

# Diffuse, Attend, and Segment: Unsupervised Zero-Shot Segmentation using Stable Diffusion

Junjiao Tian<sup>1</sup>  
jtian73@gatech.edu

Lavisha Aggarwal<sup>2</sup>  
lavishaggarwal@google.com

Andrea Colaco<sup>2</sup>  
andreacolaco@google.com

Zsolt Kira<sup>1</sup>  
zk15@gatech.edu

Mar Gonzalez-Franco<sup>2</sup>  
margon@google.com

## Abstract

Producing quality segmentation masks for images is a fundamental problem in computer vision. Recent research has explored large-scale supervised training to enable zero-shot transfer segmentation on virtually any image style and unsupervised training to enable segmentation without dense annotations. However, constructing a model capable of segmenting anything in a zero-shot manner without any annotations is still challenging. In this paper, we propose to utilize the self-attention layers in stable diffusion models to achieve this goal because the pre-trained stable diffusion model has learned inherent concepts of objects within its attention layers. Specifically, we introduce a simple yet effective iterative merging process based on measuring KL divergence among attention maps to merge them into valid segmentation masks. The proposed method does not require any training or language dependency to extract quality segmentation for any images. On COCO-Stuff-27, our method surpasses the prior unsupervised zero-shot transfer SOTA method by an absolute 26% in pixel accuracy and 17% in mean IoU. The project page is at <https://sites.google.com/view/diffseg/home>.<sup>1</sup>

## 1. Introduction

Semantic segmentation divides an image into regions of entities sharing the same semantics. It is an important foundational application for many downstream tasks such as image editing [26], medical imaging [27], and autonomous driving [14] etc. While supervised semantic segmentation, where a target dataset is available and the categories are known, has been widely studied [10, 38, 40], zero-shot

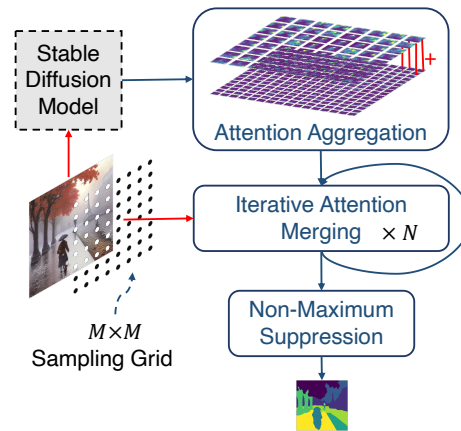


Figure 1. Overview of DiffSeg. DiffSeg is an unsupervised and zero-shot segmentation algorithm using a pre-trained stable diffusion model. Starting from  $M \times M$  anchor points, DiffSeg iteratively merges self-attention maps from the diffusion model for  $N$  iterations to segment any image without any prior knowledge and external information.

transfer segmentation<sup>2</sup> for any images with unknown categories is much more challenging. A recent popular work SAM [20], trains a neural network on 1.1B segmentation annotations and achieves impressive zero-shot transfer to any images. This is an important step towards making segmentation a more flexible foundation task for many other tasks, not just limited to a given dataset with limited categories. However, the cost of collecting per-pixel labels is high. Therefore, it is of high research and production interest to investigate *unsupervised* and *zero-shot transfer* segmentation methods under the least restrictive settings, i.e.,

<sup>2</sup>Zero-shot transfer segmentation means that the underlying model is not trained for segmentation, similar to CLIP which is not trained for classification but can be used for zero-shot classification [32].

<sup>1</sup>Georgia Institute of Technology<sup>1</sup> Google<sup>2</sup>

no form of annotations and no prior knowledge of the target.

Few works have taken up the challenge due to the combined difficulty of the unsupervised and zero-shot requirements. For example, most works in unsupervised segmentation require access to the target data for unsupervised adaptation [11, 15, 16, 19, 24, 37, 42]. Therefore, these methods cannot segment images that are not seen during the adaptation. Recently, ReCo [36] proposed a retrieve-and-co-segment strategy. It removes the requirement for supervision because it uses an unsupervised vision backbone, e.g., DINO [9], and is zero-shot since it does not need training on the target distribution of images. Still, to satisfy these requirements, ReCo needs to 1) identify the concept beforehand and 2) maintain a large pool of unlabelled images for retrieval. This is still far from the capability of SAM [20], which does not require knowledge of the target image and any auxiliary information.

To move towards the goal, we propose to utilize the power of a stable diffusion (SD) model [33] to construct a general segmentation model. Recently, stable diffusion models have been used to generate prompt-conditioned high-resolution images [33]. It is reasonable to hypothesize that there exists information on object groupings in a diffusion model. For example, DiffuMask [39] discovers that the *cross-attention* layer contains explicit object-level pixel groupings when cross-referenced between the produced attention maps and the input prompt. However, DiffuMask can only produce segmentation masks for a generated image with a corresponding text prompt and is limited to dominant foreground objects explicitly referred to by the prompt. Motivated by this, we investigate the unconditioned *self-attention* layer in a diffusion model. We observed that the attention tensors from the self-attention layers contain specific spatial relationships: **Intra-Attention Similarity** and **Inter-Attention Similarity** (Sec. 3.1). Relying on these two properties, it is possible to uncover segmentation masks for any objects in an image.

