

MixReorg: Cross-Modal Mixed Patch Reorganization is a Good Mask Learner for Open-World Semantic Segmentation

Kaixin Cai^{1*} Pengzhen Ren^{1*} Yi Zhu² Hang Xu² Jianzhuang Liu²
 Changlin Li³ Guangrun Wang⁴ Xiaodan Liang^{1,5,6 †}

¹Shenzhen Campus of Sun Yat-sen University ²Huawei Noah’s Ark Lab ³University of Technology Sydney

⁴University of Oxford ⁵MBZUAI ⁶DarkMatter AI Research

caikx7@mail2.sysu.edu.cn, pzhren@foxmail.com, {zhuyi36, xu.hang, liu.jianzhuang}@huawei.com,
 {changlinli.ai, wanggrun, xdliang328}@gmail.com

Abstract

Recently, semantic segmentation models trained with image-level text supervision have shown promising results in challenging open-world scenarios. However, these models still face difficulties in learning fine-grained semantic alignment at the pixel level and predicting accurate object masks. To address this issue, we propose **MixReorg**, a novel and straightforward pre-training paradigm for semantic segmentation that enhances a model’s ability to reorganize patches mixed across images, exploring both local visual relevance and global semantic coherence. Our approach involves generating fine-grained patch-text pairs data by mixing image patches while preserving the correspondence between patches and text. The model is then trained to minimize the segmentation loss of the mixed images and the two contrastive losses of the original and restored features. With MixReorg as a mask learner, conventional text-supervised semantic segmentation models can achieve highly generalizable pixel-semantic alignment ability, which is crucial for open-world segmentation. After training with large-scale image-text data, MixReorg models can be applied directly to segment visual objects of arbitrary categories, without the need for further fine-tuning. Our proposed framework demonstrates strong performance on popular zero-shot semantic segmentation benchmarks, outperforming GroupViT by significant margins of 5.0%, 6.2%, 2.5%, and 3.4% mIoU on PASCAL VOC2012, PASCAL Context, MS COCO, and ADE20K, respectively.

1. Introduction

Image segmentation has important applications in scenarios such as virtual presence, virtual try-on, movie post-

*Equal contribution.

†Corresponding author.

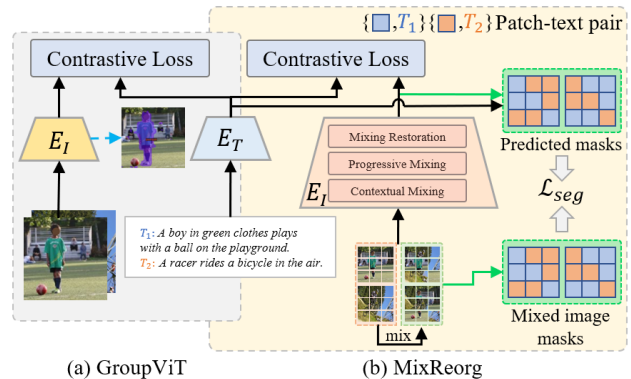


Figure 1: Comparison between GroupViT [38] and MixReorg. (a) GroupViT obtains image segmentation implicitly from image-text pairs to achieve cross-modal semantic alignment. (b) MixReorg explicitly constructs the fine-grained patch-text pairs data from the image-text pairs for free by mixing the patches from different images and preserving the correspondence between patches and text.

production, and autonomous driving. Currently, state-of-the-art semantic segmentation methods [35, 30, 6] benefit from a large number of densely annotated data. However, the assumption of this closed-world setting requires that all categories of objects that appear in the test set are included in the training set. This heavy dependence on annotations limits that they can only work well in closed-set settings. However, considering ubiquitous new concepts in real-world scenarios, learning an open-world segmentation model is more practical, but it is also more challenging. The open-world segmentation model is required to segment all entities and objects class-agnostically and exhaustively during training and be highly generalizable for aligning pixels with new semantics during testing.

Early methods achieve open-world semantic segmentation through few-shot learning [3] or unsupervised clustering [26]. The former actually still assumes that the training and testing classes are in the same latent feature space,



Figure 2: Visual comparison between MixReorg and GroupViT [38] on images from the network. Our method can better handle open-world classes for segmentation task.

while the latter cannot guarantee the consistency of segmentation semantics. Recently, GroupViT [38] achieves state-of-the-art open-world segmentation performance using only text supervision. It realizes the automatic grouping of image patches by vision-language contrastive learning (Figure 1 (a)). ViL-Seg [22] implements image segmentation by introducing additional online clustering of visual embeddings for vision-language contrast. Massive image-text pairs provide rich visual and textual semantics for open-world scenarios. Similar to other CLIP-based [28] vision-language pre-training models (VLM) [32, 41, 19], although these methods achieve local information alignment of different modalities to a certain extent, they are still a computationally-based implicit matching strategy (fine-grained matching is learned by computing patch-text [32] or token-wise [41] similarity matrices). Therefore, *how to learn more fine-grained semantic alignment from image-text pair data becomes a key challenge for text-based supervised open-world segmentation tasks.*

Inspired by the related work of mixed image modeling [27, 21], we propose a simple and novel cross-modal mixed image reconstruction mask learner. Specifically, as shown in Figure 1 (b), MixReorg mixes patches from different images to generate mixed images. Unlike previous methods for jigsaw puzzle [27] or mixed image reconstruction [21] in the single visual modality, MixReorg’s mixed patch reorganization is a cross-modal mask learner designed for semantic segmentation. MixReorg preserves the correspondence between each patch and text when mixing image patches (the legend of Figure 1 (b)). In this way, we can obtain fine-grained patch-text pairs from the image-text data for free.

However, there are still two challenges: (i) the mixed image segmentation is easily disturbed by low-level features, which makes the model unable to realize patch reorganization of mixed images by high-level semantics; (ii) each patch in mixed images are easily interfered by irrelevant patches from different images in the transformer layers, which may cause the image semantics to be difficult to match with the corresponding text.

