

Unsupervised Part Discovery via Descriptor-Based Masked Image Restoration with Optimized Constraints

Jiahao Xia¹, Yike Wu¹, Wenjian Huang², Jianguo Zhang², Jian Zhang^{*,1}

¹Faculty of Engineering and IT, University of Technology Sydney

²Dept. of Comp. Sci. and Eng., Southern University of Science and Technology

Jiahao.Xia-1@uts.edu.au, Yike.Wu@student.uts.edu.au, {huangwj, zhangjg}@sustech.edu.cn, Jian.Zhang@uts.edu.au

Abstract

Part-level features are crucial for image understanding, but few studies focus on them because of the lack of fine-grained labels. Although unsupervised part discovery can eliminate the reliance on labels, most of them cannot maintain robustness across various categories and scenarios, which restricts their application range. To overcome this limitation, we present a more effective paradigm for unsupervised part discovery, named *Masked Part Autoencoder (MPAE)*. It first learns part descriptors as well as a feature map from the inputs and produces patch features from a masked version of the original images. Then, the masked regions are filled with the learned part descriptors based on the similarity between the local features and descriptors. By restoring these masked patches using the part descriptors, they become better aligned with their part shapes, guided by appearance features from unmasked patches. Finally, MPAE robustly discovers meaningful parts that closely match the actual object shapes, even in complex scenarios. Moreover, several looser yet more effective constraints are proposed to enable MPAE to identify the presence of parts across various scenarios and categories in an unsupervised manner. This provides the foundation for addressing challenges posed by occlusion and for exploring part similarity across multiple categories. Extensive experiments demonstrate that our method robustly discovers meaningful parts across various categories and scenarios. The code is available at the project website¹.

1. Introduction

Decomposing objects into meaningful parts drives models to gain better image understanding and further improves the performance of downstream tasks, such as person search [27], fine-grained classification [2, 35] and be-

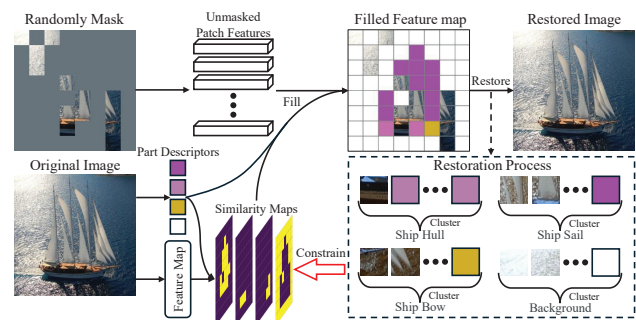


Figure 1. MPAE uses similarity maps between local image features and part descriptors to assign these descriptors to fill masked regions in a randomly masked version of the inputs. The restoration implicitly clusters the features of the restored patches that share similar appearance, further aligning low-level appearance in unmasked patches and high-level semantics of part descriptors. This drives similarity maps to closely match the actual part shapes.

havior analysis [25]. But the high cost of part-level label results in a scarcity of training samples. Consequently, part-level feature learning has received much less attention compared to instance-level feature learning.

To eliminate the reliance on part-level labels, existing methods achieve part discovery mainly through two strategies: reconstruction [12, 19, 21, 23, 32, 40, 41, 43] and clustering [4, 6]. By directly reconstructing the input images using learned part descriptors, the descriptors are expected to align with the part shapes, allowing these part descriptors to further predict part masks. However, to discover meaningful parts, these descriptors are designed to learn high-level semantics rather than low-level appearance [1], which prevents such reconstruction and alignment, especially in complex scenarios. Clustering-based methods use foreground masks to remove irrelevant regions and directly cluster foreground features, regardless of their actual appearance similarity. Some near-outlier features are identified as separate parts, leading to the discovery of unreasonable parts.

To address these limitations and extend unsupervised

*Corresponding Author

¹<https://github.com/Jiahao-UTS/MPAE>

part discovery to more complex scenarios, we propose a novel paradigm, named Mask Part Autoencoder (MPAE). As shown in Fig. 1, instead of directly using part descriptors for image reconstruction, we first calculate the similarities between these descriptors and a feature map learned from the input. Based on the similarity maps, we assign these descriptors to fill the masked regions in a randomly masked version of inputs. By utilizing both filled part descriptors and unmasked patch features within the same part regions to generate image patches with similar appearances, these features are implicitly clustered together [42]. The appearance features of unmasked patches (e.g., boundaries and textures) further align the high-level semantics of the part descriptors with the part shapes. Guided by part appearance, the similarity maps closely match the shapes of corresponding parts even in complex scenarios. Thus, they can be viewed as pixel-level masks of the discovered parts.

Moreover, without using category labels, it is difficult for unsupervised methods to determine which object parts are present or absent in an image, because different categories consist of different types of parts. As a result, most existing methods [12, 19, 23, 41] are restricted to a single category and do not consider the absence of individual parts. However, it is still impossible for all parts to appear in every image, even within a single category, because objects can be occluded. Although [40] achieves part discovery on a dataset with multiple categories, it still relies on category labels to determine which parts appear before part discovery. Therefore, we propose a looser constraint to enable MPAE to identify the presence of parts across various scenarios and categories in an unsupervised manner. Additionally, a more effective constraint is proposed to ensure that the MPAE discovers part-level masks, as the model can easily degrade to predicting only a single instance-level mask when part absence is allowed. This provides a foundation for tackling challenges posed by occlusion. As a result, these optimized constraints allow MPAE to leverage similar parts across different categories to achieve more robust part discovery without relying on any labels.

To evaluate the effectiveness of MPAE and our optimized constraints, extensive experiments are conducted on four widely used benchmarks, including two with multiple categories and two with a single category. The experimental results clearly demonstrate that MPAE robustly discovers meaningful object parts across diverse categories and achieves more competitive performance on all datasets than other state-of-the-art methods.

The main contributions of this paper can be summarized as follows:

- We present a novel paradigm named Masked Part Autoencoder (MPAE), which utilizes low-level appearance features of the unmasked patches to align the part descriptors with actual part shapes. Consequently, MPAE discovers

meaningful parts that closely match the object shapes.