To this end, we propose DiffSeg (Sec. 3.2), a simple yet effective post-processing method, for producing segmentation masks by using the attention tensors generated by the self-attention layers in a diffusion model solely. The algorithm consists of three main components: attention aggregation, iterative attention merging, and non-maximum suppression, as shown in Fig. 1. Specifically, DiffSeg aggregates the 4D attention tensors in a spatially consistent manner to preserve visual information across multiple resolutions and uses an iterative merging process by first sampling a grid of anchor points. The sampled anchors provide starting points for merging attention masks, where anchors in the same object are absorbed eventually. The merging process is controlled by measuring the similarity between two attention maps using KL divergence. Unlike, popular clustering-based unsupervised segmentation

methods [11, 16, 24], DiffSeg does not require specifying the number of clusters beforehand and is also deterministic. Given an image, without any prior knowledge, DiffSeg can produce a quality segmentation without resorting to any additional resources (e.g. as SAM requires). We benchmark DiffSeg on the popular unsupervised segmentation dataset COCO-Stuff-27 and a specialized self-driving dataset Cityscapes. DiffSeg outperforms prior works on both datasets despite requiring less auxiliary information during inference. In summary,

- We propose an unsupervised and zero-shot segmentation method, DiffSeg, using a pre-trained stable diffusion model. DiffSeg can segment images in the wild without any prior knowledge or additional resources.
- DiffSeg sets state-of-the-art performance on two segmentation benchmarks. On COCO-Stuff-27, our method surpasses a prior unsupervised zero-shot SOTA method by an absolute 26% in pixel accuracy and 17% in mean IoU.

## 2. Related Works

**Diffusion Models.** Our work builds on pre-trained stable diffusion models [18, 33, 34]. Many existing works have studied the discriminative visual features in them and used them for zero-shot classification [22], supervised segmentation [4], label-efficient segmentation [6], semantic-correspondence matching [17, 43] and open-vocabulary segmentation [41]. They rely on the high-dimensional *visual features* learned in stable diffusion models to perform those tasks requiring additional training to take full advantage of those features. Instead, we show that object grouping is also an emergent property of the *self-attention* layer in the Transformer module manifested in 4-dimensional spatial attention tensors with particular spatial patterns.

**Unsupervised Segmentation.** Our work is closely related to unsupervised segmentation [11, 15, 16, 19, 24, 37, 42]. This line of works aims to generate dense segmentation masks without using any annotations. However, they generally require unsupervised training on the target dataset to achieve good performance. We characterize the most recent works into two main categories: clustering based on invariance [11, 19] and, more recently, clustering using pre-trained models [16, 24, 30, 39, 42]. Our work falls in the second category. To produce segmentation masks, these new works utilize the discriminative features learned in those pre-trained models, i.e., features from the same semantic class are generally more similar. DiffuMasks [39] uses the cross-modal grounding between the text prompt and the image in the cross-attention layer of a diffusion model to segment the most prominent object, referred to by the text prompt in the synthetic image. However, DiffuMasks can only be applied to a generated image. In contrast, our method applies to real images without prompts.

**Zero-Shot Transfer Segmentation** Another related

field is zero-shot segmentation [7, 15, 20, 29, 36, 39, 45]. Works in this line of work possess the capability of “segmenting anything” without any training. Most recently, SAM [20], trained on 1.1B segmentation masks, demonstrated impressive zero-shot transfer to any images. Some recent works utilize auxiliary information such as additional images [15, 36] or text inputs [29, 45] to facilitate zero-shot transfer segmentation. In contrast, our method uses a diffusion model to generate segmentation without synthesizing and querying multiple images and, more importantly, without knowing the objects in the image.

### 3. Method

DiffSeg utilizes a pre-trained stable diffusion model [33] and specifically its *self-attention* layers to produce high-quality segmentation masks. We will briefly review the architecture of the stable diffusion model in Sec. 3.1 and introduce DiffSeg in Sec. 3.2.

#### 3.1. Stable Diffusion Model Review

The Stable diffusion model [33] is a popular variant of the diffusion model family [18, 34], a generative model. Diffusion models have a forward and a reverse pass. In the forward pass, at every time step, a small amount of Gaussian noise is iteratively added until an image becomes an isotropic Gaussian noised image. In the reverse pass, the diffusion model is trained to iteratively remove the Gaussian noise to recover the original clean image. The stable diffusion model [33] introduces an encoder-decoder and U-Net design with attention layers [18]. It first compress an image  $x \in \mathbb{R}^{H \times W \times 3}$  into a latent space with smaller spatial dimension  $z \in \mathbb{R}^{h \times w \times c}$  using an encoder  $z = \mathcal{E}(x)$ , which can be de-compressed through a decoder  $\hat{x} = \mathcal{D}(z)$ . All diffusion processes happen in the latent space through a U-Net architecture. The U-Net architecture will be the focus of this paper’s investigation.