For the first challenge, we propose two strategies of contextual mixing and progressive mixing to solve this prob-

lem. The contextual mixing strategy allows each patch in the mixed image to obtain the global semantics of its original image in advance by adding a transformer layer before the mixing operation, thereby forcing the model to learn the mixed image reorganization from high-level semantics. Furthermore, to further enhance the global information in the mixed image features, we propose to use the original image features to enhance the global semantics in the mixed image features. For the second challenge, we present a mixing restoration strategy. It guarantees the semantic association of each patch token in the mixed image with the text through contrastive learning between the image recovered from mixed image and the text. In this way, the mutual interference between patches from different images in the mixed image can be effectively suppressed.

In general, MixReorg constructs a set of fine-grained patch-text pairs for free from image-text pair data, and successfully builds a cross-modal mixed image patch reorganization mask learner for open-world segmentation tasks. The proposed MixReorg as a good mask learner also shows strong performance compared with the popular zero-shot semantic segmentation baselines, achieving the performance of 50.5%, 25.4%, 23.6% and 10.1% mIoU on multi-scale evaluations on PASCAL VOC2012, PASCAL Context, MS COCO and ADE20K, respectively. The visualization in Figure 2 shows that MixReorg significantly outperforms GroupViT [38] on open-world segmentation. Our contributions can be summarized as follows:

- We propose a novel and simple method that can easily construct patch-text data with fine-grained matching relationships from image-text data, thereby providing densely supervised information for open-world segmentation.
- For the constructed patch-text data, we propose a cross-modal mixed patch reorganization method. It successfully addresses the challenge of model failure due to mixed image segmentation susceptible to low-level features and irrelevant patches.
- The proposed MixReorg exhibits strong open-world segmentation performance and significantly outperforms current state-of-the-art zero-shot segmentation baselines.

2. Related Work

VLM and Segmentation. Recently, vision-language pre-training models [28, 17] have achieved great success. The models [8, 28, 38, 40] trained with VLM are flexible and versatile, and can adapt to visual [38, 25, 28, 32] and multi-modal [37, 16, 42] upstream and downstream tasks only by using the matching relationship of image-text data. This success has also been found in segmentation and has attracted the attention of lots of researchers [38, 32, 39, 44].

Because traditional semantic segmentation is limited by expensive manual dense annotation, VLMs are expected to break this limitation. Although the above methods achieve promising performance, using image-text sample-level matching relations to learn segmentation masks still faces the challenge of lacking local dense supervision information. In addition, there are some works [41, 32, 19] that explore the alignment of multi-modal local information, but they are still computationally-dominated pseudo-local information correspondence without hard fine-grained supervision from the data level. Therefore, how to obtain finer-grained local supervision information from image-text data through data-level improvement is in great demand for semantic segmentation.

Self-Supervision Strategies. Self-supervision is an effective way to avoid the limitation of expensive manual annotations. It builds a self-supervised pipeline by fully mining the properties of the data itself. For example, self-supervision strategies such as masked image reconstruction [12, 21, 18], jigsaw puzzles [27], multi-view contrast [2, 45, 5] and angle recognition [11] are widely adopted in the vision domain. Similar self-supervision strategies have been widely and successfully applied in natural language processing [7, 34]. But the above methods are all designed for the single modality. In contrast, semantic segmentation not only needs to consider the representation and segmentation of image features but also needs to consider cross-modal semantic alignment. Therefore, how to draw more supervised information from image-text data by borrowing self-supervision strategies from the vision domain is very beneficial for semantic segmentation. Especially for the extraction of cross-modal fine-grained supervision information.

Open-World Segmentation. The open-world problem has been studied in the context of recognition [1, 23], namely how to get a model trained only on a given closed-world dataset to also recognize new classes of objects. Similar settings are also used in object detection [36, 15] and segmentation [3, 26, 38]. For example, [26] proposes an unsupervised open-world semantic segmentation; however, it obtains the mask of the image by a clustering method without any network parameter update. On the other hand, VLMs [28, 19, 8] exhibit strong performance and generalization ability with the help of massive image-text pair data. Inspired by this, TSEG [32] and ICILP [19] attempt to obtain fine-grained semantic alignment from image-text pairs to achieve image segmentation. Similarly, GroupViT [38] introduces a set of learnable group tokens for ViT to group patches and uses the generated segment tokens to align with text embeddings. The massively available image-text pair data provides rich visual and textual semantics for open-world scenarios [40]. Therefore, open-world semantic segmentation based on text supervision can achieve more refined segmentation results at a lower cost of annotation.

Based on the above observations, this paper follows the open-world semantic segmentation setting based on text supervision to further improve the performance of semantic segmentation.

3. Methodology

The overall framework of the MixReorg mask learner is shown in Figure 3. MixReorg is CLIP-based [28] and mainly consists of an image encoder and a text encoder. We use the text encoder from CLIP [28]. MixReorg’s image encoder mainly has three stages: contextual mixing, progressive mixing, and mixing restoration (Sec. 3.1). Then, the loss composition of MixReorg is described in detail in Sec. 3.2. Finally, the total loss is introduced in Sec. 3.3.

3.1. MixReorg

Contextual Mixing. As shown in Figure 3(a), in the contextual image patch mixing stage, patches from different images are randomly mixed to construct a set of mixed images with known segmentation masks. According to the original image-text pairs, the patch-text correspondence of the mixed images is preserved, and the mixed image masks are used as the semantic segmentation labels of the mixed images. Similar to [21], the mixed images only have randomly mixed patch tokens from different images at their same locations. Specially, unlike [21], MixReorg adopts a contextual information image patch mixing strategy. We add a transformer layer before the image patch mixing operation to provide each patch the global image semantic which is closer to the text to create the coherence between patch and text. Meanwhile, it can preliminarily force the model to learn the mixed image patch reorganization from high-level features, thus effectively avoiding the interference of low-level features with the semantic learning of the model.

