

Compositor: Bottom-up Clustering and Compositing for Robust Part and Object Segmentation

Ju He^{1*} Jieneng Chen^{1*} Ming-Xian Lin² Qihang Yu¹ Alan Yuille¹
¹Johns Hopkins University ²Chinese Academy of Sciences

Abstract

*In this work, we present a robust approach for joint part and object segmentation. Specifically, we reformulate object and part segmentation as an optimization problem and build a hierarchical feature representation including pixel, part, and object-level embeddings to solve it in a bottom-up clustering manner. Pixels are grouped into several clusters where the part-level embeddings serve as cluster centers. Afterwards, object masks are obtained by compositing the part proposals. This bottom-up interaction is shown to be effective in integrating information from lower semantic levels to higher semantic levels. Based on that, our novel approach **Compositor** produces part and object segmentation masks simultaneously while improving the mask quality. Compositor achieves state-of-the-art performance on PartImageNet and Pascal-Part by outperforming previous methods by around 0.9% and 1.3% on PartImageNet, 0.4% and 1.7% on Pascal-Part in terms of part and object mIoU and demonstrates better robustness against occlusion by around 4.4% and 7.1% on part and object respectively.*

1. Introduction

Detecting objects and parsing them into semantic parts is a fundamental ability of human visual system. When viewing images, humans not only detect, segment, and classify objects but also segment their semantic parts and identify them. This gives a hierarchical representation that enables a detailed and interpretable understanding of the object which is useful for downstream tasks. For example, humans can estimate the pose of a tiger based on the spatial configuration of its parts and hence judge whether it is about to attack or if it is peacefully sleeping. It is conjectured by cognitive psychologists [3, 27] that these hierarchical representations are constructed in a *bottom-up* manner where humans first perceive parts and then group them together to form objects.

By contrast, the computer vision literature on semantic

segmentation mostly concentrates on object-level, neglecting intermediate part representations, although object and part segmentation have been shown to be mutually beneficial to each other [13, 43]. We emphasize that parts help many other tasks such as pose estimation [11, 46], detection [2, 7], and fine-grained recognition [49]. In addition, exploiting local features or part information can increase robustness of object models against occlusion [1, 25, 40].

Recently, He et al. [20] proposed PartImageNet, where both part and object annotations are provided. Meanwhile, their studies showed that naively using part annotation as deep supervision can improve object segmentation. This motivates us to further design a better interaction pipeline between objects and parts for high-quality segmentation.

In this work, we present a strategy for jointly segmenting parts and objects in a bottom-up process. Specifically, we consider a hierarchical representation of images in terms of pixels, parts, and objects. We learn feature embeddings which enables us to reformulate semantic segmentation as an optimization problem whose goal is to find feature centroids that represent parts and objects. As shown in Figure 1, our method uses a bottom-up strategy where pixels are grouped to form part embeddings which, in turn, are grouped to form object embeddings. We implement this in two steps. First, we cluster image pixels to make proposals for object parts. Here the feature embeddings are learned so that pixels belonging to the same part have similar features. Second, we use a similar approach to compose these part proposals to segment the whole object which involves selecting some part proposals and rejecting others. Our complete algorithm, *Compositor*, for segmenting parts and objects consists of these *clustering* and *compositing* steps. This novel algorithm not only helps us to build a hierarchical segmentation model but also increases the robustness of the model against occlusion since our parts are clustered based on the similarity of pixel features, which are less affected by occlusion compared to other context-based methods. Moreover, objects are constructed using parts that helps minimize the influence of occlusion.

We verify Compositor’s effectiveness on both PartImageNet [20] and Pascal-Part [7], where the former focuses on

*These authors contributed equally to this work.