The U-Net architecture [33] consists of a stack of modular blocks (a schematic is provided in Appendix 9.1). Among those blocks, there are 16 specific blocks composed of two major components: a ResNet layer and a Transformer layer. The Transformer layer uses two attention mechanisms: self-attention to learn the global attention across the image and cross-attention to learn attention between the image and optional text input. The component of interest for our investigation is the *self-attention* layer in the Transformer layer. Specifically, there are a total of 16 self-attention layers distributed in the 16 composite blocks, giving 16 *self-attention tensors*:

$$\mathcal{A} \in \{\mathcal{A}_k \in \mathbb{R}^{h_k \times w_k \times h_k \times w_k} | k = 1, \dots, 16\}. \quad (1)$$

Each attention tensor  $\mathcal{A}_k$  is 4-dimensional<sup>3</sup>. Inspired by DiffuMask [39], which demonstrates object grouping in the

<sup>3</sup>There is a fifth dimension due to the multi-head attention mechanism.

*cross-attention* layer, where salient objects are given more attention when referred by its corresponding text input, we hypothesize that the unconditional self-attention also contains inherent object grouping information and can be used to produce segmentation masks without text inputs.

Specifically, for each spatial location  $(I, J)$  in the attention tensor, the corresponding 2D *attention map*  $\mathcal{A}_k[I, J, :, :] \in \mathbb{R}^{h_k \times w_k}$ <sup>4</sup> captures the semantic correlation between all locations and the location  $(I, J)$ . Each location  $(I, J)$  corresponds to a region in the original image pixel space, the size of which depends on the receptive field of the tensor.

To illustrate this, we visualize the self-attention tensors  $\mathcal{A}_7$  and  $\mathcal{A}_8$  in Fig. 2. They have a dimension of  $\mathcal{A}_7 \in \mathbb{R}^{8^4}$  and  $\mathcal{A}_8 \in \mathbb{R}^{16^4}$ <sup>5</sup> respectively. Two important observations will motivate our method in Sec. 3.2.

- **Intra-Attention Similarity:** Within a 2D attention map  $\mathcal{A}_k[I, J, :, :]$ , locations tend to have strong responses if they correspond to the same *object group* as  $(I, J)$  in the original image space, e.g.,  $\mathcal{A}_k[I, J, I + 1, J + 1]$  is likely to be a large value.
- **Inter-Attention Similarity:** Between two 2D attention maps, e.g.,  $\mathcal{A}_k[I, J, :, :]$  and  $\mathcal{A}_k[I + 1, J + 1, :, :]$ , they tend to share similar activations if  $(I, J)$  and  $(I + 1, J + 1)$  belong to the same *object group* in the original image space.

The second observation is a direct consequence of the first one. In general, we found that attention maps tend to focus on *object groups*, i.e., groups of individual objects sharing similar visual features (see segmentation examples in Fig. 6). In this example (Fig. 2), the two people are grouped as a single *object group*. The  $8 \times 8$  resolution map from a location inside the group highlights much of the entire group. In contrast, the  $16 \times 16$  resolution map from a location inside the object highlights a smaller portion.

Theoretically, the resolution of the attention map dictates the size of its receptive field w.r.t the original image. A smaller resolution leads to a larger respective field, corresponding to a larger portion of the original image. Practically, lower resolution maps, e.g.,  $8 \times 8$ , provide a better grouping of large objects as a whole, and larger resolution maps, e.g.,  $16 \times 16$ , give a more fine-grained grouping of components in larger objects and potentially are better for identifying small objects. The current stable diffusion model has attention maps in 4 resolutions  $(h_k, w_k) \in \{8 \times 8, 16 \times 16, 32 \times 32, 64 \times 64\}$ . Building on these observations, in the next section, we will propose a simple heuristic to aggregate weights from different resolutions and introduce an iterative method to merge all attention maps into a

Every attention layer has 8 multi-head outputs. We average the attention tensors along the multi-head axis to reduce to 4 dimensions because they are very similar across this dimension. Please see Appendix 9.2 for details.

<sup>4</sup> $\sum \mathcal{A}_k[I, J, :, :] = 1$  is a valid probability distribution.

<sup>5</sup> $\mathbb{R}^{16^4}$  denotes  $\mathbb{R}^{16 \times 16 \times 16 \times 16}$ .




Figure 2. Visualization of Segmentation Masks and Self-Attentions Tensors. Left: Overlay of segmentation and the original image. Right: Attention maps from a stable diffusion model have two properties: Intra-Attention Similarity and Inter-Attention Similarity. Maps of different resolutions have varying receptive fields w.r.t the original image.

valid segmentation.

**A Note on Extracting Attention Maps.** In our experiments, we use the stable diffusion pre-trained models from Huggingface. These prompt-conditioned diffusion models typically run for  $\geq 50$  diffusion steps to generate new images. However, in our case, we want to efficiently extract attention maps for an existing clean image without conditional prompts. We achieve this by 1) using only the unconditioned latent and 2) only running the diffusion process once. The unconditional latent is calculated using an unconditioned text embedding. Since we only do one pass through the diffusion model, we also set the time-step variable  $t$  to a large value, e.g.,  $t = 300$ , in our experiments such that the real images are viewed as primarily denoised generated images from the perspective of the diffusion model.