Specifically, given a batch of image-text pairs $\{(x_i^I, x_i^T)\}_{i=1}^B$. Following the design in ViT [9], we first split each input image into N non-overlapping patches and linearly project each patch into a latent space. These projected patches are denoted as $\{p_i\}_{i=1}^N$. For M image-text pairs $\{(x_i^I, x_i^T)\}_{i=1}^M$, MixReorg randomly mixes the patches from M different images to construct M mixed images, and the corresponding patch composition of M mixed images can be denoted as

$$\text{mix}(\{\{p_i\}_{i=1}^N\}_{m=1}^M) = \{\{p_{m,i}^j\}_{i=1}^N\}_{m=1}^M, 1 \leq j \leq M, \quad (1)$$

where $p_{m,i}^j$ denotes that the i -th patch of the m -th mixed image comes from the j -th image. Correspondingly, we keep the correspondence between each image patch and the text corresponding to the original image, resulting in a semantic segmentation dataset with patch-text correspondence. The patch-text correspondence of the m -th mixed image can be expressed as $\{p_{m,i}^j, x_j^T\}_{i=1}^N$. Therefore, we obtain a set of

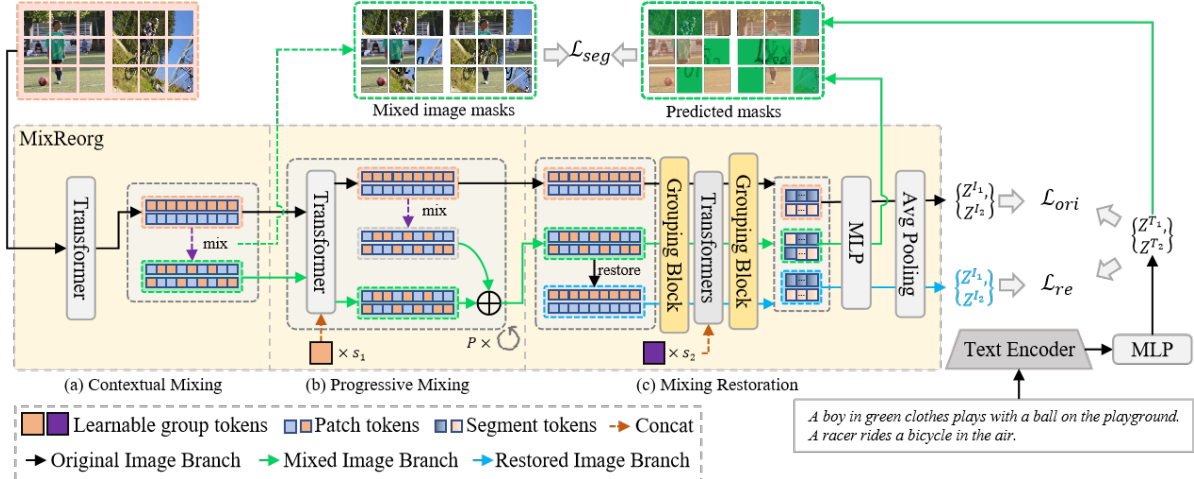


Figure 3: The training pipeline and framework of MixReorg (take two images as an example). MixReorg’s image encoder can be divided into three stages: (a) contextual mixing stage: a set of additional patch-text pairs with known segmentation mask is obtained by randomly mixing contextual patches from different images; (b) progressive mixing stage: the original image features are used to enhance the global information of the mixed image features after mixing; (c) mixing restoration stage: the original features, mixed features, and restored features are segmented through a two-stage grouping block [38], and the corresponding segment tokens are obtained. Note that we omit group tokens in the forward process for simplicity. During testing, MixReorg only needs to execute the original image branch.

sample-level image-text pairs $\{(x_i^I, x_i^T)\}_{i=1}^B$ and a set of patch-level patch-text pairs $\{\{p_{m,i}^j, x_j^T\}_{i=1}^N\}_{m=1}^M$.

Progressive Mixing. Mixed patches’ features cannot improve after Contextual Mixing because of the semantic mixing in mixed images. Since more layers in Contextual Mixing will lead to more parameters, we propose Progressive Mixing to enhance mixed features without additional parameters. As shown in Figure 3(b), in the progressive mixing phase, the patch tokens of normal and mixed images are concatenated with s_1 learnable group tokens $\{g_i\}_{i=1}^{s_1}$ respectively and fed to the multi-layer transformers independently. At the same time, the original features are used to enhance the contextual information of the mixed features. The above process of an original image going through the l -th transformer layer can be represented as

$$\{\{\hat{g}_i\}_{i=1}^{s_1}, \{\hat{p}_i\}_{i=1}^N\} = \text{Trans}_l(\{\{g_i\}_{i=1}^{s_1}; \{p_i\}_{i=1}^N\}), \quad (2)$$

where $[\ ; \]$ denotes the concatenation operator. Similarly, the output of the m -th mixed image through the l -th transformer layer can be expressed as

$$\{\{\hat{g}_i\}_{i=1}^{s_1}, \{\hat{p}_{m,i}^j\}_{i=1}^N\} = \text{Trans}_l(\{\{g_i\}_{i=1}^{s_1}; \{p_{m,i}^j\}_{i=1}^N + \text{mix}(\{\{p_i\}_{i=1}^N\}_{m=1}^M\}l-1)\}). \quad (3)$$

Mixing Restoration. Although we can achieve patch-text alignment by mixed images, there are a lot of different semantics in one mixed image which will interfere with each other. Therefore, this requires the patches from different images in a mixed image still need to maintain a semantic match with the corresponding text.