- Several looser yet effective constraints are proposed to enable MPAE to identify the presence of parts across various scenarios and categories in an unsupervised manner. As a result, MPAE can leverage similar parts across various categories to achieve more robust part discovery.
- Extensive experiments are carried out on four widely-used benchmarks. The results show that our method can robustly discover reasonable object parts in complex scenarios and achieve highly competitive performance.

2. Related works

2.1. Unsupervised / Self-supervised Learning

Many works in natural language processing (NLP) [10, 28] show unsupervised / self-supervised learning learns more effective features for downstream tasks than supervised learning. Inspired by their success, recent works [4, 5, 8, 13, 14, 17, 18, 26, 30, 45] also adopt unsupervised / self-supervised learning to improve the effectiveness of visual features. In addition to performance improvements, some new properties also emerge in Vision Transformer (ViT) [11] with the use of unsupervised / self-supervised learning. Caron et al. [4] found using momentum encoder [16] and multi-crop augmentation [3] drives the ViT features from the same category to have high similarity. Based on this property, Wang et al. [39] and Wang et al. [37, 38] achieve foreground segmentation, instance segmentation, and video instance segmentation, respectively, in an unsupervised manner by using normalized cut [31] to directly group self-similar regions. Although learned ViT features within the same part exhibit higher similarity [30, 45], discovering object parts from images is still much more challenging than discovering instances. The main reason is that features from different parts also have very high similarity, resulting in less consistent semantics and boundaries in the discovered parts. Unfortunately, there are relatively few works that focus on using unsupervised / self-supervised features to discover meaningful and consistent object parts.

2.2. Part Discovery

Part discovery aims to identify meaningful object parts with consistent semantics and predict their pixel-level masks. Existing works mainly achieve part discovery through weakly-supervised or unsupervised learning. Weakly-supervised part discovery utilizes instance-level labels, such as category [2, 20, 35] and identity [27, 44], to discover relevant parts with pixel-level masks during training. Then, part-level features are extracted using these masks and further boost the performance of downstream tasks. Nevertheless, without aligning their semantics from part embeddings with actual object appearances through reconstruction, their part masks do not follow the actual object bound-

aries well, making them insufficiently fine-grained.

Unsupervised part discovery is more challenging than weakly-supervised part discovery because no labels are provided during training. Early works discover meaningful object parts with rough part masks by factorizing deep features [7] or spatial embeddings [34]. Other methods [1, 6] directly cluster the foreground features across all images to discover meaningful parts. Nevertheless, the semantics and boundaries of the predicted part masks are not sufficiently stable due to appearance biases in each image. Most state-of-the-art methods [12, 21, 23, 32, 40, 41, 43] employ image reconstruction as a learning objective to align the part descriptors with the actual part shapes. Therefore, these part descriptors can further predict the part masks with consistent semantics. However, these descriptors are designed to capture high-level semantics rather than low-level appearance, which hinders reconstruction and alignment, ultimately degrading part discovery, especially in complex scenarios. Amir et al. [1] demonstrate that directly clustering self-supervised features in foreground regions can also discover meaningful parts. Nevertheless, they only consider feature similarity, disregarding actual appearance similarity. As a result, some near-outlier features are identified as independent parts, leading to unreasonably segmented parts. Therefore, robustly discovering meaningful parts in an unsupervised manner remains a significant challenge.

3. Methodology

Unsupervised part discovery aims to learn a model that can parse all objects into a total of K parts without any labels. Given an input image $I \in \mathbb{R}^{H_1 \times W_1 \times 3}$, the learned model predicts pixel-level mask $M_p \in \{0, 1\}^{H_1 \times W_1 \times (K+1)}$ to assign each pixel into one of K parts or background, where (H_1, W_1) is the size of the input image.

To achieve this goal, we present a novel paradigm, named **Masked Part Autoencoder (MPAE)**, and further incorporate it with several optimized constraints. In what follows, we first detail the training and inference pipeline of MPAE, followed by an introduction to our optimized constraints and their corresponding functions.

3.1. Masked Part Autoencoder

3.1.1. Training pipeline

The overall training pipeline of MPAE is shown in Fig. 2. It consists of four blocks: (a) random masking block, (b) descriptor extraction block, (c) matching block, (d) image restoration block.

Random masking block first divides the input image I into patches of size p . The patchified image I^P is denoted as $I^P \in \mathbb{R}^{\frac{H_1}{p} \times \frac{W_1}{p} \times 3p^2}$. A binary mask $M = \{0, 1\}^{\frac{H_1}{p} \times \frac{W_1}{p}}$ is then generated to randomly mask out a specified proportion r of the image patches. As MAE [18], these masked

patches will not be used in the following training. Only visible patches are finally fed into the MPAE encoder to learn unmasked patch features $F^U \in \mathbb{R}^{\frac{H_1 \times W_1}{p \times p} (1-r) \times C}$, where C is the dimension of the unmasked patch feature.

Descriptor extraction block is designed to learn part-level descriptors from input images. Specifically, as [40], we replace the class embedding of a classical ViT with $K + 1$ embeddings to extract $K + 1$ descriptors. By incorporating MPAE with appropriate constraints, each embedding learns to specialize in one potential part, adaptively aggregating image features to produce the corresponding part descriptor. The $K + 1$ part descriptors $D \in \mathbb{R}^{(K+1) \times C}$ include K descriptors for foreground parts and 1 for the background, where the descriptor dimension C is the same as that of F^U .

To determine how to assign the sparse part descriptors D to fill the masked regions, we employ a frozen pre-trained ViT, followed by a trainable 1×1 convolutional layer, to extract a dense feature map $F \in \mathbb{R}^{H_F \times W_F \times C}$, where (H_F, W_F) is the size of the feature map. By matching F with D , we can obtain a similarity map for each part (as detailed in the following paragraphs). Based on the similarity maps, we can fill the masked patches using D . Moreover, it enables PartFormer to leverage knowledge from the ViT, pre-trained through unsupervised or self-supervised learning, to discover more reasonable object parts.