### 3.2. DiffSeg

Since the self-attention layers capture inherent object grouping information in spatial attention (probability) maps, we propose DiffSeg, a simple post-processing method, to aggregate and merge attention tensors into a valid segmentation mask. The pipeline consists of three components: attention aggregation, iterative attention merging, and non-maximum suppression, as shown in Fig. 1. We will introduce each element in detail next.

**Attention Aggregation.** Given an input image passing through the encoder and U-Net, a stable diffusion model generates 16 attention tensors. Specifically, there are 5 tensors for each of the dimensions:  $64 \times 64 \times 64 \times 64$ ,  $32 \times 32 \times 32 \times 32$ ,  $16 \times 16 \times 16 \times 16$ ,  $8 \times 8 \times 8 \times 8$ ; and 1 tensor for the dimension  $8 \times 8 \times 8 \times 8$ . The goal is to aggregate attention tensors of different resolutions into the highest resolution tensor. To achieve this, we need to carefully treat the 4 dimensions differently.

As discussed in the previous section, the 2D map  $\mathcal{A}_k[I, J, :, :]$  corresponds to the correlation between all spatial locations and the location  $(I, J)$ . Therefore, the last 2 dimensions in the attention tensors are spatially consistent despite different resolutions. Therefore, we upsample (bi-

linear interpolation) the last 2 dimensions of all attention maps to  $64 \times 64$ , their highest resolution. Formally, for  $\mathcal{A}_k \in \mathbb{R}^{h_k \times w_k \times h_k \times w_k}$ ,

$$\tilde{\mathcal{A}}_k = \text{Bilinear-upsample}(\mathcal{A}_k) \in \mathbb{R}^{h_k \times w_k \times 64 \times 64}. \quad (2)$$

On the other hand, the first 2 dimensions indicate the locations to which attention maps are referenced. Therefore, we need to aggregate attention maps accordingly. For example, as shown in Fig. 3a, the attention map in the  $(0, 0)$  location in  $\mathcal{A}_k \in \mathbb{R}^{8^4}$  is first upsampled and then repeatedly aggregated pixel-wise with the 4 attention maps  $(0, 0)$ ,  $(0, 1)$ ,  $(1, 0)$ ,  $(1, 1)$  in  $\mathcal{A}_z \in \mathbb{R}^{16^4}$ . Formally, the final aggregated attention tensor  $\mathcal{A}_f \in \mathbb{R}^{64^4}$  is,


$$\mathcal{A}_f[I, J, :, :] = \sum_{k \in \{1, \dots, 16\}} \tilde{\mathcal{A}}_k[I/\delta_k, J/\delta_k, :, :] * R_k, \quad (3)$$

$$\text{where } \delta_k = 64/w_k, \quad \sum_k R_k = 1.$$

where  $/$  denotes floor division here. Furthermore, to ensure that the aggregated attention map is also a valid distribution, i.e.,  $\sum \mathcal{A}_f[I, J, :, :] = 1$ , we normalize  $\mathcal{A}_f[I, J, :, :]$  after aggregation.  $R_k$  is the aggregation importance ratio for each attention map. We assign every attention map of different resolutions a weight proportional to its resolution  $R_k \propto w_k$ . We call this the proportional aggregation scheme (Propto.). The weights  $R$  are important hyper-parameters, and we conduct a detailed study of them in Appendix 9.6.

**Iterative Attention Merging.** In the previous step, the algorithm computes an attention tensor  $\mathcal{A}_f \in \mathbb{R}^{64^4}$ . In this section, the goal is to merge the  $64 \times 64$  attention maps in the tensor  $\mathcal{A}_f$  to a stack of *object proposals* where each proposal likely contains the activation of a single object or stuff category. The naive solution is to run a K-means algorithm on  $\mathcal{A}_f$  to find clusters of objects following existing works in unsupervised segmentation [11, 16, 24]. However, the K-means algorithm requires the specification of the number of clusters. This is not an intuitive hyper-parameter to tune because we aim to segment any image in the wild. Moreover, the K-means algorithm is stochastic depending on the





(a) Example of Attention Aggregation. An attention map from a lower resolution is first upsampled and then duplicated to match the receptive field of the higher-resolution maps.

(b) Example of Non-Maximum Suppression (NMS). NMS looks up the maximum activation across the  $\mathcal{L}_p$  proposals for each pixel.

Figure 3

initialization. Each run can have a drastically different result for the same image. To highlight these limitations, we compare them with K-Means baselines in our experiments.