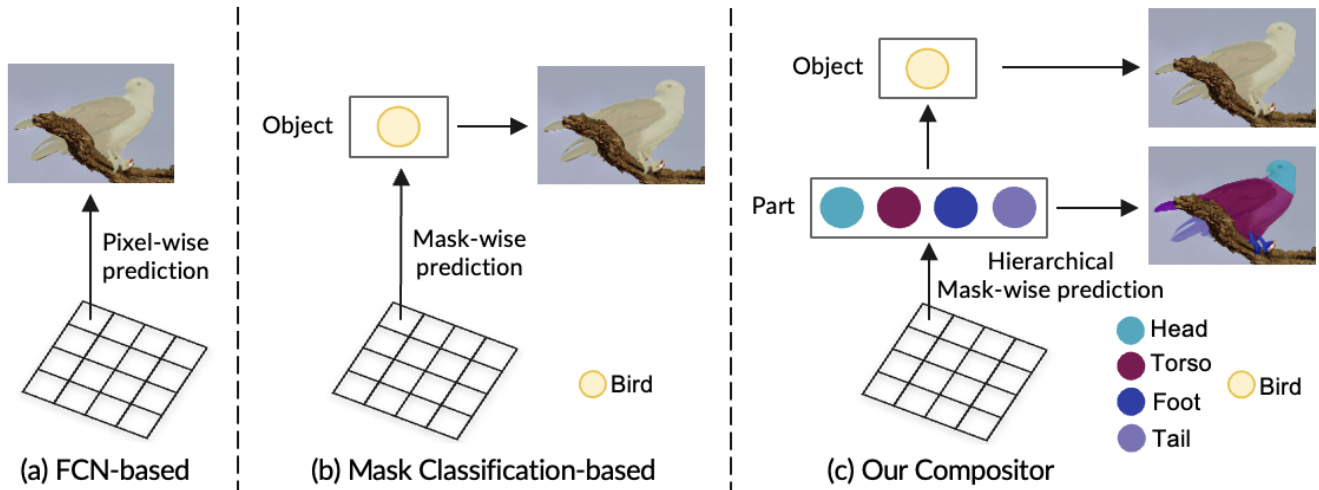


Figure 1. Paradigm comparison among traditional FCN-based method, Mask Classification-based method, and our proposed Composer for object segmentation. We show example with single object instance here for simplicity.

single-instance and the latter contains more multi-instances scenarios. We show that Composer generates high-quality semantic parts from pixels which further benefits object segmentation. Quantitatively, Composer achieves 61.44% and 71.78% mIoU on part and object segmentation with the ResNet-50 [22], outperforming single-task specialized MaskFormer [9] by 1.1% and 1.6% respectively. We get consistent improvement on Pascal-Part by surpassing MaskFormer by 0.4% and 1.7% in terms of part and object mIoU.

We further show the robustness of Composer against occlusion with Occluded-PartImageNet, which is obtained by appending artificial occluders on the original images in PartImageNet following the protocol of OccludedPASCAL3D+ [40]. As a result, Composer outperforms MaskFormer by around 4.4% and 7.1% on part and object mIoU respectively. Ablation studies are conducted to validate the effectiveness of our key designs. Qualitative visualization results on both clean images and occluded images are presented. Error analysis is conducted to better understand the model and guide future work. In summary, we make the following contributions in this work:

1. We propose a bottom-up strategy for segmentation, where we first generate parts from pixels followed by compositing parts into objects. This strategy gives us a joint solution for part and object segmentation.
2. We validate Composer on PartImageNet and Pascal-Part by extensive experiments showing that interactions between parts and objects help each other and result in state-of-the-art performance on both tasks.
3. We create Occluded-PartImageNet by adding occluders enabling us to demonstrate the innate robustness of Composer against occlusion.

2. Related Works

2.1. Object Parsing

Parsing objects into parts is a long-standing problem in computer vision and there is a rich literature on the topic. Pictorial Structure was first proposed in the early 1970's [17]. After that, plenty of different methods [15, 16, 19, 44, 51] have been proposed to explicitly model parts and their spatial relations to the whole object. These methods share a common theme that the object-part models provide rich representations of objects and help interpretability. In the era of deep learning with data-driven models, research on part-based models gets hindered due to the lack of large-scale datasets. Huang et al. [23] proposed a self-supervised co-segmentation method for generating semantically consistent part segmentation results on certain objects. Liu et al. [30] disentangled object appearance and shape information to learn a part segmentation model in an unsupervised manner. However, these works mainly focus on unsupervised part discovery in specific classes instead of finding parts and evaluating them precisely in more general classes.

2.2. Semantic Segmentation

Semantic segmentation has been extensively studied and evaluated on multiple benchmarks. Classic works in semantic segmentation [5, 6, 21, 24] adopt per-pixel classification setting. With the recent progress in transformers [4], a new paradigm named mask classification [9, 41] has been proposed, where segmentation predictions are represented by a set of binary masks with its class label, which is generated through the conversion of object queries to mask embedding vectors followed by multiplying with the image features. The predicted masks are trained by Hungarian matching

with ground truth masks. Thus the essential component of mask transformers is the decoder which takes object queries as input and gradually transfers them into mask embedding vectors. Most of the recent works [8, 9, 38, 41, 42, 50] adopt this setting and the major difference lies in the design of the decoder, while in this work, we propose a novel decoder with *clustering* and *compositing* steps, inspired by the transformer decoder [39] and its variants [8, 28, 32, 47, 48], which gives a joint solution for part and object segmentation.