Matching block calculates similarity maps P between D and F , which can be defined as:

$$P_{i,j,k} = \frac{\exp(D_k \cdot F_{i,j})}{\sum_{k'=1}^{K+1} \exp(D_{k'} \cdot F_{i,j})}, \quad (1)$$

where $F_{i,j}$ indicates the pixel-level feature at (i, j) of F , and D_k is the part descriptor for the k -th part, with $i \in \{1, \dots, W_F\}$, $j \in \{1, \dots, H_F\}$, and $k \in \{1, \dots, K + 1\}$. $P_{i,j,k}$ represents the normalized similarity between the D_k and $F_{i,j}$. Based on the similarity map, the part descriptors are assigned to fill the masked regions and reform a dense feature map R , which can be formulated as:

$$R_{i,j} = \begin{cases} F_{i,j}^U & , M_{i,j} = 0, \\ \sum_{k=1}^{K+1} (P_{i,j,k} D_k) & , M_{i,j} = 1, \end{cases} \quad (2)$$

where $R_{i,j}$ is the feature vector of the filled feature map R at position (i, j) . $M_{i,j}$ is the value of the binary mask M , indicating that the position (i, j) is masked or not in random masking block.

Image restoration block feeds R into the MPAE decoder to restore the masked patches, which can be written as:

$$\mathcal{L}_r = \frac{1}{2} \|I - I'\|_1 + \frac{1}{2} \|\Phi(I) - \Phi(I')\|_1, \quad (3)$$

where I' is the restored image and Φ is a frozen VGG-19 [33]. Directly using the L1 loss constraint for image

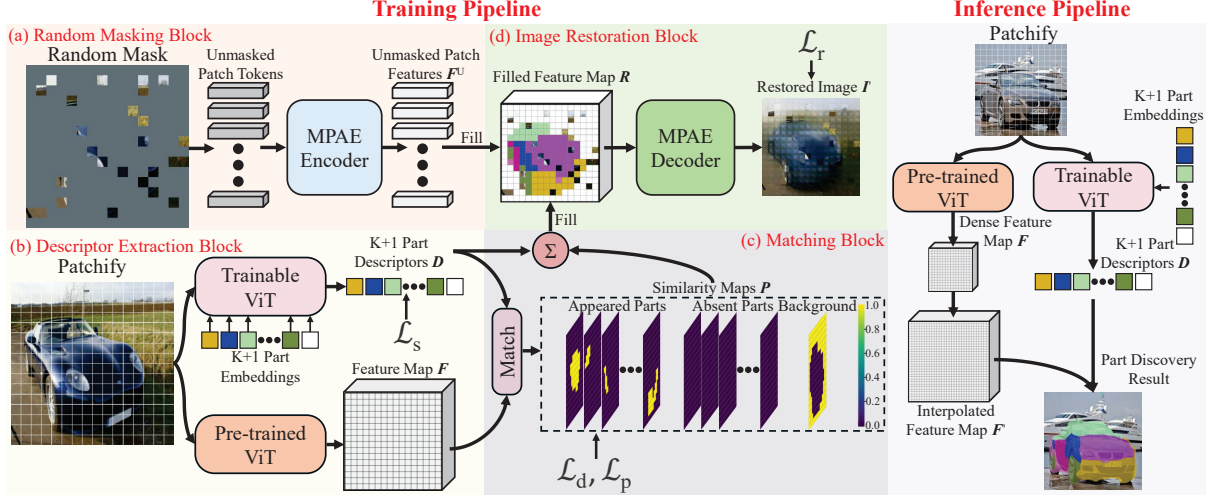


Figure 2. Training and inference pipelines of MPAE. **The training pipeline** consists of four blocks: **(a)** random masking block, **(b)** descriptor extraction block, **(c)** matching block, and **(d)** image restoration block. **(a)** randomly masks out a large proportion of patches and extracts patch features F^U from the unmasked patches using MPAE encoder. The input image I is also fed into **(b)**, after being patchified, to extract part descriptors D and a feature map F . By matching D with F using Hadamard product and a Softmax function, we generate a similarity map for each descriptor in **(c)**. Based on these similarity maps P , we use D to fill the masked regions and reconstruct a feature map R . In **(d)**, R is further utilized for image restoration, aligning the high-level semantics of the D with the corresponding part shapes. **The inference pipeline** directly employs D as mask prototypes to predict part-level masks from the Interpolated feature map F' .

restoration, as in [18], can lead to structure deviations between the raw and restored images [18], which in turn causes the features in different parts to be poorly clustered, regardless of their actual appearance similarity. Therefore, in addition to using the L1 loss, we employ VGG-19 to encourage greater structural similarity between the restored image and the input image as the perceptual loss [22].

The functions of this image restoration are threefold: 1) It retains low-level appearance features in MPAE through the unmasked patch features F^U , thereby compensating for the lack of low-level appearance features in the part descriptors D . 2) It implicitly clusters F^U and D within the same part region together in latent space by utilizing them to produce the image patches with similar appearance. Guided by low-level appearance features from F^U , the high-level semantics from D are aligned more accurately with the actual part shapes compared to directly using D from image reconstruction. 3) This alignment further constrains the similarity maps P , encouraging them to more closely match the actual shapes of their corresponding parts, even in highly complex scenarios. Consequently, P can serve as pixel-level masks for the discovered parts.

3.2. Learning Constraints

To ensure robust part discovery, we incorporate MPAE with three types of optimized constraints: the presence constraint \mathcal{L}_p , the semantic consistency constraint \mathcal{L}_s , and the distribution constraint \mathcal{L}_d .

Presence constraint $\mathcal{L}_p = \mathcal{L}_f + \mathcal{L}_b$, where \mathcal{L}_f and \mathcal{L}_b

represent the foreground and background presence losses respectively. \mathcal{L}_p is directly applied to P to ensure the presence of K foreground parts across the entire data and the presence of background in every image.