Instead, we propose to generate a sampling grid from which the algorithm can iteratively merge attention maps to create the proposals. Specifically, as shown in Fig. 1, in the *sampling grid generation* step, a set of  $M \times M$  ( $1 \leq M \leq 64$ ) evenly spaced anchor points are generated  $\mathcal{M} = \{(i_m, j_m) | m = 1, \dots, M^2\}$ . We then sample the corresponding attention maps from the tensor  $\mathcal{A}_f$ . This operation yields a list of  $M^2$  2D attention maps as anchors,

$$\mathcal{L}_a = \{\mathcal{A}_f[i_m, j_m, :, :] \in \mathbb{R}^{64 \times 64} | (i_m, j_m) \in \mathcal{M}\}. \quad (4)$$

Since we aim to merge attention maps to find the maps most likely containing an object, we rely on the two observations in Sec. 3.1. **Specifically, when iteratively merging “similar” maps, Intra-Attention Similarity reinforces the activation of an object and Inter-Attention Similarity grows the activation to include as many pixels of the same object within the merged map.** To measure similarity, we use KL divergence to calculate the “distance” between two attention maps, since each attention map  $\mathcal{A}_f[i, j, :, :]$  (abbrev.  $\mathcal{A}_f[i, j]$ ) is a valid probability distribution. Formally,

$$\begin{aligned} 2 * D(\mathcal{A}_f[i, j], \mathcal{A}_f[y, z]) \\ = (\text{KL}(\mathcal{A}_f[i, j] || \mathcal{A}_f[y, z]) + \text{KL}(\mathcal{A}_f[y, z] || \mathcal{A}_f[i, j])). \end{aligned} \quad (5)$$

We use the forward and reverse KL to address the asymmetry in KL divergence. Intuitively, a small  $\mathcal{D}(\cdot)$  indicates a high similarity between two attention maps and that their union is likely to represent better the object they both belong to.

Then, we start the  $N$  iterations of the merging process. **In the first iteration**, we compute the pair-wise distance between each element in the anchor list and all attention maps using  $\mathcal{D}(\cdot)$ . We introduce a threshold hyper-parameter  $\tau$ . For each element in the list, we average all attention

maps with a distance smaller than the threshold, effectively taking the union of attention maps likely belonging to the same object that the anchor point belongs to. All merged attention maps are stored in a new proposal list  $\mathcal{L}_p$ .

Note that the first iteration does *not* reduce the number of proposals compared to the number of anchors. **From the second iteration onward**, the algorithm starts merging maps and reducing the number of proposals simultaneously by computing the distance between an element from the proposal list  $\mathcal{L}_p$  and all elements from the *same* list, and merging elements with distance smaller than  $\tau$  *without* replacement. The entire iterative attention merging algorithm is provided in Alg. 1 (Appendix 9.3).

**Non-Maximum Suppression.** The iterative attention merging step yields a list  $\mathcal{L}_p \in \mathbb{R}^{N_p \times 64 \times 64}$  of  $N_p$  object proposals in the form of attention maps (probability maps). Each proposal in the list potentially contains the activation of a single object. To convert the list into a valid segmentation mask, we use non-maximum suppression (NMS). This can be easily done since each element is a probability distribution map. We can take the index of the largest probability at each spatial location across all maps and assign that location’s membership to the corresponding map’s index. An example is shown in Fig. 3b. Note that, before NMS, we upsample all elements in  $\mathcal{L}_p$  to the original resolution. Formally, the final segmentation mask  $\mathcal{S} \in \mathbb{R}^{512 \times 512}$  is

$$\begin{aligned} \bar{\mathcal{L}}_p &= \text{Bilinear-upsample}(\mathcal{L}_p) \in \mathbb{R}^{N_p \times 512 \times 512} \\ \mathcal{S}[i, j] &= \arg \max \bar{\mathcal{L}}_p[:, i, j] \quad \forall i, j \in \{0, \dots, 511\}. \end{aligned} \quad (6)$$

## 4. Experiments

**Datasets.** Following existing works in unsupervised segmentation [8, 11, 16, 19, 36, 45], we use two popular segmentation benchmarks, COCO-Stuff-27 [11, 12, 19] and Cityscapes [12]. COCO-stuff-27 [11, 19] is a curated version of the original COCO-stuff dataset [12]. Specifically, COCO-stuff-27 merges the 80 things and 91 stuff categories in COCO-stuff into 27 mid-level categories. We evaluate our method on the validation split curated by prior

Model	Requirements			Metrics	
	LD	AX	UA	ACC.	mIoU
IIC [19]	✗	✗	✓	21.8	6.7
MDC [8]	✗	✗	✓	32.3	9.8
PiCLE [11]	✗	✗	✓	48.1	13.8
PiCLE+H [11]	✗	✗	✓	50.0	14.4
STEGO [16]	✗	✗	✓	56.9	28.2
ACSeg [24]	✓	✗	✓	-	28.1
MaskCLIP [45]	✓	✗	✗	32.2	19.6
ReCo [36]	✓	✓	✗	46.1	26.3
K-Means-C (512)	✗	✗	✗	58.9	33.7
K-Means-S (512)	✗	✗	✗	62.6	34.7
DBSCAN (512)	✗	✗	✗	57.7	27.2
Ours: DinoSeg (224)	✗	✗	✗	68.2	39.1
Ours: DiffSeg (320)	✗	✗	✗	<b>72.5</b>	43.0
Ours: DiffSeg (512)	✗	✗	✗	<b>72.5</b>	<b>43.6</b>

Table 1. Evaluation on COCO-Stuff-27. Language Dependency (LD), Auxiliary Images (AX), Unsupervised Adaptation (UA). We additionally present baseline results using K-Means. K-Means-C (constant) uses a constant number of clusters, 6. K-means-S (specific) uses a specific number of clusters for each image based on the ground truth. The K-Means results use  $512 \times 512$  resolution.

works [11, 19]. Cityscapes [12] is a self-driving dataset with 27 categories. We evaluate the official validation split. For both datasets, we report results for the diffusion-model-native  $512 \times 512$  input resolution and a lower  $320 \times 320$  resolution following prior works.