2.3. Hierarchical learning of objects and parts

Learning objects through the intermediate representation - parts, is a challenging but attractive research topic as it provides a more robust and interpretable understanding of objects. Morabia et al. [37] first proposed to solve part and object detection simultaneously through an attention mechanism. Recently, Ziegler and Asano [52] shows that learning object parts serves as a good pretext task for self-supervised semantic segmentation as it can provide spatially diverse representation. As far as we've concerned, the most related work to ours is Wang et al. [43] on the PASCAL-Part [7], which predicted parts and objects simultaneously. However, in most of these works, object representation still comes through the pixel features, while we show that object representation can be updated based only on part proposals.

2.4. Robustness against occlusion

To ensure good performance in real-world conditions it is crucial to evaluate the robustness of vision algorithms in out-of-distribution scenarios. In particular, as proposed by [40], it is important to test robustness to occluders either by creating a new dataset by superimposing objects onto images or by carefully estimating the amount of occlusion that exists in the dataset. Several works focus on improving robustness through architecture changes. For example, analysis-by-synthesis approaches [1, 25, 35] using compositional networks show much stronger robustness against occlusion based on the generative nature of the model. In this work, we follow the same strategy as [40] to create Occluded-PartImageNet and use it to verify the innate robustness of our model in this out-of-distribution setting.

3. Method

In this section, we first provide a view that regards the object segmentation problem as an optimization problem of learning a feature vector as a centroid for grouping corresponding pixels. Meanwhile, we introduce a hierarchical representation to tackle the problem from part to whole. Afterwards, we illustrate the proposed Compositor, which clusters pixel embedding to part embedding and further composites part embedding to obtain object embedding. In the end, we discuss how we train the model and how to obtain both part and object predictions during the inference.

3.1. Segmentation as Clustering and Compositing

An image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ can be viewed as a set of non-overlapping regions with associated labels:

$$\{y_i\}_{i=1}^M = \{(d_i, c_i)\}_{i=1}^M \quad (1)$$

where $d_i \in \{0, 1\}^{H \times W}$ indicates whether a pixel belongs to the region and c_i denotes the class label of region d_i . M is the number of non-overlapping regions.

Segmentation aims to segment the image into such regions. Two mainstream methods exist to tackle this problem. One directly gives the per-pixel prediction with a prediction head in FCN manner [5, 33]. Another one, inspired by the transformers [39], proposes to disentangle the problem into class-agnostic mask segmentation and mask classification [4, 9, 41]. In this paper, we present a more general view of the problem, by considering the task of segmentation is essentially the same as learning a feature embedding, which can serve as a centroid to group the corresponding pixels together and thus formulates a mask. We first define a bottom-up representation, consisting of pixels, parts, and objects as:

Pixel embedding. Given an image $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$, we use a CNN or Transformer backbone to extract the image features, which are then fed into a FPN [29] to obtain the feature embedding in a higher-resolution. We further add a learnable positional encoding on it and forward it into a MLP layer to obtain the final feature embedding $\mathbf{F} \in \mathbb{R}^{HW \times C}$ for pixels, where C is the feature dimension. Formally, we have $\mathbf{F} = f(\mathbf{I}; \theta)$, where θ is the parameters of the backbone, FPN, and MLP.

Part embedding. Part is a natural intermediate representation for objects [25]. We consider a feature embedding $\mathbf{P} \in \mathbb{R}^{N \times C}$ for object parts with N specifying the number of part proposals, which groups pixels belonging to the same part region. Thus, each \mathbf{P}_i can be viewed as the clustering centroid of the corresponding part pixel features.

Object embedding. Object offers a higher level abstraction of the scene and is usually the most important notion that we care in the scene. Similar to parts, a feature vector $\mathbf{O} \in \mathbb{R}^{M \times C}$ is used to group pixels into object masks.

With the above notations, we reformulate the energy function to minimize for part and object segmentation as:

$$\begin{aligned} \underset{w.r.t \mathbf{F}, \mathbf{P}, \mathbf{O}}{\text{minimize}} & \sum_{i=1}^{HW} \sum_{a=1}^N W_{ia} \cdot \text{Dis}(\mathbf{F}_i, \mathbf{P}_a) \\ & + \sum_{i=1}^{HW} \sum_{b=1}^M W_{ib} \cdot \text{Dis}(\mathbf{F}_i, \mathbf{O}_b), \end{aligned} \quad (2)$$

where $W_{ia}, W_{ib} \in \{0, 1\}$ denote the ground truth mask annotation in a format whether pixel i belongs to the corresponding part a and the object b . Dis is a distance function