\mathcal{L}_f divides each mini-batch into N_g mini-groups with G samples. Then, \mathcal{L}_f can be calculated as:

$$\mathcal{L}_f = \frac{1}{N_g} \sum_{n_g=1}^{N_g} \left(2 - \frac{1}{G} \sum_{g=1}^G \max_{i,j,k} P_{i,j,k}^{g,n_g} - \frac{1}{K} \sum_{k=1}^K \max_{i,j,g} P_{i,j,k}^{g,n_g} \right), \quad (4)$$

where $P_{i,j,k}^{g,n_g}$ is the similarity $P_{i,j,k}$ of g -th sample in n_g -th mini-group. Instead of requiring all parts to appear in each image as previous unsupervised methods, \mathcal{L}_f only encourages them to appear at least once in each mini-group and each image to contain at least one part. This looser constraint provides the basis for MPAE to identify the presence of parts across various scenarios and categories. The influence of the mini-group will be discussed in the **Appendix**.

\mathcal{L}_b encourages the background to appear near the boundaries of each image. Inspired by [2], we define \mathcal{L}_b as:

$$\mathcal{L}_b = -\frac{1}{N_g G} \sum_{n_g=1}^{N_g} \sum_{g=1}^G \log(\max_{i,j} d_{i,j} P_{i,j,K+1}^{g,n_g}), \quad (5)$$

where $d_{i,j}$ is the distance from position (i, j) to the image center. $d_{i,j}$ can be calculated as:

$$d_{i,j} = 2 \left(\frac{i-1}{W_F-1} - \frac{1}{2} \right)^2 + 2 \left(\frac{j-1}{H_F-1} - \frac{1}{2} \right)^2. \quad (6)$$

Incorporated with \mathcal{L}_f and \mathcal{L}_b , our model also learns to distinguish the K foreground parts and the background region.

Semantic constraint \mathcal{L}_s : because \mathcal{L}_p only requires each image to contain at least one part, the model tends to simply segment out the entire foreground object, thereby degrading into an instance-discovery model. By encouraging each part descriptor to align only with the local features that exhibit very high similarity, \mathcal{L}_s ensures that the model decomposes the target into multiple object parts with consistent semantics. \mathcal{L}_s can be formulated as:

$$\mathcal{L}_s = -\frac{1}{K} \sum_{k=1}^K \log \frac{M^k e^{s(\cos(\theta_{(k,k)}+m))}}{M^k e^{s(\cos(\theta_{(k,k)}+m))} + \sum_{t=1, t \neq k}^K M^t e^{s \cos \theta_{(t,k)}}, \quad (7)$$

where $M_k \in \{0, 1\}$ represents the presence of the k -th part. When $\sum_{i=1}^{W_F} \sum_{j=1}^{H_F} P_{i,j,k} \leq 0.001$, $M_k = 0$, else, $M_k = 1$. m and s are two hyperparameters. $\theta_{(k,t)}$ is the angle between the k -th part descriptor \mathbf{D}_k and the mean feature vector of t -th part region $\bar{\mathbf{F}}_t$ in latent space. $\bar{\mathbf{F}}_t$ and $\cos(\theta_{(k,t)})$ are defined as:

$$\bar{\mathbf{F}}_t = \frac{\sum_{i=1}^{W_1} \sum_{j=1}^{H_1} P_{i,j,t} \mathbf{F}_{i,j}}{\sum_{i=1}^{W_1} \sum_{j=1}^{H_1} P_{i,j,t}}, \quad (8)$$

$$\cos(\theta_{(t,k)}) = \frac{\mathbf{D}_k \cdot \bar{\mathbf{F}}_t}{\|\mathbf{D}_k\| \|\bar{\mathbf{F}}_t\|}. \quad (9)$$

As [9], m enables \mathcal{L}_s to still have a certain gradient value for further maximizing $\cos(\theta_{(k,k)})$ when $\theta_{(k,k)}$ is very small, and s controls the smoothness of \mathcal{L}_s . By minimizing $\theta_{(k,k)}$ and maximizing $\theta_{(t,k)}$, the discovered parts across different objects have more consistent semantics. Moreover, the use of M^k enables \mathcal{L}_s to only consider the semantic consistency of those appeared parts, providing the basis to constrain the semantic consistency across multiple categories.

Distribution constraint $\mathcal{L}_d = \mathcal{L}_v + \mathcal{L}_e$, where \mathcal{L}_v and \mathcal{L}_e are total variation loss [29] and entropy loss respectively. \mathcal{L}_d encourages the predicted probability map \mathbf{P} to better follow the shapes of object parts.

\mathcal{L}_v minimizes the spatial gradient of the predicted \mathbf{P} , which can be formulated as:

$$\mathcal{L}_v = \frac{1}{H_F W_F} \sum_{k=1}^{K+1} \sum_{i=1}^{W_F} \sum_{j=1}^{H_F} |\nabla P_{i,j,k}|, \quad (10)$$

where $\nabla P_{i,j,k}$ represents the spatial image gradient at position (i, j) on the probability map for the k -th part. The function of \mathcal{L}_v is the same as that of the concentration loss [21], encouraging pixels with high confidence for the same part to form a connected region. However, the concentration loss penalizes pixels that are far from the part center much more heavily than those that are close. This imbalance further causes the discovered parts to take on an approximately

square shape. \mathcal{L}_v successfully eliminates this imbalance by calculating the gradient of \mathbf{P} , enabling the discovered parts to more accurately follow the shapes of objects, even when they are very complex.

\mathcal{L}_e ensures that the distribution at any position on \mathbf{P} has high similarity in only one part descriptor by penalizing its entropy:

$$\mathcal{L}_e = -\frac{1}{K+1} \sum_{k=1}^{K+1} \sum_{i=1}^{W_F} \sum_{j=1}^{H_F} P_{i,j,k} \log(P_{i,j,k}). \quad (11)$$

This drives the similarity map \mathbf{P} to exhibit clear and stable boundaries, resulting in more precise part discovery.

Overall training target: by minimizing $\mathcal{L} = \mathcal{L}_r + \lambda_p \mathcal{L}_p + \lambda_s \mathcal{L}_s + \lambda_d \mathcal{L}_d$, MPAAE learns to discover meaningful parts in complex scenarios, where λ_p , λ_s , and λ_d are the weights of \mathcal{L}_p , \mathcal{L}_s , \mathcal{L}_d respectively.