**Metrics.** Following prior works [8, 11, 16, 19, 36, 45], we use pixel accuracy (ACC) and mean intersection over union (mIoU) to measure segmentation performance. Because our method does not provide a semantic label, we use the Hungarian matching algorithm [21] to assign each predicted mask to a ground truth mask. When there are more predicted masks than ground truth masks, the unmatched predicted masks will be considered false negatives. In addition to the metrics, we highlight the requirements other baselines need in inference. Specifically, we emphasize unsupervised adaptation (UA), language dependency (LD), and auxiliary image (AX). UA means that the specific method requires unsupervised training on the target dataset. This is common in the unsupervised segmentation literature. Methods without the UA requirement are considered zero-shot. LD means that the method requires text input, such as a descriptive sentence for the image, to facilitate segmentation. AX means that the method requires an additional pool of reference images or synthetic images.

**Models.** DiffSeg builds on pre-trained stable diffusion models. We use stable diffusion V1.4 [3].

Model	Requirements			Metrics	
	LD	AX	UA	ACC.	mIoU
IIC [19]	✗	✗	✓	47.9	6.4
MDC [8]	✗	✗	✓	40.7	7.1
PiCLE [11]	✗	✗	✓	65.5	12.3
STEGO [16]	✗	✗	✓	73.2	21.0
MaskCLIP [45]	✓	✗	✗	35.9	10.0
ReCo [36]	✓	✓	✗	74.6	19.3
Ours: DiffSeg (320)	✗	✗	✗	67.3	15.2
Ours: DiffSeg (512)	✗	✗	✗	<b>76.0</b>	<b>21.2</b>

Table 2. Evaluation on Cityscapes. Language Dependency (LD), Auxiliary Images (AX), Unsupervised Adaptation (UA).

## 4.1. Main Results

On the COCO benchmark, we include two K-means baselines, K-Means-C and K-Means-S. K-Means-C uses a constant number of clusters, 6, calculated by averaging over the number of objects in all evaluated images. K-Means-S uses a specific number for each image based on the number of objects in the ground truth of that image. We also include another classic unsupervised clustering algorithm DBSCAN [13] that automatically discovers clusters just like DiffSeg. In Tab. 1, we observe that both K-Means variants outperform prior works, demonstrating the advantage of using self-attention tensors. Furthermore, the K-Means-S variant outperforms the K-Means-C variant. This shows that the number of clusters is an important hyper-parameter and should be tuned for each image. Nevertheless, DiffSeg, despite relying on the same attention tensors  $\mathcal{A}_f$ , significantly outperforms the K-Means baselines. This proves that our algorithm avoids the disadvantages of K-Means and provides better segmentation. Even though DBSCAN does not require the number of clusters as input, DiffSeg significantly outperforms DBSCAN in our experiments. We hypothesize that this is because DiffSeg encodes spatial prior in its clustering process, i.e., spatially distributed anchors. Please refer to Appendix 9.4 for more details on implementing the K-Means and DBSCAN baselines. Moreover, DiffSeg significantly outperforms prior works by an absolute 26% in accuracy and 17% in mIoU compared to the prior SOTA zero-shot method, ReCo on COCO-Stuff-27 for both resolutions (320 and 512).

On the more specialized self-driving segmentation task (Cityscapes), our method is on par with prior works using the smaller 320-resolution input. It outperforms prior works using the larger 512-resolution input in accuracy and mIoU. The resolution of the inputs affects the performance of Cityscapes more severely than that of COCO because Cityscapes has more small classes such as light poles and traffic signs. See discussion on limitations in Appendix 9.7.




Figure 4. Effects of using Different Aggregation Weights ( $R$ ). DiffSeg uses a proportional aggregation strategy to balance consistency and detailedness. Higher-resolution maps produce more detailed but fractured segmentation while lower-resolution maps produce more consistent but coarse segmentation.

Compared to prior works, DiffSeg achieves this level of performance in a pure zero-shot manner without any language dependency or auxiliary images. Therefore, DiffSeg can segment any image.

DiffSeg is a generic clustering algorithm, theoretically applicable to any Transformer based model. We further show the performance of a variant (DinoSeg) using DINO [9] as the backbone in Tab. 1. We observe worse performance due to a smaller attention map size and smaller pre-training dataset. Please refer to Appendix 9.5 for more details and discussion.