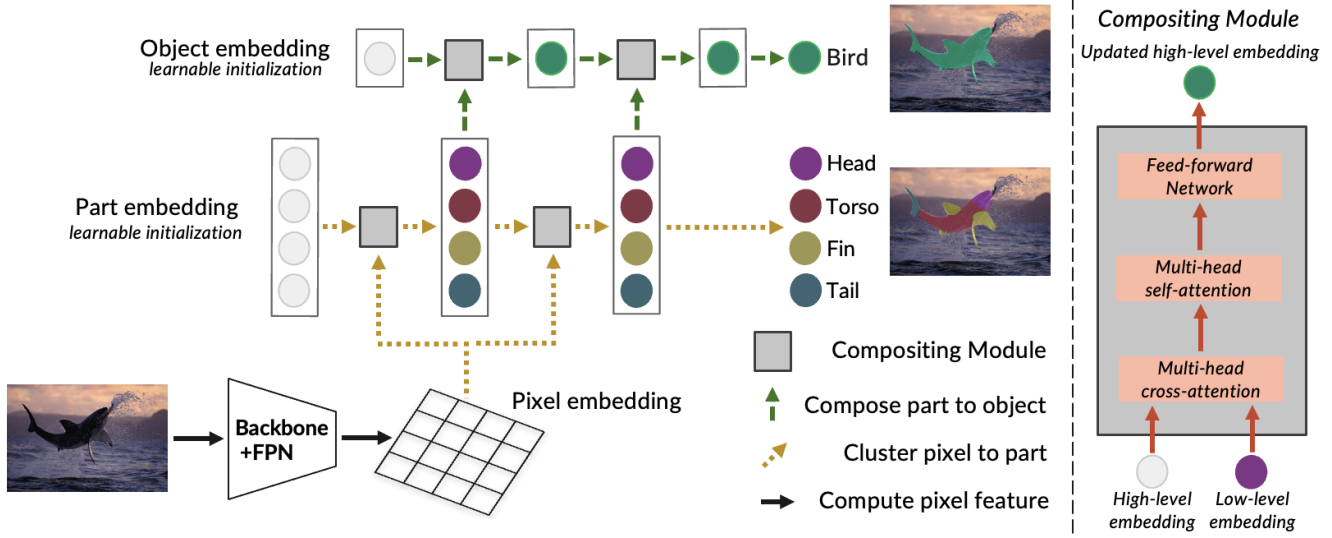


Figure 2. Overview of the proposed framework Compositor, which clusters pixel embedding to part embedding and further composites part embedding to obtain object embedding. Part embedding and object embedding will be iteratively updated and used to predict part/object categories and segmentation maps.

for measuring the similarity between pixel feature vector \mathbf{F} and part feature embedding \mathbf{P} and object feature embedding \mathbf{O} . Note that $\sum_a W_{ia} = 1, \sum_b W_{ib} = 1$ as each pixel is assigned to only one part/object. In short, the problem has been transferred to learning good feature representations for pixels, parts, and objects, which can then be used to perform clustering and compositing for the segmentation problem.

3.2. Compositor: Hierarchical Segmentation from Part to Whole

In prior arts, objects are usually directly derived from pixels while an interpretable intermediate representation is missing. Though [20] proposes a task to handle object parsing and object segmentation at the same time, object parsing is only used as a deep supervision to help the object segmentation. Nonetheless, we note that this strategy does not make full use of the great potential and relationship between the two tasks. Therefore, we propose our hierarchical framework Compositor, which groups pixels into semantic parts first, and then composites them into objects.

Grouping Pixels into Parts. The aim is to decompose an image into a subset of regions where each region is assigned a part class label. Formally, the model decomposes I into disjoint regions \hat{D} , such that $\bigcup_a \hat{d}_a = I$. Each \hat{d}_a is assigned a class label $\hat{c}_a \in \{1, \dots, K\}$ corresponding to either object parts or background.

We obtain such decomposition by computing the similarity between pixel features \mathbf{F} and part centroids \mathbf{P} , thus acquiring the soft-assignment \hat{W} between pixels and corresponding parts. Formally, we estimate the \mathbf{P} and \hat{W} by the

following update rules:

$$\begin{aligned} \hat{W}_{ia} &= \frac{\exp(\mathbf{P}_a \times (\mathbf{F}_i)^T)}{\sum_n \exp(\mathbf{P}_n \times (\mathbf{F}_i)^T)} \\ \mathbf{P}_a &= \sum_i \hat{W}_{ia} \cdot \mathbf{F}_i, \end{aligned} \quad (3)$$

which minimize the the first term $\sum_{i=1}^{HW} \sum_{a=1}^N W_{ia} \cdot \text{Dis}(\mathbf{F}_i, \mathbf{P}_a)$ in Eq. 2 with respect to W and \mathbf{P} separately. We further revise the update rules by computing linear projections of the embedding and utilizing a skip-connection on \mathbf{P} inspired by recent success in Transformers:

$$\begin{aligned} \hat{W}_{ia} &= \frac{\exp(\text{Linear}(\mathbf{P}_a) \times (\text{Linear}(\mathbf{F}_i))^T)}{\sum_n \exp(\text{Linear}(\mathbf{P}_n) \times (\text{Linear}(\mathbf{F}_i))^T)} \\ \mathbf{P}'_a &= \mathbf{P}_a + \sum_i \hat{W}_{ia} \cdot \text{Linear}(\mathbf{F}_i), \end{aligned} \quad (4)$$

where \mathbf{P}'_a is the updated version of part embedding. Note that our part embedding initialization \mathbf{P}_{init} is a set of learnable parameters to encourage the part embedding to have the ability to generate part proposals.