3.3. Inference pipeline

The overall inference pipeline is shown on the right in Fig. 2. In this phase, we retain only PartFormer and the pre-trained ViT. However, the low-resolution of \mathbf{F} is insufficient to produce high-resolution part masks. Therefore, we resize \mathbf{F} to match the size of the input images using bilinear interpolation. By accurately aligning with part shapes with the guide of low-level appearance features, \mathbf{D} can still accurately respond to the corresponding part regions in the \mathbf{F}' with much higher resolution. Thus, the similarity maps calculated from \mathbf{D} and \mathbf{F}' can serve as high-resolution part masks for the parts discovered in an unsupervised manner.

4. Experimental Setup

4.1. Datasets

To comprehensively validate the effectiveness of MPAAE, we evaluate it on four widely-used benchmarks: two with multiple categories (PartImageNet OOD dataset [15] and PartImageNet Segmentation dataset [15]) and two with a single category (CUB dataset [36] and CelebA dataset [24]). More details about these four datasets and the implementation of our method are provided in **Appendix A.1** and **Appendix A.2** respectively.

4.2. Evaluation Metrics

Normalized Mutual Information (NMI) and Adjusted Rand Index (ARI) are primarily used to measure the similarity between two clusters. In this task, the NMI and ARI scores between the annotated and discovered parts directly reflect how closely the semantics of the discovered parts align with human conception.

Normalized Mean Error (NME): NME is calculated using the centroids of the predicted parts and the annotated keypoints, representing the semantic consistency of the discovered parts across the entire test set.

Method	Label	K	PartImage_O (%)		PartImage_S (%)	
			NMI \uparrow	ARI \uparrow	NMI \uparrow	ARI \uparrow
Weakly-supervised part discovery						
Huang et al. [20]	Category label	8	5.88	1.53	-	-
		25	7.56	1.25	-	-
		50	10.19	1.05	-	-
PDiscoNet [35]	Category label	8	27.13	8.76	10.68	32.43
		25	32.41	10.69	18.44	40.70
		50	41.49	14.17	21.98	39.40
PDiscoNet + ViT-B/14	Category label	8	19.28	34.72	16.68	32.01
		25	28.23	25.93	46.81	40.70
		50	29.48	28.68	34.13	39.40
PDiscoFormer [2]+ frozen ViT-B/14*	Category label	8	28.84	55.66	-	-
		25	43.36	62.82	-	-
		50	44.48	57.91	-	-
PDiscoFormer + ViT-B/14 \dagger	Category label	8	29.00	52.40	20.29	38.90
		25	44.71	59.27	43.18	64.61
		50	46.29	62.21	44.06	60.10
Unsupervised part discovery						
DINO* [1]	N/A	8	19.17	7.59	-	-
		25	31.46	14.16	-	-
		50	37.81	16.50	-	-
Xia et al.* [40]	N/A	4 N_c	27.85	40.77	25.17	37.50
MPAE* (ours)	N/A	8	32.90	66.17	33.51	65.05
		25	39.28	68.87	37.22	68.04
		50	53.65	74.22	55.10	73.52

Table 1. Quantitative comparisons with state-of-the-art methods on PartImageNet OOD and Segmentation datasets. N_c represents the number of categories, which is 109 and 158 for the OOD and Segmentation datasets respectively. \dagger = using a partially fine-tuned pretrained backbone, * =using a frozen pretrained backbone.

4.3. Comparisons with State-of-the-art Methods

PartImageNet OOD: the quantitative results of MPAE and other state-of-the-art methods on PartImageNet OOD dataset are reported in Table 1, and some randomly chosen qualitative results are shown in Fig. 3. MPAE achieves highly competitive performance among both unsupervised and weakly-supervised methods. In the second row of Fig. 3, we observe that some small, round areas are identified as separate parts by DINO. These unreasonable parts result from directly clustering foreground features, which identifies certain near-outlier features as separate parts. By considering part appearance into clustering through our paradigm, we observe significant improvements in NMI of 71.62%, 24.86% and 41.89% compared to DINO when $K = 8, 25, 50$ respectively. Due to the complex scenarios, it is difficult to align high-level semantics from part descriptors with object shapes through reconstruction. As a result, the parts discovered by Xia et al. are often incomplete. Although PDiscoFormer and PDiscoNet employ classification loss for part discovery in a weakly supervised manner, they do not further align high-level semantics with object shapes. Consequently, their predicted masks do not match objects as closely as those predicted by MPAE. With the same backbone (frozen ViT-B/14), MPAE still outperforms them.

PartImageNet Segmentation increases category number N_c from 109 to 158, which means larger variance in ob-




Figure 3. Some qualitative results of MPAE (ours) and other state-of-the-art unsupervised and weakly-supervised methods on PartImageNet OOD dataset ($K = 50$).

ject appearance and shape. Therefore, as shown in Table 1, the metrics in both NMI and ARI of PDiscoFormer and Xia et al. degrade slightly. However, by using image restoration to constrain the learning and by exploring part similarities across different categories, MPAE successfully maintains strong robustness on this more challenging dataset. Compared to PDiscoFormer, MPAE still achieves an impressive improvement of 65.16% and 25.06% respectively in NMI when $K=8$ and 50, even though PDiscoFormers employ both partially fine-tuned backbone and category labels.

CUB: MPAE also achieves highly competitive performance on CUB (Table 2 and Fig. 4), demonstrating its effectiveness in part discovery within a single category. By implicitly clustering patch features with similar appearances and allowing for the absence of specific parts, MPAE can robustly discover meaningful bird parts in cases involving shadows, self-occlusion, and reflections (see first, second and fourth columns in Fig. 4). With the *same frozen* pretrained ViT-B/14, MPAE outperforms PDiscoFormer in all metrics, despite not using any labels for training. Compared to PDiscoFormer with a partially fine-tuned ViT-B/14, MPAE still outperforms it by 10.45%, 7.66% and 19.22% in NME, NMI and ARI respectively when $K = 16$.

