in our method, as shown in Table 3. To intuitively show the retrieval performance, we also show some concepts and their five nearest neighbors in Figure 6. The proposed ACSeg surpasses the previous methods and baselines, showing the promising performance of fully unsupervised region-level representations.

Unsupervised semantic segmentation with text. The emergence of vision-language pre-training models helps the unsupervised classification in visual tasks by constructing classifiers from prompt texts. MaskCLIP [58] proposes to modify the visual encoder of CLIP [36] and apply the text-based classifiers to pixel level. However, the pixel-level classification results are relatively coarse due to the image-level pre-training task. Here we demonstrate the effectiveness of combining the localization ability of the discovered concepts and the classification ability of CLIP through Table 4. With the help of accurate localization of concepts, we get higher performance than the related works MaskCLIP [58], GroupViT [52] and ReCo [39].

4.4. Ablation Study

Effectiveness of the ACG. As the crucial part of ACSeg, we discuss the effectiveness of the ACG here on PASCAL VOC dataset. Since the ACSeg can be regarded as a kind of clustering, we adopt some commonly used clustering methods k-means [16], spectral clustering [46], affinity propagation [10], and agglomerative clustering [33]. The results in Table 5 demonstrate the effectiveness of the ACG. The ACG is more suitable for clustering pixel-level representations of produced by self-supervised ViT due to adaptively considering the attributes in different images, so that achieving better performance. Thanks to the adaptiveness of the ACG, it can process all images through a single network.

Method	mIoU	Speed
K-Means [16]	28.6	2.4
Spectral [46]	28.3	3.4
AP [10]	11.0	6.8
Agglomerative [33]	13.9	15.8
ACG (<i>Ours</i>)	47.1	149.2

Table 5. Effectiveness and efficiency of the ACG. Speed indicates images per second.

Num	Clustering	Retrieval
2	35.6	43.7
7	46.9	56.9
10	42.1	54.4
15	34.5	51.1
5	47.1	57.8

Table 6. The effect of choosing different number of prototypes.

Therefore, it is faster and can be highly accelerated by GPU compared to the classical iterative clustering algorithms.

The number of prototypes is the main design choice in our method. Although the ACG generates an adaptive number of concepts, the number of prototypes impacts on the granularity of conceptualization. Specifically, it determines the upper bound on the concepts detected in each image, and more prototypes lead to finer-grained conceptualization for all images. We show the clustering and retrieval results on VOC with different numbers of prototypes in Table 6. When setting the number to 2, the ACG degenerates into foreground-background segmentation, resulting in performance degradation. A large value also degrades the performance since over-clustering may occur on some objects with finer-grained division. Our method is robust when the value is within a reasonable range.

5. Conclusion and Discussion

In this work, we propose the ACSeg, an approach that efficiently transfers the pixel-level knowledge of self-supervised ViT for unsupervised semantic segmentation. We design the ACG to achieve adaptive conceptualization, *i.e.*, generating concepts from pixel-level representations while considering the semantic distributions of different images. Meanwhile, a novel loss function called modularity loss is designed to optimize the ACG without relying on any annotations and enable an adaptive number of concepts in different images. Finally, we turn the USS task into classifying the discovered concepts in an unsupervised manner. Qualitative and quantitative experimental results under different settings demonstrate the effectiveness and superiority of the proposed method as a directly transfer for self-supervised ViT without learning new representations.