With the above update rules, we exploit the part centroids \mathbf{P}_a to generate the mask regions by applying a threshold on the soft-assignment \hat{W}_{ia} to obtain a hard assignment between pixels and part mask \hat{d}_a . We further learn a MLP (multi-layer perceptron) with parameter ϕ to produce the class label \hat{c}_a for each part mask \hat{d}_a .

Compositing Parts into Objects. Nonetheless, unlike pixels to parts where each pixel will eventually correspond

to some part embeddings, the part embeddings may not all be used to obtain the object embedding (*e.g.*, part centroid \mathbf{P}_a that corresponds to bicycle wheel might not have any mask prediction in an image that only contains animals.). To this end, part to object is more like compositing instead of clustering, where each part may be detected or dropped for the current image. Therefore, we propose to first filter out those undetected parts (based on the previous results of part detection) beforehand to only keep a subset \mathbf{A} of the part centroids \mathbf{P} in order to ease the compositing process. We then estimate the object centroids \mathbf{O} based on \mathbf{A} as:

$$\begin{aligned} \hat{W}_{ab} &= \frac{\exp(\text{Linear}(\mathbf{O}_b) \times (\text{Linear}(\mathbf{A}_a))^T)}{\sum_m \exp(\text{Linear}(\mathbf{O}_m) \times (\text{Linear}(\mathbf{A}_a))^T)} \\ \mathbf{O}'_b &= \mathbf{O}_b + \sum_a \hat{W}_{ab} \cdot \text{Linear}(\mathbf{A}_a) \\ \mathbf{A} &= \{\mathbf{P}_a \mid \mathbf{P}_a \text{ is detected}\}, \end{aligned} \quad (5)$$

which minimize $\sum_{a=1}^N \sum_{b=1}^M W_{ab} \cdot \text{Dis}(\mathbf{P}_a, \mathbf{O}_b)$. As the part embeddings serve as the cluster centroids of the pixels belonging to the parts and the object embeddings are composited from the part ones, this optimization goal essentially achieves a similar effect compared to minimizing the second term in Eq. 2. We obtain our pixel-object hard assignment by multiplying \mathbf{O} with the image features \mathbf{F} followed by applying a threshold to obtain a binary mask. A MLP with parameter ψ is learned to produce the object class label \hat{c}_b for each predicted object mask \hat{d}_b . Object embedding initialization \mathbf{O}_{init} is a set of learnable parameters as well.

In summary, Compositor resolves the object segmentation problem in a hierarchical manner, by clustering pixel embedding into part embedding and then compositing part embedding into object embedding, thus achieving part segmentation simultaneously. The model outputs $\{\mathbf{P}_a, W_{ia}\}$ and $\{\mathbf{O}_b, W_{ab}\}$, which specify the mask regions $\{\hat{d}_a\}$ and $\{\hat{d}_b\}$ with their mask classification labels $\{\hat{c}_a\}$ and $\{\hat{c}_b\}$.

Learning the model. We aim at optimizing the model parameters θ, ϕ, ψ along with \mathbf{P}_{init} and \mathbf{O}_{init} jointly. Given the mask predictions $\{\hat{d}, \hat{c}\}$ and the groundtruth mask $\{d, c\}$ of the image, we compute an assignment V between the predictions and groundtruth, which can be formulated as a correspondence problem with energy:

$$\underset{w.r.t V}{\text{minimize}} \sum_{i,j} V_{ij} \mathcal{L}((d_i, c_i), (\hat{d}_j, \hat{c}_j)), \quad (6)$$

where \mathcal{L} denotes the loss between the predictions and groundtruth. Note that since we typically have much more mask predictions than the groundtruth masks (*i.e.*, $i > j$), we have an additional \emptyset for mapping those unmatched predictions. Except for this \emptyset , we have constraints for Eq. 6 as $\forall i, \sum_j V_{ij} = 1$ & $\forall j, \sum_i V_{ij} = 1$. Inspired by previous works [4, 41], this matching process is implemented