CelebA: as shown in Fig. 4, the masks predicted by MPAE accurately follow the entire facial region, resulting in significantly better qualitative results. Despite less challenging scenarios, the masks predicted by Xia et al. remain incomplete and fail to cover the entire face region. Other methods

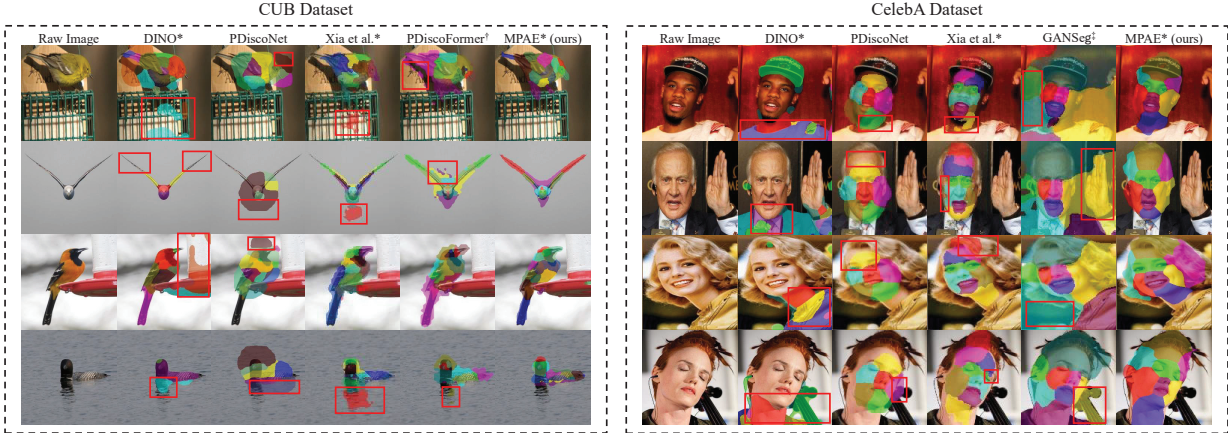


Figure 4. Some visualized part discovery results of MPAE (ours) and state-of-the-art methods on CUB dataset ($K = 16$) and CelebA dataset ($K = 8$). We use red boxes to highlight the regions with poor part discovery results.

Method	Label	$K = 4$ (%)			$K = 8$ (%)			$K = 16$ (%)		
		NME↓	NMI↑	ARI↑	NME↓	NMI↑	ARI↑	NME↓	NMI↑	ARI↑
Weakly-supervised part discovery										
Huang et al. [20]	Species label	11.51	29.74	14.04	11.60	35.72	15.90	12.60	43.92	21.01
PDiscoNet [35]	Species label	9.12	37.82	15.26	8.52	50.08	26.96	7.64	56.87	38.05
Choudhury et al. [6]	Foreground mask	9.20	43.50	19.60	-	51.50	28.30	-	-	-
PDiscoNet + ViT-B/14	Species label	7.70	52.59	26.66	6.34	65.01	37.90	5.95	68.63	43.41
PDiscoFormer + ViT-B/14* [2]	Species label	8.19	52.88	23.22	6.23	67.59	41.35	6.44	69.54	49.99
PDiscoFormer + ViT-B/14†	Species label	7.41	58.13	25.11	6.54	69.87	42.76	5.74	73.38	55.83
Unsupervised part discovery										
DINO* [1]	N/A	-	31.18	11.21	-	47.21	19.76	-	50.57	26.14
Xia et al.* [40]	N/A	9.92	45.22	19.40	8.58	52.42	26.42	8.27	53.52	33.95
MPAE* (ours)	N/A	7.75	58.93	33.73	5.53	68.24	45.85	5.14	79.00	66.56

Table 2. Quantitative comparisons with state-of-the-art methods on CUB. †= using partially fine-tuned backbone, *=using a frozen pre-trained backbone.

Method	Label	$K = 4$ (%)			$K = 8$ (%)			$K = 16$ (%)		
		NME↓	NMI↑	ARI↑	NME↓	NMI↑	ARI↑	NME↓	NMI↑	ARI↑
Weakly-supervised part discovery										
Huang et al. [20]	Attribute label	8.75	56.69	34.74	7.96	54.80	34.74	7.62	62.22	41.01
PDiscoNet [35]	Identity label	11.11	75.97	69.53	9.82	62.61	51.89	9.46	77.43	70.58
Unsupervised part discovery										
DINO* [1]	N/A	11.36	1.38	0.01	10.74	1.12	0.01	-	3.29	0.06
GANSeg‡ [19]	N/A	12.26	41.71	28.06	6.18	67.28	56.23	-	-	-
Xia et al.* [40]	N/A	12.28	50.69	34.57	10.24	58.35	46.78	-	-	-
MPAE* (ours)	N/A	11.44	56.13	37.20	7.55	59.64	41.72	7.10	70.21	60.33

Table 3. Quantitative comparisons with state-of-the-art methods on CelebA. *=using a frozen pretrained backbone, ‡=using a multi-stage self-training strategy.

also misidentify some background regions as foreground parts. However, CelebA provides only five keypoints near the eyes, nose, and mouth for metric calculation, which is insufficient to reflect the improvements achieved by MPAE, as presented in Table 3.

4.4. Ablation Studies

Influence of loss functions: to study the effectiveness of the loss functions used in MPAE, we conduct ablation stud-

ies on PartImageNet Segmentation dataset. The quantitative and qualitative results are shown in Table 4 and Fig. 5 respectively. **Foreground loss** \mathcal{L}_f and **background loss** \mathcal{L}_b are crucial for preventing the model from identifying all pixels as background or foreground. Otherwise, the learned model converges to a trivial solution. Without the prior knowledge provided by **Total variation loss** \mathcal{L}_v , the pixels identified as foreground regions cannot form connected and concentrated parts. **Entropy loss** \mathcal{L}_e leads to more sta-




Figure 5. Qualitative results from the ablation studies of the loss functions on PartImageNet Segmentation dataset ($K = 50$). From left to right: (1) raw images, models trained without (2) \mathcal{L}_f , (3) \mathcal{L}_b , (4) \mathcal{L}_v , (5) \mathcal{L}_e , (6) \mathcal{L}_r , (7) \mathcal{L}_s and (8) nothing (full model).