To this end, as shown in Figure 3(c), in the mixing restoration phase, MixReorg also restores the mixed image

according to the patch position of the image before mixing. The original features $\{p_i\}_{i=1}^N$, mixed features $\{p_{m,i}^j\}_{i=1}^N$, and restored features $re(\{p_{m,i}^j\}_{i=1}^N)$ are segmented through a two-stage grouping block [38], and the corresponding segment tokens $\{\text{seg}_i\}_{i=1}^{s_2}$ are obtained. These segment tokens are fed into multiple transformer layers and then projected to the same dimensionality D as text embeddings $Z^T \in \mathbb{R}^{B \times D}$ through an MLP.

3.2. Cross-Modal Mixed Patch Reorganization Loss

Mixed Segmentation Loss. Cross-modal mixed patch reorganization is a core module designed in MixReorg for semantic segmentation. It provides the model with more refined local alignment information by using a constructed mixed image dataset with patch-text correspondence. We expect the model to learn more fine-grained semantic alignment with the help of the constructed patch-text pairs. And the calculation process of the mixed image segmentation mask is shown in Figure 4, where every B_I (e.g., $B_I = M$) images are mixed. The mixed image segment tokens and the text embeddings of the whole batch are used to compute the similarity $S \in \mathbb{R}^{B_I \times s_2 \times B}$ for the text, where \otimes denotes matrix multiplication. Since we adopt a two-stage grouping block [38], we have the attention map $A \in \mathbb{R}^{B_I \times HW \times s_2}$ ($HW = N$), which contains the grouping relationship between HW patches and s_2 segment tokens. Further, by $A \otimes S$, we can predict the segmentation mask $M_p \in \mathbb{R}^{B_I \times HW \times B}$ of the mixed image. Finally, we compute the cross-entropy loss between the mixed image masks M_{mix} and the prediction masks M_p , formulated as

$$\mathcal{L}_{seg} = \mathcal{L}_{CE}(M_p, M_{mix}), \quad (4)$$




Figure 4: Cross-modal mixed patch reorganization, which combines attention maps and segmentation tokens from the image encoder, and text embeddings to reorganize and predict segmentation masks for mixed images. Where $B_I = M$ means that every B_I images are mixed. For simplicity, we take a mixed image generated by mixing two images as an example.

where $\mathcal{L}_{CE}(\mathbf{p}, \mathbf{q}) = -\sum_i \mathbf{q}_i \log(\mathbf{p}_i)$ is the cross-entropy loss of output \mathbf{p} and target \mathbf{q} .

Restoration Contrastive Loss. Furthermore, we consider cross-modal semantic contrastive learning of mixing restoration features and text embeddings. Specifically, in order to take full advantage of the correspondence provided by the image-text pair and enhance the model’s ability to align semantically across modalities, MixReorg computes the contrastive loss and the multi-label image-text contrastive loss [38] between the output of the restored features branch and the text embeddings, respectively. The calculation of the total image-text contrastive loss is as follows

$$\mathcal{L}_{I \leftrightarrow T}^{re} = \mathcal{L}_{I \rightarrow T}^{re} + \mathcal{L}_{I \leftarrow T}^{re}, \quad (5)$$

where the image-to-text contrastive loss is defined as

$$\mathcal{L}_{I \rightarrow T}^{re} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(z_i^I \cdot z_i^T / \tau)}{\sum_{j=1}^B \exp(z_i^I \cdot z_j^T / \tau)}, \quad (6)$$

and the text-to-image contrastive loss is defined as

$$\mathcal{L}_{I \leftarrow T}^{re} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\exp(z_i^T \cdot z_i^I / \tau)}{\sum_{j=1}^B \exp(z_i^T \cdot z_j^I / \tau)}, \quad (7)$$

where τ is a learnable temperature parameter, and z_i^I and z_i^T are image and text embedding for image-text pairs $\{(x_i^I, x_i^T)\}_{i=1}^B$. In addition, we calculate the multi-label contrastive loss of the restored features branch as follows

$$\mathcal{L}_{I \leftrightarrow \{T_k\}_{k=1}^K}^{re} = \mathcal{L}_{I \rightarrow \{T_k\}_{k=1}^K}^{re} + \mathcal{L}_{I \leftarrow \{T_k\}_{k=1}^K}^{re}, \quad (8)$$

where $\{T_k\}_{k=1}^K$ is an additional K text labels generated using the “prompting engineering” mechanism in [28],

$$\mathcal{L}_{I \rightarrow \{T_k\}_{k=1}^K}^{re} = -\frac{1}{B} \sum_{i=1}^B \log \frac{\sum_{k=1}^K \exp(z_i^I \cdot z_i^{T_k} / \tau)}{\sum_{k=1}^K \sum_{j=1}^B \exp(z_i^I \cdot z_j^{T_k} / \tau)} \quad (9)$$

and

Arch.	Method	Mask mIoU (%)
ViT [9]*	pixel-wise	20.1
	K-means	25.0
	Mean-shift	20.7
	Spectral clustering	19.7
GroupViT [38]	-	41.1
ViewCo [31]	-	45.7
MixReorg (ours)	-	47.9

Table 1: Comparison with zero-shot semantic segmentation baselines on PASCAL VOC. GroupViT[38] and MixReorg are trained on CC12M. The superscript * means the results are from [38].

$$\mathcal{L}_{I \leftarrow \{T_k\}_{k=1}^K}^{re} = -\frac{1}{KB} \sum_{k=1}^K \sum_{i=1}^B \log \frac{\exp(z_i^{T_k} \cdot z_i^I / \tau)}{\sum_{j=1}^B \exp(z_i^{T_k} \cdot z_j^I / \tau)}. \quad (10)$$

The total contrastive loss of the restored features and the text embeddings is as follows

$$\mathcal{L}_{re} = \mathcal{L}_{I \leftrightarrow T}^{re} + \mathcal{L}_{I \leftrightarrow \{T_k\}_{k=1}^K}^{re}. \quad (11)$$