through Hungarian Matching [26] for efficiently computing the matched pairs. We define our loss function between the matched pairs as $\mathcal{L} = \lambda_{ce} \mathcal{L}_{ce} + \lambda_{dice} \mathcal{L}_{dice} + \lambda_{cls} \mathcal{L}_{cls}$, where \mathcal{L}_{ce} and \mathcal{L}_{dice} compute the binary cross-entropy loss and dice loss [36] between d and \hat{d} , \mathcal{L}_{cls} computes the classification loss on the label c, \hat{c} , $\lambda_{ce}, \lambda_{dice}, \lambda_{cls}$ sets the loss weight respectively. Our final loss consists of loss on both parts and objects and can thus be formulated as:

$$\begin{aligned} \mathcal{L} &= \lambda_{ce} \mathcal{L}_{ce_p} + \lambda_{dice} \mathcal{L}_{dice_p} + \lambda_{cls} \mathcal{L}_{cls_p} \\ &+ \beta (\lambda_{ce} \mathcal{L}_{ce_o} + \lambda_{dice} \mathcal{L}_{dice_o} + \lambda_{cls} \mathcal{L}_{cls_o}), \end{aligned} \quad (7)$$

where the subscript p and o denote parts and objects respectively and β is an additional hyper-parameter for controlling the tendency to whether parts or objects of the model.

Post-processing. Ideally, for several categories of objects, each of their corresponding parts should be connected to at least one other object part unless occlusion (*e.g.*, all animal parts should connect). We exploit this weak object knowledge prior as a post-processing protocol to effectively remove portions of false part predictions in the background. To be specific, during inference time, we first search our predicted segmentation mask to find all connected components. We then go over all the components to check if they are connected to other parts and change the label of those isolated ones into background.

4. Experiments

Datasets. We study Compositor on both PartImageNet [20] and Pascal-Part [7]. Both datasets offer large-scale and high-quality per-pixel part annotations on a wide range of objects. To be specific, PartImageNet consists of 158 classes from ImageNet [10] with 24,095 images. Pascal-Part is an additional annotation of VOC [14] which contains 10,103 images from 20 classes. For Pascal-Part, we only consider the 16 classes which have part-level annotations and ignore the rest. We manually merge the provided labels to a higher-level definition of parts (*e.g.* left wing & right wing \rightarrow wing) on Pascal-Part since the original parts are too fine-grained. Note that images in PartImageNet usually only contain one object while Pascal-Part scenes are more complicated with multiple objects. We follow the official train/val/test split and report performance on the val set.

4.1. Implementation Details

Pixel feature extractor. To verify the generality of Compositor across classic and more advanced feature extractor backbones, we experiment with both ResNet-50 [22] and Swin Transformer [31] (tiny variants). For simplicity, a lightweight FPN [24] is adopted as the default pixel feature decoder for further enhancing the pixel embedding. Specifically, we upsample the low-resolution feature map and sum it with the corresponding upper feature map to produce a

Table 1. PartImageNet *val* set results. mIoU, mACC on parts and objects are reported. †: Models on parts and objects are trained separately, ‡: Models on parts and objects are trained jointly, *: Models on objects are trained with parts as deep supervision.

method	backbone	params	Part		Object	
			mIoU	mACC	mIoU	mACC
Deeplab v3+ [†] [6]	ResNet-50 [22]	84M (42M×2)	60.57	71.07	68.38	81.00
Deeplab v3+ [*] [20]	ResNet-50 [22]	42M	-	-	69.82	81.96
MaskFormer [†] [9]	ResNet-50 [22]	90M (45M×2)	60.34	72.75	70.21	81.99
MaskFormer-Dual [‡]	ResNet-50 [22]	50M	58.02	70.42	70.44	81.81
Compositor[‡]	ResNet-50 [22]	50M	61.44	73.41	71.78	83.01
SegFormer [†] [45]	MiT-B2 [45]	48M (24M×2)	61.97	73.77	74.55	85.24
MaskFormer [†] [9]	Swin-T [31]	92M (46M×2)	63.96	77.37	77.92	87.44
MaskFormer-Dual [‡]	Swin-T [31]	51M	61.69	75.64	77.24	87.12
Compositor (Ours)[‡]	Swin-T [31]	51M	64.64	78.31	78.98	87.80

Table 2. Pascal-Part *val* set results. mIoU, mACC on parts and objects are reported. †: Models on parts and objects are trained separately, ‡: Models on parts and objects are trained jointly.

method	backbone	params	Part		Object	
			mIoU	mACC	mIoU	mACC
MaskFormer [†] [9]	ResNet-50 [22]	86M (43M×2)	47.61	58.59	72.69	81.89
MaskFormer-Dual [‡]	ResNet-50 [22]	50M	46.60	57.96	72.13	81.06
Compositor[‡]	ResNet-50 [22]	50M	48.01	58.83	74.35	83.83
MaskFormer [†] [9]	Swin-T [31]	92M (46M×2)	55.42	67.21	81.37	89.29
MaskFormer-Dual [‡]	Swin-T [31]	51M	54.21	66.42	81.02	88.71
Compositor (Ours)[‡]	Swin-T [31]	51M	55.92	67.63	83.10	90.42

multi-scale pixel embedding at output stride 32, 16, and 8 respectively. 1×1 convolution layers are used to guarantee that all per-pixel features share the same dimension.