Method	$K = 8$ (%)		$K = 25$ (%)		$K = 50$ (%)	
	NMI \uparrow	ARI \uparrow	NMI \uparrow	ARI \uparrow	NMI \uparrow	ARI \uparrow
Without \mathcal{L}_f	0.00	0.00	0.00	0.00	0.00	0.00
Without \mathcal{L}_b	30.91	73.20	33.67	67.85	39.36	55.24
Without \mathcal{L}_v	18.30	20.62	13.52	4.54	14.37	4.08
Without \mathcal{L}_e	31.21	63.27	27.48	54.93	52.51	73.08
Without \mathcal{L}_r	25.60	57.39	31.53	63.39	22.32	49.73
Without \mathcal{L}_s	55.69	76.54	57.70	76.97	56.42	74.03
Full model	33.51	65.05	37.22	68.04	55.10	73.52

Table 4. Performance comparisons of MPAE without different loss functions on the PartImageNet Segmentation dataset in the setting of $K = 8, 25, 50$.

Masking Ratio r	50%	75%	80%	85%	90%	95%	100%
NMI (%) \uparrow	45.50	50.11	52.20	52.70	55.10	54.63	43.39
ARI (%) \uparrow	70.25	70.89	70.90	71.85	73.52	73.23	41.29

Table 5. Performance comparisons of MPAE with different masking ratio r on the PartImageNet Segmentation dataset in the setting of $K = 50$.

ble and smooth part boundaries by enforcing that each pixel highly responds to a unique part descriptor. The **absence of reconstruction loss** \mathcal{L}_r results in a failure to align part descriptors with actual part shapes. Consequently, the predicted masks contain only rough regions without clear part boundaries. **Without semantic constraint** \mathcal{L}_s , the model identifies the regions with less semantic similarity as the same part. In some cases, the learned model even degrades into a foreground segmentation model. Although this degradation may improve the numerical metrics in NMI and ARI, the learned model cannot achieve the goal of part discovery. Therefore, all constraints used in MPAE are essential for learning a robust model for unsupervised part discovery.

Influence of the masking ratio r : the quantitative results of MPAE with different masking ratios r are reported in Table 5. When $r < 90\%$, a higher r results in more features in the filled feature map \mathbf{R} comes from part descriptors \mathbf{D} . Therefore, by aligning the part descriptors \mathbf{D}

Change from MPAE model:	NMI (%) \uparrow	ARI (%) \uparrow
$\mathcal{L}_f \rightarrow$ Area loss in [40]	1.90	-0.14
$\mathcal{L}_p \rightarrow$ Presence loss in [2]	42.45	60.24
$\mathcal{L}_d \rightarrow$ Concentration loss in [21]	17.29	18.74
Part number $K \rightarrow 4N_c$	47.59	70.71
MPAE (our best model)	55.10	73.52

Table 6. Comparisons of our optimized losses and designs with previous ones on the PartImageNet Segmentation dataset in the setting of $K = 50$.

with object part shapes, a stronger constraint can be applied to the similarity maps \mathbf{P} , requiring them to better match the actual part shapes. As a results, we observe an improvement in performance. When $r > 90\%$, a higher r results in insufficient unmasked patches containing low-level appearance information to guide the alignment between descriptors and part shapes. Therefore, the optimal masking ratio r is 90%.

Improvements of our optimized constraints and designs: The **area loss in [40]** requires all parts to appear in each image. But, it is impossible for every object to contain all 50 parts. Consequently, \mathcal{L}_a leads to a collapse during training. The image classification loss in [2] ensures that each image to have at least one discovered part, as the embedding for classification cannot be a zero vector. Without any labels, **presence loss in [2]** cannot guarantee that unsupervised methods discover at least one part in each image. Therefore, some images may not have any discovered parts, which further leads to a degradation in metrics. The **concentration loss in [21]** penalizes pixels that are farther from the part center more heavily, which prevents the predicted parts from accurately following their actual shapes, whereas \mathcal{L}_d addresses this imbalance. As [40], we also implement a model that simply decomposes each category into four parts without exploring their semantic similarities ($K \rightarrow 4N_c$). Consequently, the number of training samples for each part significantly decreases, leading to reduced robustness.

More ablation studies and visualized results are included in the **Appendix**.

5. Conclusion

In this paper, we present a novel paradigm (MPAE) with looser yet effective constraints for more robust unsupervised part discovery. By restoring masked image patches using the corresponding part descriptors, MPAE effectively aligns their high-level semantics with the part shapes under the guidance of low-level appearance features from unmasked patches. As a result, MPAE predicts masks that closely match the actual shapes of corresponding parts. The optimized constraints enable MPAE to identify the presence of parts across various scenarios and categories in an unsupervised manner, effectively addressing challenges posed by occlusion and diverse object categories.