In summary, the total cross-modal mixed image patch reorganization loss is as follows

$$\mathcal{L}_{mixed} = \mathcal{L}_{seg} + \mathcal{L}_{re}. \quad (12)$$

3.3. Overall Loss Function

Similar to Eq. (11), the total contrastive loss between the original image features and the text embeddings is as follows

$$\mathcal{L}_{ori} = \mathcal{L}_{I \leftrightarrow T}^{ori} + \mathcal{L}_{I \leftrightarrow \{T_k\}_{k=1}^K}^{ori}. \quad (13)$$

Finally, the total loss of MixReorg is

$$\mathcal{L} = \mathcal{L}_{mixed} + \mathcal{L}_{ori}. \quad (14)$$

When testing, MixReorg only needs to execute the original image branch (the solid black line in Figure 3), so it does not add any extra testing time.

4. Experiments

4.1. Implementation Details

Architecture. The image encoder of MixReorg is based on a 2-stage GroupViT [38] with 12 transformer layers, while adding one transformer layer before the mix operation. The size of the input image is 224×224 , the patch size is 16×16 and the hidden dimensionality is 384. The model outputs 32 segment tokens (*i.e.*, $s_2 = 32$ and $s_1 = 64$). Following [28, 38], the text encoder of MixReorg consists of 12 layers of transformer with the hidden feature dimensionality of 256.

Training and Inference. During the training phase, we use CC12M [4] as the training dataset, which contains 12M image-text pairs. We apply the mix operation for every 16 images (*i.e.*, $M = 16$). Following [38, 28], our batch size is 4096. We set the weight for each loss function to 1. During the inference phase, only the original image branch is executed. More details can be seen in Appendix.

Arch.	Pre-training			Zero-Shot	Transfer (mIoU (%))	
	Model	Dataset	Supervision		PASCAL VOC	PASCAL Context
ViT	DeiT [33]	ImageNet	class	✗	53.0	35.9
	DINO [2]	ImageNet	self	✗	39.1	20.4
	DINO	CC12M+YFCC	self	✗	37.6	22.8
	MoCo [13]	ImageNet	self	✗	34.3	21.3
	MoCo	CC12M+YFCC	self	✗	36.1	23.0
CLIP	SLIP* [25]	LAION-20M	text & self	✓	-	12.3
	CLIP-MAE* [8]	LAION-20M	text & self	✓	-	16.8
	MaskCLIP [8]	LAION-20M	text & self	✓	-	17.7
	ViewCo [31]	CC12M	text & self	✓	45.7	20.8
	MaskCLIP [44]	CLIP-400M	text	✓	-	21.7
	CLIP* [28]	LAION-20M	text	✓	-	13.5
	GroupViT[38]	CC12M+YFCC	text	✓	51.2	22.3
	GroupViT	CC12M	text	✓	41.1 45.5 [†]	18.2 19.2 [†]
	MixReorg (ours)	CC12M	text	✓	47.9 (6.8 \uparrow) 50.5[†] (5.0 \uparrow)	23.9 (5.7 \uparrow) 25.4[†] (6.2 \uparrow)

Table 2: Performance comparison on PASCAL VOC [10] and PASCAL Context [24]. Zero-shot means that the model is directly transferred to the semantic segmentation task without any fine-tuning on the target dataset. The superscript * denote the results are from [8]. [†] indicates the results of the multi-scale evaluation.

Model	Pre-training Dataset	Transfer mIoU (%)
ViewCo [31]	CC12M	20.6
GroupViT [38]	CC12M+YFCC	20.9
GroupViT	CC12M	18.4 21.1 [†]
MixReorg (ours)	CC12M	21.3 (2.9 \uparrow) 23.6[†] (2.5 \uparrow)

Table 3: Performance comparison on COCO [20]. [†] indicates the results of the multi-scale evaluation.

Open-World Semantic Segmentation. We evaluate the performance of MixReorg on the open-world segmentation task on four commonly used open semantic segmentation datasets PASCAL VOC 2012 [10], PASCAL Context [24], COCO [20], and ADE20K [43]. They contain 20, 59, 80, and 150 foreground classes, respectively, with validation images of 1.5K, 5K, 5K, and 2K, respectively. MixReorg is transferred to the target dataset in a zero-shot manner without any fine-tuning. Following GroupViT [38], MixReorg obtains the corresponding segmentation of the image through the learned group token.

4.2. Comparisons with Existing Methods

Comparison with Zero-Shot Baselines. In Table 1, we present a comparison of four ViT-based baselines, which utilize the image-text contrastive loss defined in CLIP [28] to train the vision and text encoders. These baselines employ pixel-wise, k-means, mean-shift, and spectral clustering strategies, respectively. Additionally, we in-

Model	Pre-training Dataset	Transfer mIoU (%)
ALIGN ^a [14]	ALIGN-1800M	9.7
ALIGN ^a	HQITP-134M	7.5
CLIP ^a [28]	HQITP-134M	5.1
CLIP	CLIP-400M	5.8
CLIP ^b	LAION-20M	7.7
SLIP ^b [25]	LAION-20M	6.8
GroupViT [38]	CC12M	5.8 6.7
MixReorg (ours)	CC12M	8.7 (2.9 \uparrow) 10.1[†] (3.4 \uparrow)

Table 4: Performance comparison on ADE20K [43]. The superscript ^a and ^b denote the results are from [29] and [8], respectively.

clude GroupViT [38], which employs a bottom-up grouping method. The results in Table 1 demonstrate that MixReorg outperforms both the ViT baselines and GroupViT (41.1% vs. 47.9%), indicating that MixReorg is effective in enhancing the segmentation ability of the model.