Acknowledgements. This work was supported in part by Natural Science Foundation of China (No. 61972217, 32071459, 62176249, 62006133, 62271465, 62202014), and the Natural Science Foundation of Guangdong Province in China (No. 2019B1515120049). Li Yuan was supported in part by the National Key R&D Program of China (2022ZD0118101) and also sponsored by CCF Tencent Open Research Fund.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4981–4990, 2018. 1
- [2] Adam Bielski and Paolo Favaro. Move: Unsupervised movable object segmentation and detection. *arXiv preprint arXiv:2210.07920*, 2022. 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in Neural Information Processing Systems*, 33:1877–1901, 2020. 2
- [4] 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*, pages 9650–9660, 2021. 2, 6, 7
- [5] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 7
- [6] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16794–16804, 2021. 2, 7
- [7] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016. 1
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 2
- [9] 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. In *International Conference on Learning Representations*, 2020. 2, 6
- [10] Delbert Dueck. *Affinity propagation: clustering data by passing messages*. University of Toronto Toronto, ON, Canada, 2009. 8
- [11] Mark Everingham and John Winn. The pascal visual object classes challenge 2012 (voc2012) development kit. *Pattern Anal. Stat. Model. Comput. Learn., Tech. Rep.*, 2007:1–45, 2012. 1, 2, 6
- [12] Deng-Ping Fan, Ge-Peng Ji, Tao Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Pranet: Parallel reverse attention network for polyp segmentation. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, pages 263–273. Springer, 2020. 1
- [13] Santo Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010. 5
- [14] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1
- [15] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. In *International Conference on Learning Representations*, 2021. 2, 3
- [16] John A Hartigan. *Clustering algorithms*. John Wiley & Sons, Inc., 1975. 8
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 7
- [18] Haiyang Huang, Zhi Chen, and Cynthia Rudin. Segdiscover: Visual concept discovery via unsupervised semantic segmentation. *arXiv preprint arXiv:2204.10926*, 2022. 2, 3
- [19] Jyh-Jing Hwang, Stella X Yu, Jianbo Shi, Maxwell D Collins, Tien-Ju Yang, Xiao Zhang, and Liang-Chieh Chen. Segsort: Segmentation by discriminative sorting of segments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7334–7344, 2019. 1, 2
- [20] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019. 1, 2, 6, 7
- [21] Peng Jin, JinFa Huang, Fenglin Liu, Xian Wu, Shen Ge, Guoli Song, David A. Clifton, and Jie Chen. Expectation-maximization contrastive learning for compact video-and-language representations. In *Thirty-Sixth Conference on Neural Information Processing Systems*, 2022. 2
- [22] Tsung-Wei Ke, Jyh-Jing Hwang, Yunhui Guo, Xudong Wang, and Stella X Yu. Unsupervised hierarchical semantic segmentation with multiview cosegmentation and clustering transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2571–2581, 2022. 2
- [23] Harold W Kuhn. The hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97, 1955. 7
- [24] Neeraj Kumar, Ruchika Verma, Sanuj Sharma, Surabhi Bhargava, Abhishek Vahadane, and Amit Sethi. A dataset and a technique for generalized nuclear segmentation for computational pathology. *IEEE Transactions on Medical Imaging*, 36(7):1550–1560, 2017. 1
- [25] Hao Li, Jinfa Huang, Peng Jin, Guoli Song, Qi Wu, and Jie Chen. Toward 3d spatial reasoning for human-like text-based visual question answering. *arXiv preprint arXiv:2209.10326*, 2022. 2
- [26] Hao Li, Xu Li, Belhal Karimi, Jie Chen, and Mingming Sun. Joint learning of object graph and relation graph for visual question answering. *arXiv preprint arXiv:2205.04188*, 2022. 2
- [27] Di Lin, Jifeng Dai, Jiaya Jia, Kaiming He, and Jian Sun. Scribblesup: Scribble-supervised convolutional networks for

- semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3159–3167, 2016. 1
- [28] 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*, pages 740–755. Springer, 2014. 1, 2, 6
- [29] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3431–3440, 2015. 1
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2018. 6
- [31] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Finding an unsupervised image segmenter in each of your deep generative models. In *International Conference on Learning Representations*, 2021. 3
- [32] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Deep spectral methods: A surprisingly strong baseline for unsupervised semantic segmentation and localization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8364–8375, 2022. 2, 3, 7, 8
- [33] Fionn Murtagh and Pedro Contreras. Algorithms for hierarchical clustering: an overview. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 2(1):86–97, 2012. 8
- [34] Mark EJ Newman and Michelle Girvan. Finding and evaluating community structure in networks. *Physical Review E*, 69(2):026113, 2004. 2, 3, 4, 5
- [35] Youngmin Oh, Beomjun Kim, and Bumsub Ham. Background-aware pooling and noise-aware loss for weakly-supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6913–6922, 2021. 1
- [36] 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*, pages 8748–8763. PMLR, 2021. 2, 3, 6, 8
- [37] Wei Shen, Zelin Peng, Xuehui Wang, Huayu Wang, Jiazhong Cen, Dongsheng Jiang, Lingxi Xie, Xiaokang Yang, and Qi Tian. A survey on label-efficient deep segmentation: Bridging the gap between weak supervision and dense prediction. *arXiv preprint arXiv:2207.01223*, 2022. 1
- [38] Gyungin Shin, Weidi Xie, and Samuel Albanie. Named-mask: Distilling segmenters from complementary foundation models. *arXiv:2209.11228*, 2022. 3
- [39] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. *arXiv preprint arXiv:2206.07045*, 2022. 3, 8
- [40] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. *arXiv preprint arXiv:2109.14279*, 2021. 2
- [41] Korsuk Sirinukunwattana, Josien PW Pluim, Hao Chen, Xiaojuan Qi, Pheng-Ann Heng, Yun Bo Guo, Li Yang Wang, Bogdan J Matuszewski, Elia Bruni, Urko Sanchez, et al. Gland segmentation in colon histology images: The glas challenge contest. *Medical Image Analysis*, 35:489–502, 2017. 1
- [42] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9(11), 2008. 6
- [43] 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*, pages 10052–10062, 2021. 2, 7, 8
- [44] Wouter Van Gansbeke, Simon Vandenhende, and Luc Van Gool. Discovering object masks with transformers for unsupervised semantic segmentation. *arXiv preprint arXiv:2206.06363*, 2022. 2, 3, 7
- [45] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in Neural Information Processing Systems*, 30, 2017. 2, 3, 4
- [46] Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and Computing*, 17(4):395–416, 2007. 8
- [47] Xinlong Wang, Zhiding Yu, Shalini De Mello, Jan Kautz, Anima Anandkumar, Chunhua Shen, and Jose M Alvarez. Freesolo: Learning to segment objects without annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14176–14186, 2022. 3
- [48] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14543–14553, 2022. 2
- [49] Yude Wang, Jie Zhang, Meina Kan, Shiguang Shan, and Xilin Chen. Self-supervised equivariant attention mechanism for weakly supervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12275–12284, 2020. 1
- [50] Yuan Wang, Wei Zhuo, Yucong Li, Zhi Wang, Qi Ju, and Wenwu Zhu. Fully self-supervised learning for semantic segmentation. *arXiv preprint arXiv:2202.11981*, 2022. 2
- [51] Xin Wen, Bingchen Zhao, Anlin Zheng, Xiangyu Zhang, and Xiaojuan Qi. Self-supervised visual representation learning with semantic grouping. *arXiv preprint arXiv:2205.15288*, 2022. 2
- [52] 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*, pages 18134–18144, 2022. 3, 8

- [53] Zhaoyuan Yin, Pichao Wang, Fan Wang, Xianzhe Xu, Hanling Zhang, Hao Li, and Rong Jin. Transfgu: a top-down approach to fine-grained unsupervised semantic segmentation. In *European Conference on Computer Vision*, pages 73–89. Springer, 2022. [2](#), [3](#), [7](#)
- [54] Andrii Zadaianchuk, Matthaeus Kleindessner, Yi Zhu, Francesco Locatello, and Thomas Brox. Unsupervised semantic segmentation with self-supervised object-centric representations. *arXiv preprint arXiv:2207.05027*, 2022. [2](#)
- [55] Feihu Zhang, Philip Torr, René Ranftl, and Stephan Richter. Looking beyond single images for contrastive semantic segmentation learning. *Advances in Neural Information Processing Systems*, 34:3285–3297, 2021. [2](#)
- [56] Xiao Zhang and Michael Maire. Self-supervised visual representation learning from hierarchical grouping. *Advances in Neural Information Processing Systems*, 33:16579–16590, 2020. [2](#)
- [57] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 633–641, 2017. [1](#)
- [58] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision (ECCV)*, 2022. [3](#), [6](#), [8](#)
- [59] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14502–14511, 2022. [2](#), [7](#)