References

- [1] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. In *ECCV Workshops*, 2022. 1, 3, 6, 7
- [2] Ananthu Aniraj, Cassio F Dantas, Dino Ienco, and Diego Marcos. Pdiscoformer: Relaxing part discovery constraints with vision transformers. In *ECCV*, 2024. 1, 2, 4, 6, 7, 8
- [3] Mathilde Caron, Ishan Misra, Julien Mairal, Priya Goyal, Piotr Bojanowski, and Armand Joulin. Unsupervised learning of visual features by contrasting cluster assignments. In *NIPS*, pages 9912–9924. Curran Associates, Inc., 2020. 2
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jegou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, pages 9630–9640, 2021. 1, 2
- [5] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, pages 9620–9629, 2021. 2
- [6] Subhabrata Choudhury, Iro Laina, Christian Rupprecht, and Andrea Vedaldi. Unsupervised part discovery from contrastive reconstruction. In *NIPS*, pages 28104–28118, 2021. 1, 3, 7
- [7] Edo Collins, Radhakrishna Achanta, and Sabine Süsstrunk. Deep feature factorization for concept discovery. In *ECCV*, pages 352–368, 2018. 3
- [8] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. In *ICLR*, 2024. 2
- [9] Jiankang Deng, Jia Guo, Jing Yang, Niannan Xue, Irene Kotsia, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. *IEEE TPAMI*, 44(10): 5962–5979, 2022. 5
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, pages 4171–4186, 2019. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 2
- [12] Qingzhe Gao, Bin Wang, Libin Liu, and Baoquan Chen. Unsupervised co-part segmentation through assembly. In *ICML*, pages 3576–3586, 2021. 1, 2, 3
- [13] Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, Bilal Piot, koray kavukcuoglu, Remi Munos, and Michal Valko. Bootstrap your own latent - a new approach to self-supervised learning. In *NIPS*, pages 21271–21284, 2020. 2
- [14] Agrim Gupta, Jiajun Wu, Jia Deng, and Fei-Fei Li. Siamese masked autoencoders. In *NIPS*, pages 40676–40693. Curran Associates, Inc., 2023. 2
- [15] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, Qihang Yu, and Alan Yuille. Partimagenet: A large, high-quality dataset of parts. In *ECCV*, pages 128–145, 2022. 5
- [16] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020. 2
- [17] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, pages 9726–9735, 2020. 2
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 15979–15988, 2022. 2, 3, 4
- [19] Xingzhe He, Bastian Wandt, and Helge Rhodin. Ganseg: Learning to segment by unsupervised hierarchical image generation. In *CVPR*, pages 1215–1225, 2022. 1, 2, 7
- [20] Zixuan Huang and Yin Li. Interpretable and accurate fine-grained recognition via region grouping. In *CVPR*, pages 8659–8669, 2020. 2, 6, 7
- [21] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *CVPR*, pages 869–878, 2019. 1, 3, 5, 8
- [22] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, pages 694–711, 2016. 4
- [23] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. Unsupervised part segmentation through disentangling appearance and shape. In *CVPR*, pages 8351–8360, 2021. 1, 2, 3
- [24] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *ICCV*, pages 3730–3738, 2015. 5
- [25] Xun Long Ng, Kian Eng Ong, Qichen Zheng, Yun Ni, Si Yong Yeo, and Jun Liu. Animal kingdom: A large and diverse dataset for animal behavior understanding. In *CVPR*, pages 19001–19012, 2022. 1
- [26] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel HAZIZA, Francisco Massa, Alaaeldin El-Nouby, Mido Assran, Nicolas Ballas, Wojciech Galuba, Russell Howes, Po-Yao Huang, Shang-Wen Li, Ishan Misra, Michael Rabbat, Vasu Sharma, Gabriel Synnaeve, Hu Xu, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. DINOv2: Learning robust visual features without supervision. *TMLR*, 2024. 2
- [27] Jicheol Park, Dongwon Kim, Boseung Jeong, and Suha Kwak. Plot: Text-based person search with part slot attention for corresponding part discovery. In *ECCV*, pages 474–490, 2025. 1, 2
- [28] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. *OpenAI blog*, 2018. 2
- [29] Leonid I. Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1):259–268, 1992. 5
- [30] Oindrila Saha and Subhransu Maji. Particle: Part discovery and contrastive learning for fine-grained recognition. In *ICCV Workshops*, pages 167–176, 2023. 2

- [31] Jianbo Shi and J. Malik. Normalized cuts and image segmentation. *IEEE TPAMI*, 22(8):888–905, 2000. [2](#)
- [32] Aliaksandr Siarohin, Subhankar Roy, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Motion-supervised co-part segmentation. In *ICPR*, pages 9650–9657, 2021. [1](#), [3](#)
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. [3](#)
- [34] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *ICCV*, pages 3229–3238, 2017. [3](#)
- [35] R. van der Klis, S. Alaniz, M. Mancini, C. F. Dantas, D. Ienco, Z. Akata, and D. Marcos. Pdisconet: Semantically consistent part discovery for fine-grained recognition. In *ICCV*, pages 1866–1876, 2023. [1](#), [2](#), [6](#), [7](#)
- [36] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The Caltech-UCSD Birds-200-2011 Dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. [5](#)
- [37] Xudong Wang, Rohit Girdhar, Stella X. Yu, and Ishan Misra. Cut and learn for unsupervised object detection and instance segmentation. In *CVPR*, pages 3124–3134, 2023. [2](#)
- [38] Xudong Wang, Ishan Misra, Ziyun Zeng, Rohit Girdhar, and Trevor Darrell. Videocutler: Surprisingly simple unsupervised video instance segmentation. In *CVPR*, pages 22755–22764, 2024. [2](#)
- [39] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L. Crowley, and Dominique Vaufreydaz. Self-supervised transformers for unsupervised object discovery using normalized cut. In *CVPR*, pages 14523–14533, 2022. [2](#)
- [40] Jiahao Xia, Wenjian Huang, Min Xu, Jianguo Zhang, Haimin Zhang, Ziyu Sheng, and Dong Xu. Unsupervised part discovery via dual representation alignment. *IEEE TPAMI*, pages 1–18, 2024. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [41] Chang Yu, Xiangyu Zhu, Xiaomei Zhang, Zidu Wang, Zhaoxiang Zhang, and Zhen Lei. Hp-capsule: Unsupervised face part discovery by hierarchical parsing capsule network. In *CVPR*, pages 4022–4031, 2022. [1](#), [2](#), [3](#)
- [42] Qi Zhang, Yifei Wang, and Yisen Wang. How mask matters: Towards theoretical understandings of masked autoencoders. In *NIPS*, pages 27127–27139, 2022. [2](#)
- [43] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *CVPR*, pages 2694–2703, 2018. [1](#), [3](#)
- [44] Kuan Zhu, Haiyun Guo, Tianyi Yan, Yousong Zhu, Jinqiao Wang, and Ming Tang. Pass: Part-aware self-supervised pre-training for person re-identification. In *ECCV*, pages 198–214, 2022. [2](#)
- [45] Adrian Ziegler and Yuki M. Asano. Self-supervised learning of object parts for semantic segmentation. In *ICCV*, pages 14482–14491, Los Alamitos, CA, USA, 2022. [2](#)