Comparison with SoTA Methods. We conduct a comprehensive evaluation of MixReorg against fully-supervised methods [33], self-supervised methods [2, 13], vision-language contrastive learning baselines [44, 28, 38], and baselines combining vision-language contrastive learning with self-supervised learning [25, 8, 31]. Table 2 and Table 3 summarizes the comparison results on PASCAL VOC, PASCAL Context, and COCO datasets, both in single-scale and multi-scale evaluation settings. The results show that

MixReorg outperforms all the methods by a significant margin, demonstrating its effectiveness in semantic segmentation. In comparison to GroupViT, MixReorg that pre-trains on CC12M yields substantial performance improvements of 6.8% mIoU, 5.7% mIoU, and 2.9% mIoU for single-scale evaluation, and 5.0% mIoU, 6.2% mIoU, and 2.5% mIoU for multi-scale evaluation on PASCAL VOC, PASCAL Context, and COCO, respectively. MixReorg also outperforms GroupViT which pre-trains on CC12M and YFCC with less data. Additionally, MixReorg has a clear advantage over the methods which rely on additional supervision information in the vision branch. Furthermore, we evaluate the performance of MixReorg on the ADE20K dataset, as shown in Table 4. MixReorg outperforms GroupViT by a significant margin (8.7% vs. 5.8%), demonstrating its superior performance in complex segmentation tasks.

Image Classification. We evaluate the zero-shot classification performance of MixReorg on ImageNet. As shown in Table 5, MixReorg significantly outperforms GroupViT, indicating that MixReorg achieves better image-text alignment through fine-grained mask learning.

Arch.	Zero-shot	
	Acc@1 (%)	Acc@5 (%)
GroupViT [38]	37.5	65.5
MixReorg(ours)	38.8	66.7

Table 5: Zero-shot classification performance on ImageNet.

4.3. Ablation Study

Contextual Mixing. In Table 6, we perform ablation on contextual mixing (CM) strategies. Specifically, we first ablate the parameter variation of the model. As the results show, by adding a transformer layer to GroupViT’s image encoder (*i.e.*, GroupViT+), the performance of the model remains consistent (18.4% vs. 18.2%). This shows that simply increasing the number of parameters does not improve the performance of the model. Further, we also compared MixReorg with only CM (*i.e.*, row 3) and GroupViT+. They have the same amount of parameters, but MixReorg significantly outperforms GroupViT+ with the help of CM (19.3% vs. 18.2%). This shows that it is beneficial to help the model obtain more global semantic information in the early stage of the model. For MixReorg, it can be found that MixReorg without CM assistance (row 4) will degenerate to a similar performance to GroupViT, which indicates that CM plays a crucial role in MixReorg’s mixed segmentation module (19.3% vs. 18.0% vs. 18.4%). This is mainly because CM can help MixReorg acquire global semantic information early in the model, thus forcing the model to learn mixed image reorganization from high-level semantic, which frees the model from low-level semantic information (*e.g.*, texture and color, etc.) to obtain trivial solutions.

Ablation of Losses. In Table 6, we also conduct an ablation study on each loss that MixReorg uses. Specifically, \mathcal{L}_{seg} improves the performance by 1.2% mIoU (20.5% vs.

Method	CM	\mathcal{L}_{seg}	\mathcal{L}_{re}	mIoU (%)
GroupViT	-	-	-	18.4
GroupViT+	-	-	-	18.2
MixReorg	✓			19.3
		✓		18.0
	✓	✓		20.5
	✓		✓	20.3
	✓	✓	✓	21.3

Table 6: On COCO, MixReorg’s ablation study on contextual mixing (CM) and the loss functions. GroupViT+ means adding one transformer layer at the 1-st stage of GroupViT. For MixReorg, without CM means images are mixed prior to passing through the transformer layer that we add.




Figure 5: On COCO, MixReorg’s ablation study on the number of progressive mixings and the number of images for the contextual mixing operation. (a) Yellow line: Ablation study on the number P of the progressive mixing modules. We replace one progressive mixing module with one transformer layer to maintain the model size. (b) Red line: Ablation study on the number M of images for each contextual mixing operation.

19.3%). This means that \mathcal{L}_{seg} plays a significant role. The model learns to distinguish different semantics in the image through \mathcal{L}_{seg} . Additionally, \mathcal{L}_{re} improves the performance of the model (20.3% vs. 19.3%). It helps the patches of the mixed image to maintain consistency between its original image semantics and corresponding text. Furthermore, by combining \mathcal{L}_{seg} and \mathcal{L}_{re} based on CM, the performance of MixReorg can be further improved (21.3% vs. 20.5% vs. 20.3%), which illustrates that CM and two loss functions are strongly related. CM is fundamental to achieving patch-text alignment since it provides global information to each patch, while \mathcal{L}_{seg} provides fine-grained semantic alignment ability and \mathcal{L}_{re} assisting \mathcal{L}_{seg} keeping the patches’ original semantic, free from the interference from different images.

Number of Images for Mixing. In Figure 5 (red line), we observe the performance impact of the number of images M used for contextual mixing operation. It can be observed that $M = 16$ is optimal. As M increases, mixed images contain more semantic categories, which is helpful for the model in learning semantic grouping (20.5% vs. 17.1%). However, increasing M beyond a certain threshold (*e.g.*: $M = 32$) causes semantic representation in the mixed image to be insufficient due to resolution constraints, thereby interfering with model learning (20.5% vs. 18.2%).