Clustering and Compositing Steps. The proposed *Clustering* and *Compositing* operations can be easily implemented with multi-head attention as illustrated in Fig. 2. Three continuous proposed *Clustering* and *Compositing* operations are grouped into a block to enhance the model capability. In total, three such blocks are applied to process the multi-scale pixel embeddings. The numbers of part embedding and object embedding (*i.e.*, N, M) are set to be larger than the number of actual parts and objects to offer more proposals to the model. The specific numbers vary from different datasets and we set it to 30/5, 50/20 in default for PartImageNet and Pascal-Part respectively.

Training settings. We instantiate Compositor training setting based on MaskFormer [9]. To be specific, AdamW [34] is adopted with an initial learning rate of 0.0002 and a weight decay of 0.05. The ImageNet-pretrained backbone has a lower learning rate with a multiplier 0.1. We decay the learning rate by a factor of 10 at 0.9 and 0.95 fractions of the total training process. Without an additional statement, we train our models for 50k iterations on PartImageNet and 10k iterations on Pascal-Part with a batch size of 128. We adopt random cropping and large-scale jittering [12, 18] for data augmentations. We set $\lambda_{ce} = 5.0$, $\lambda_{dice} = 5.0$, $\lambda_{cls} = 2.5$ respectively. β is set to be 1/2 as default.

4.2. Main Results

We compare our Compositor with a few task-specialized classic CNN-based segmentation methods as well as more advanced methods with transformer architectures. Besides, we establish a dual-task version of MaskFormer by adding additional separate cross-attention modules and prediction heads on the pixel encoder, which produces both part and object segmentation results from pixels simultaneously.

Table 1 summarizes our experimental results on PartImageNet [20]. With ResNet-50 as the backbone, Compositor achieves 61.44% and 71.78% on part and object mIoU respectively, surpassing the task-specialized MaskFormer by around 1.6%. With a stronger backbone Swin-T, Compositor boosts the performance of mIoU on parts and objects to 64.64% and 78.98%, which outperforms MaskFormer by 0.7% on part and 1.1% on object. However, we discover that the dual-task strategy does not bring help to these tasks, instead, it hurts the performance consistently, especially on part. We argue that this is because the dual-task pipeline ignores the connections between parts and objects, which makes these two tasks no longer beneficial to each other and increases the learning difficulty of the backbone as it needs to provide a good pixel embedding for both tasks.

Table 2 shows our experimental results on the Pascal-Part [7]. Compositor achieves 48.01% and 74.35% in terms of part and object mIoU with ResNet-50. Compositor further increases the performance to 55.92% and 83.10% with



Figure 3. Qualitative comparison for different methods on Pascal-Part. Note that our Composer produces much more accurate part segmentation results with correct labels (e.g., row 2&3) and fewer artifacts (e.g., row 4&5).

Swin-T, which leads MaskFormer by around 1.7%. Note that under multi-instances scenarios, the performance of the dual-branch model still drops by a non-trivial margin.

We also conduct qualitative comparison as illustrated in Figure 3. We show that our Composer is able to parse multiple different objects into corresponding parts even in complicated scenes. In general, Composer produces more accurate part segmentation results with the correct label (e.g., row 2&3) and fewer artifacts (e.g., row 4&5).

4.3. Ablation Studies

Loss ratio of parts to objects. Table 3 summarizes the influence of different loss ratio β regarding the performance on Pascal-Part. Decreasing β at a small range can improve the performance on both tasks as better parts would bring better candidate components for objects thus also improving object segmentation. Moreover, we find that if we increase β which stands for paying more attention to the object side, the part segmentation score drops by a large margin.

Table 3. Ablation study on Table 4. Ablation study on part and object loss ratio β with numbers of queries (i.e., N/M) Swin-T on Pascal-Part.

β	Part	Object	N/M	Part	Object
	mIoU	mIoU		mIoU	mIoU
1/4	55.04	81.97	15/40	54.47	81.74
1/3	55.36	82.65	15/50	54.68	82.03
1/2	55.92	83.10	20/50	55.92	83.10
1	53.38	82.03	25/50	55.30	82.48
2	51.41	81.74	20/60	54.71	82.21

Number of queries. Table 4 shows the influence of the number of queries. Generally speaking, the more queries we have, the more possible proposals for the part and object can be offered. Experimentally, we find that increasing the number of queries within a range boosts the performance while increasing more does not provide additional gain.