## 4.2. Hyper-Parameter Study

Name	COCO	Cityscapes
Aggregation weights ( $R$ )	Propto.	Propto.
Time step ( $t$ )	300	300
Num. of anchors ( $M^2$ )	256	256
Num. of merging iterations ( $N$ )	3	3
KL threshold ( $\tau$ )	1.1	0.9

Table 3. Hyper-parameters for Diffseg. The values are used for producing the results in Tab. 1 and Tab. 2.

There are several hyper-parameters in DiffSeg as listed in Tab. 3. The listed numbers are the exact parameters used in Sec. 4.1. This section provides a sensitivity study for these hyper-parameters to illustrate their expected behavior. Specifically, we show the effects of aggregation weights in the main paper and the impact of other hyper-parameters in Appendix 9.6.

**Aggregation weights ( $R$ ).** The first step of DiffSeg is *attention aggregation*, where attention maps of 4 resolutions are aggregated together. We adopt a proportional aggregation scheme. Specifically, the aggregation weight for a map of a certain resolution is proportional to its resolution, i.e., high-resolution maps are assigned higher importance. This is motivated by the observation that high-resolution maps have a smaller receptive field w.r.t the original image, thus giving more details. To illustrate this, we show the ef-

fects of using attention maps of different resolutions for segmentation in Fig. 4 while keeping other hyper-parameters constant ( $t = 100, M^2 = 256, N = 3, \tau = 1.0$ ). We observe that high-resolution maps, e.g.,  $64 \times 64$  in Fig. 4b, yield the most detailed, however, fractured segmentation. Lower-resolution maps, e.g.,  $32 \times 32$  in Fig. 4c, give more coherent segmentation but often over-segment details, especially along the edges. Finally, too low resolutions fail to generate any segmentation in Fig. 4d because the entire image is merged into one object given the current hyper-parameter settings. Notably, our proportional aggregation strategy (Fig. 4a) balances consistency and detailedness.

## 5. Adding Semantics

While DiffSeg generates high-quality segmentation masks, like SAM [20], it does not label each mask. We propose a simple extension in Appendix 9.8 to produce labeled segmentation masks.

## 6. Visualization

To demonstrate the generalization capability of DiffSeg, we provide examples of segmentation on images of different styles. In Fig. 5 and Fig. 6, we show segmentation on several sketches and paintings. The images are taken from the DomainNet dataset [31]. In Fig. 7, we show the segmentation results of real images captured by a smartphone. We also provide more segmentation examples Cityscapes in Appendix Fig. 17 and SUN-RGBD [44] in Appendix Fig. 15. In Fig. 8, we show segmentation for satellite images and CT scans. Naturally, DiffSeg can segment synthetic images of diverse styles, generated by a stable diffusion model. We show examples in Fig. 9 and more in Appendix Fig. 18.

## 7. Conclusion

Unsupervised and zero-shot segmentation is a very challenging setting and only a few papers have attempted to solve it. Most of the existing work either requires unsupervised adaptation (not zero-shot) or external resources.






Figure 5. Examples of Segmentation on DomainNet Sketch. Overlay (left), input (middle), and segmentation (right)




Figure 7. Examples of Segmentation on real-world images captured by a smartphone. Overlay (left), input (middle), and segmentation (right).




Figure 6. Examples of Segmentation on DomainNet Painting. Overlay (left), input (middle), and segmentation (right)




Figure 8. Segmentation on satellite and brain tumor CT scan.



Figure 9. Examples of Segmentation on generated images. Overlay (left), input (middle), and segmentation (right).

In this paper, we proposed DiffSeg to segment an image without any prior knowledge or external resources, using a pre-trained stable diffusion model, without any additional training. Specifically, the algorithm relies on Intra-Attention Similarity and Inter-Attention Similarity to iteratively merge attention maps into a valid segmentation mask. DiffSeg achieves state-of-the-art performance on popular benchmarks and demonstrates superior generalization to images of diverse styles.

## 8. Acknowledgements

The work was partially done during the first author's internship at Google. This work was also partially supported by ONR grant N00014-18-1-2829 and Google Gift Grant.