Progressive Mixing Module. In Figure 5 (yellow line), we study the number P of the progressive mixing modules. We

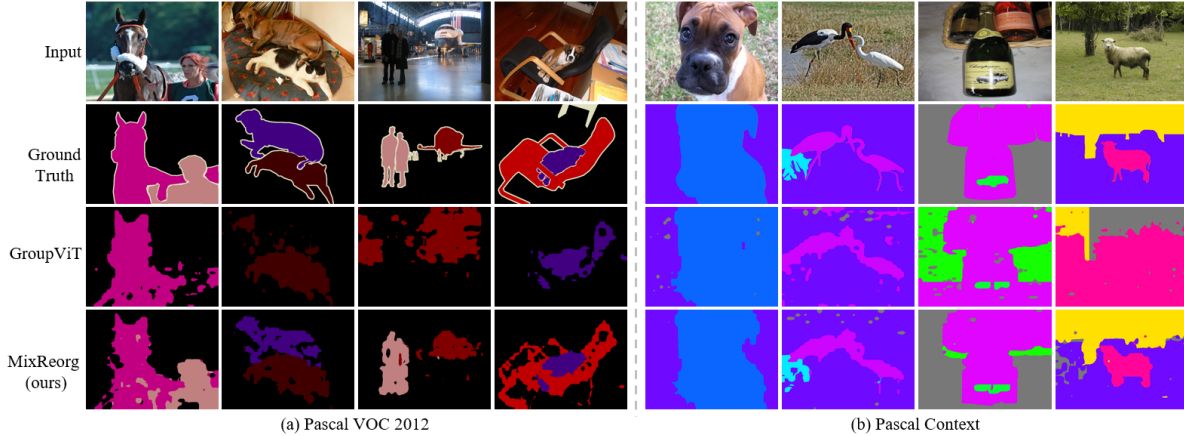


Figure 6: Comparison of semantic segmentation results on PASCAL VOC 2012 and PASCAL Context.

add one transformer layer when removing one progressive mixing module to maintain the model size. It can be seen that the model is optimal when the number P of the progressive mixing modules is 6. The progressive mixing improves over $P = 0$ by about 7% mIoU ($P = 3$ vs. $P = 0$). When $P = 0$, the original image is not used to enhance the mixed image after the mixing operation. In this case, the lack of global information on the original image hinders the learning of the model. Obviously, with the increase of the progressive mixing modules, the semantics of the mixed image features become clearer, which is thus more conducive for model learning to distinguish different semantics, thus improving the model segmentation ability.

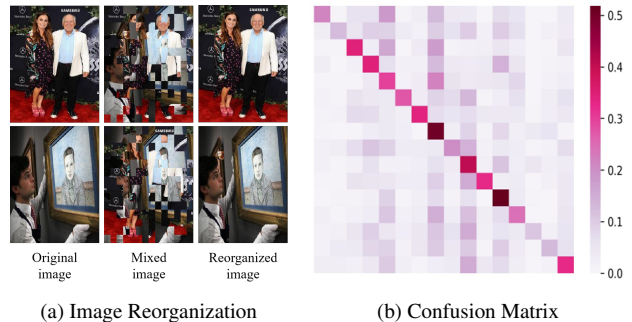
4.4. Visualization

Qualitative Results. In Figure 6, we illustrate zero-shot semantic segmentation examples predicted by GroupViT and MixReorg to verify the segmentation capability of our method. As shown in Figure 6(a), MixReorg can handle more complex segmentation examples which have different classes in one image, showing that our method can better perceive fine-grained semantics. In addition, as shown in Figure 6(b), MixReorg’s segmentation quality of stuff classes is significantly better than GroupViT. In a word, MixReorg has a stronger ability of high-level semantic understanding and segmentation.

Mixed Patch Reorganization. We visualize the reorganized images from mixed images according to the mask predictions from two mixed images in Figure 7(a). It can be seen that except for a few patches, MixReorg can correctly segment most image patches into their corresponding original semantics. In Figure 7(b), the confusion matrix of the prediction for one mixed image indicates that MixReorg can effectively align patches with text.

5. Discussion

Conclusion. We propose a patch-text data construction method with dense matching for image-text data and a cross-modal mixed image patch reorganization mask learner for mixed images to achieve fine-grained semantic



(a) Image Reorganization (b) Confusion Matrix
Figure 7: Mixed image reorganization and the confusion matrix. (a) We use the segmentation mask predicted by MixReorg on the mixed image to obtain the reorganized image. (b) Taking $M = 16$ as an example, the confusion matrix CM of the patch segmentation of the mixed images. CM_{ij} represents the proportion of patches belonging to the i -th image in the mixed image that are classified into the j -th image category.

alignment in open-world segmentation. MixReorg shows superior performance in open-world scenarios.

Limitations. There are two issues that we should explore to improve MixReorg. First, since we use contextual mixing to create additional data, the computational budget is increased during the training phase. Second, although Mixreorg successfully constructs patch-text data for semantic segmentation, there is still a gap between it with pixel-level data.

6. Acknowledgment

This work was supported in part by National Key R&D Program of China under Grant No. 2020AAA0109700, Guangdong Outstanding Youth Fund (Grant No. 2021B1515020061), Shenzhen Science and Technology Program (Grant No. RCYX20200714114642083), Shenzhen Fundamental Research Program (Grant No. JCYJ20190807154211365), Nansha Key RD Program under Grant No.2022ZD014 and Sun Yat-sen University under Grant No. 22lgqb38 and 76160-12220011. We thank MindSpore for the partial support of this work, which is a new deep learning computing framework*.

*<https://www.mindspore.cn/>

References

- [1] Abhijit Bendale and Terrance Boult. Towards open world recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1893–1902, 2015. 3
- [2] 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. 3, 6
- [3] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Deep metric learning for open world semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15333–15342, 2021. 1, 3
- [4] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3558–3568, 2021. 5
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International Conference on Machine Learning*, pages 1597–1607, 2020. 3
- [6] Bowen Cheng, Ishan Misra, Alexander G. Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022. 1
- [7] 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. 3
- [8] Xiaoyi Dong, Yinglin Zheng, Jianmin Bao, Ting Zhang, Dongdong Chen, Hao Yang, Ming Zeng, Weiming Zhang, Lu Yuan, Dong Chen, et al. Maskclip: Masked self-distillation advances contrastive language-image pretraining. *arXiv preprint arXiv:2208.12262*, 2022. 2, 3, 6
- [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*, 2021. 3, 5
- [10] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010. 6
- [11] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *International Conference on Learning Representations*, 2018. 3
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3
- [13] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 6
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc V. Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916, 2021. 6
- [15] KJ Joseph, Salman Khan, Fahad Shahbaz Khan, and Vineeth N Balasubramanian. Towards open world object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5830–5840, 2021. 3
- [16] Wonjae Kim, Bokyung Son, and Ildoo Kim. Vilt: Vision-and-language transformer without convolution or region supervision. In *International Conference on Machine Learning*, pages 5583–5594, 2021. 2
- [17] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *Advances in Neural Information Processing Systems*, 34:9694–9705, 2021. 2
- [18] Xiang Li, Wenhai Wang, Lingfeng Yang, and Jian Yang. Uniform masking: Enabling mae pre-training for pyramid-based vision transformers with locality. *arXiv preprint arXiv:2205.10063*, 2022. 3
- [19] Yi Li, Hualiang Wang, Yiqun Duan, Hang Xu, and Xiaomeng Li. Exploring visual interpretability for contrastive language-image pre-training. *arXiv preprint arXiv:2209.07046*, 2022. 2, 3
- [20] 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 *Proceedings of the European Conference on Computer Vision*, pages 740–755, 2014. 6
- [21] Jihao Liu, Xin Huang, Yu Liu, and Hongsheng Li. Mixmim: Mixed and masked image modeling for efficient visual representation learning. *arXiv preprint arXiv:2205.13137*, 2022. 2, 3
- [22] Quande Liu, Youpeng Wen, Jianhua Han, Chunjing Xu, Hang Xu, and Xiaodan Liang. Open-world semantic segmentation via contrasting and clustering vision-language embedding. In *Proceedings of the European Conference on Computer Vision*, pages 275–292, 2022. 2
- [23] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 3
- [24] Roozbeh Mottaghi, Xianjie Chen, Xiaobai Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 891–898, 2014. 6

- [25] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *Proceedings of the European Conference on Computer Vision*, pages 529–544, 2021. 2, 6
- [26] Yoshikatsu Nakajima, Byeongkeun Kang, Hideo Saito, and Kris Kitani. Incremental class discovery for semantic segmentation with rgbd sensing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 972–981, 2019. 1, 3
- [27] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *Proceedings of the European Conference on Computer Vision*, pages 69–84, 2016. 2, 3
- [28] 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, 2021. 2, 3, 5, 6
- [29] Kanchana Ranasinghe, Brandon McKinzie, Sachin Ravi, Yinfei Yang, Alexander Toshev, and Jonathon Shlens. Perceptual grouping in vision-language models. *arXiv preprint arXiv:2210.09996*, 2022. 6
- [30] Pengzhen Ren, Changlin Li, Guangrun Wang, Yun Xiao, Qing Du, Xiaodan Liang, and Xiaojun Chang. Beyond fixation: Dynamic window visual transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11987–11997, 2022. 1
- [31] Pengzhen Ren, Changlin Li, Hang Xu, Yi Zhu, Guangrun Wang, Jianzhuang Liu, Xiaojun Chang, and Xiaodan Liang. Viewco: Discovering text-supervised segmentation masks via multi-view semantic consistency. In *International Conference on Learning Representations*, 2023. 5, 6
- [32] Robin Strudel, Ivan Laptev, and Cordelia Schmid. Weakly-supervised segmentation of referring expressions. *arXiv preprint arXiv:2205.04725*, 2022. 2, 3
- [33] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357, 2021. 6
- [34] Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for all vision and vision-language tasks. *arXiv preprint arXiv:2208.10442*, 2022. 3
- [35] Wenhui Wang, Jifeng Dai, Zhe Chen, Zhenhang Huang, Zhiqi Li, Xizhou Zhu, Xiaowei Hu, Tong Lu, Lewei Lu, Hongsheng Li, Xiaogang Wang, and Yu Qiao. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023. 1
- [36] Weiyao Wang, Matt Feiszli, Heng Wang, and Du Tran. Unidentified video objects: A benchmark for dense, open-world segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10776–10785, 2021. 3
- [37] Zekun Wang, Wenhui Wang, Haichao Zhu, Ming Liu, Bing Qin, and Furu Wei. Distilled dual-encoder model for vision-language understanding. *arXiv preprint arXiv:2112.08723*, 2021. 2
- [38] 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. 1, 2, 3, 4, 5, 6, 7
- [39] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*, 2021. 2
- [40] Lewei Yao, Jianhua Han, Youpeng Wen, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, Chunjing Xu, and Hang Xu. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. *arXiv preprint arXiv:2209.09407*, 2022. 2, 3
- [41] Lewei Yao, Runhui Huang, Lu Hou, Guansong Lu, Minzhe Niu, Hang Xu, Xiaodan Liang, Zhenguo Li, Xin Jiang, and Chunjing Xu. Filip: Fine-grained interactive language-image pre-training. In *International Conference on Learning Representations*, 2022. 2, 3
- [42] Haotian Zhang, Pengchuan Zhang, Xiaowei Hu, Yen-Chun Chen, Liunian Harold Li, Xiyang Dai, Lijuan Wang, Lu Yuan, Jenq-Neng Hwang, and Jianfeng Gao. Glipv2: Unifying localization and vision-language understanding. *arXiv preprint arXiv:2206.05836*, 2022. 2
- [43] 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. 6
- [44] Chong Zhou, Chen Change Loy, and Bo Dai. Extract free dense labels from CLIP. In *Proceedings of the European Conference on Computer Vision*, pages 696–712, 2022. 2, 6
- [45] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021. 3