## References

- [1] Sklearn dbscan algorithm. <https://scikit-learn.org/stable/modules/generated/sklearn.cluster.DBSCAN.html>. 11
- [2] Sklearn k-means algorithm. <https://scikit-learn.org/stable/modules/generated/sklearn.cluster.KMeans.html>. 11
- [3] Stable diffusion v1-4 model card huggingface. <https://huggingface.co/CompVis/stable-diffusion-v1-4>. 6
- [4] Tomer Amit, Tal Shaharbany, Eliya Nachmani, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. 2
- [5] David Arthur and Sergei Vassilvitskii. K-means++ the advantages of careful seeding. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 1027–1035, 2007. 11
- [6] Dmitry Baranchuk, Ivan Rubachev, Andrey Voynov, Valentin Khruikov, and Artem Babenko. Label-efficient semantic segmentation with diffusion models. *arXiv preprint arXiv:2112.03126*, 2021. 2
- [7] Maxime Bucher, Tuan-Hung Vu, Matthieu Cord, and Patrick Pérez. Zero-shot semantic segmentation. *Advances in Neural Information Processing Systems*, 32, 2019. 3
- [8] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018. 5, 6
- [9] 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, 7, 12
- [10] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2, 2019. 1
- [11] 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, 4, 5, 6
- [12] 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 conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 5, 6
- [13] Martin Ester, Hans-Peter Kriegel, Jörg Sander, Xiaowei Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, pages 226–231, 1996. 6, 11
- [14] Di Feng, Christian Haase-Schütz, Lars Rosenbaum, Heinz Hertlein, Claudius Glaeser, Fabian Timm, Werner Wiesbeck, and Klaus Dietmayer. Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges. *IEEE Transactions on Intelligent Transportation Systems*, 22(3):1341–1360, 2020. 1
- [15] Qianli Feng, Raghudeep Gadde, Wentong Liao, Eduard Ramon, and Aleix Martinez. Network-free, unsupervised semantic segmentation with synthetic images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23602–23610, 2023. 2, 3
- [16] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. *arXiv preprint arXiv:2203.08414*, 2022. 2, 4, 5, 6
- [17] Eric Hedlin, Gopal Sharma, Shweta Mahajan, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. Unsupervised semantic correspondence using stable diffusion. *arXiv preprint arXiv:2305.15581*, 2023. 2
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3
- [19] 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. 2, 5, 6
- [20] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 1, 2, 3, 7, 13
- [21] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 6
- [22] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion model is secretly a zero-shot classifier. *arXiv preprint arXiv:2303.16203*, 2023. 2
- [23] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning*, pages 12888–12900. PMLR, 2022. 13
- [24] Kehan Li, Zhennan Wang, Zesen Cheng, Runyi Yu, Yian Zhao, Guoli Song, Chang Liu, Li Yuan, and Jie Chen. Acseg: Adaptive conceptualization for unsupervised semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7162–7172, 2023. 2, 4, 6
- [25] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. 13
- [26] Huan Ling, Karsten Kreis, Daiqing Li, Seung Wook Kim, Antonio Torralba, and Sanja Fidler. Editgan: High-precision semantic image editing. *Advances in Neural Information Processing Systems*, 34:16331–16345, 2021. 1
- [27] Xiangbin Liu, Liping Song, Shuai Liu, and Yudong Zhang. A review of deep-learning-based medical image segmentation methods. *Sustainability*, 13(3):1224, 2021. 1

- [28] Edward Loper and Steven Bird. Nltk: The natural language toolkit. *arXiv preprint cs/0205028*, 2002. [13](#)
- [29] Timo Lüddecke and Alexander Ecker. Image segmentation using text and image prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7086–7096, 2022. [3](#)
- [30] Tam Nguyen, Maximilian Dax, Chaithanya Kumar Mummadi, Nhung Ngo, Thi Hoai Phuong Nguyen, Zhongyu Lou, and Thomas Brox. Deepusps: Deep robust unsupervised saliency prediction via self-supervision. *Advances in Neural Information Processing Systems*, 32, 2019. [2](#)
- [31] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. [7](#)
- [32] 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. [1](#)
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. [2](#), [3](#)
- [34] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. [2](#), [3](#)
- [35] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *Advances in Neural Information Processing Systems*, 35:25278–25294, 2022. [12](#)
- [36] Gyungin Shin, Weidi Xie, and Samuel Albanie. Reco: Retrieve and co-segment for zero-shot transfer. *Advances in Neural Information Processing Systems*, 35:33754–33767, 2022. [2](#), [3](#), [5](#), [6](#), [12](#)
- [37] Gyungin Shin, Weidi Xie, and Samuel Albanie. Named-mask: Distilling segmenters from complementary foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4960–4969, 2023. [2](#)
- [38] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5463–5474, 2021. [1](#)
- [39] Weijia Wu, Yuzhong Zhao, Mike Zheng Shou, Hong Zhou, and Chunhua Shen. Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. *arXiv preprint arXiv:2303.11681*, 2023. [2](#), [3](#), [13](#)
- [40] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. [1](#)
- [41] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. [2](#), [13](#)
- [42] Andrii Zadaianchuk, Matthaues Kleindessner, Yi Zhu, Francesco Locatello, and Thomas Brox. Unsupervised semantic segmentation with self-supervised object-centric representations. *arXiv preprint arXiv:2207.05027*, 2022. [2](#)
- [43] Junyi Zhang, Charles Herrmann, Junhwa Hur, Luisa Polania Cabrera, Varun Jampani, Deqing Sun, and Ming-Hsuan Yang. A tale of two features: Stable diffusion complements dino for zero-shot semantic correspondence. *arXiv preprint arXiv:2305.15347*, 2023. [2](#)
- [44] Bolei Zhou, Agata Lapedriza, Jianxiong Xiao, Antonio Torralba, and Aude Oliva. Learning deep features for scene recognition using places database. *Advances in neural information processing systems*, 27, 2014. [7](#)
- [45] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from clip. In *European Conference on Computer Vision*, pages 696–712. Springer, 2022. [3](#), [5](#), [6](#)