Figure 4. Qualitative comparison on Occluded-PartImageNet-v1 where we append out-of-distribution objects as occluders to randomly mask 20% ~ 40% foreground region. Compositor exhibits stronger robustness and largely avoids four kinds of common errors under occlusion: 1) wrong classification of object, 2) missing detection of object, 3) false prediction on background, 4) partial detection of object.

Table 5. Results on Occluded-PartImageNet-v1 where around 20% ~ 40% region of the object is occluded. Numbers in brackets indicate performance drop compared to that on clean images.

method	Part	Object
	mIoU	mIoU
MaskFormer [9]	50.23 (-13.73)	56.73 (-21.19)
Compositor (Ours)	54.63 (-10.01)	63.79 (-15.19)

4.4. Occlusion

To validate the robustness of Compositor, we propose Occluded-PartImageNet following the protocol of OccludedPASCAL3D+ [40]. Specifically, we append artificial occluders to the PartImageNet images. We create three different levels of occlusion based on the occluded ratio of the object. Compositor achieves 54.63% and 63.79% on part and object mIoU on Occluded-PartImageNet-v1, while the performance of MaskFormer drops more and only gets 50.23% and 56.73% in terms of part and object mIoU, which is around 4.4% and 7.1% behind Compositor. Sample visualizations are given in Fig. 4, where we show how Compositor successfully segments the object through parts to avoid typical errors. We hypothesize that the robustness of Compositor comes from both the clustering idea to form parts which depends on the pixel feature similarity and the hierarchical design to composite parts into objects.

4.5. Error Analysis

Since our model not only predicts segmentation maps but also generates part proposals, we could further evaluate it from the perspective of part detection. Specifically, we first compute the overlap of part proposals and the groundtruth part masks. We use 0.5 as the IoU threshold to determine whether a part is detected or not (*i.e.*, detected parts). We examine those undetected parts to see whether our model completely ignores the parts and does not make any valid prediction (*i.e.*, missing parts) or produces part proposals at

Table 6. Compositor part detection error pattern analysis.

Error Type/Dataset	PartImageNet	Pascal-Part
GT Parts	3697	18392
Detected parts	3137	13414
Missing parts	403	4192
Wrong parts	All	786
	Top-3	135

the correct region but with a wrong label (*i.e.*, wrong parts). For those wrong parts, we check whether the correct label lies in the top-3 of the classification prediction (*i.e.*, Top-3). As can be seen from Table 6, our model successfully detects around 85% and 73% of parts on PartImageNet and Pascal-Part respectively. For the undetected parts, our model has valid part proposals for one fourth of them. Note that among these part proposals with wrong labels, more than 80% have the correct label within top-3 prediction.

5. Conclusion

In this work, we present a bottom-up algorithm *Compositor*, for jointly segmenting parts and objects by estimating part proposals from pixels which are then composed into objects. Compositor formulates semantic segmentation as an optimization problem which involves learning feature embeddings so that parts and objects can be represented by centroids in this embedding space. They are estimated, respectively, by *Clustering* and *Compositing* lower semantic level embeddings. Compositor achieves state-of-art performance on both part and object segmentation on PartImageNet and Pascal-Part and shows stronger robustness against occlusion on our proposed Occluded-PartImageNet. Error analysis is provided to better understand our model.

Acknowledgements. The authors gratefully acknowledge supports from ONR with award N00014-21-1-2812 and IAA with award 80052272.

References

- [1] Wang Angtian, Adam Kortylewski, and Alan Yuille. Nemo: Neural mesh models of contrastive features for robust 3d pose estimation. In *Proceedings International Conference on Learning Representations (ICLR)*, 2021. 1, 3
- [2] Hossein Azizpour and Ivan Laptev. Object detection using strongly-supervised deformable part models. In *European Conference on Computer Vision*, pages 836–849. Springer, 2012. 1
- [3] Irving Biederman. Recognition-by-components: a theory of human image understanding. *Psychological review*, 94(2):115, 1987. 1
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 2, 3, 5
- [5] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. 2, 3
- [6] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-decoder with atrous separable convolution for semantic image segmentation. In *Proceedings of the European conference on computer vision (ECCV)*, pages 801–818, 2018. 2, 6
- [7] Xianjie Chen, Roozbeh Mottaghi, Xiaobai Liu, Sanja Fidler, Raquel Urtasun, and Alan Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1971–1978, 2014. 1, 3, 5, 6
- [8] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. *arXiv preprint arXiv:2112.01527*, 2021. 3
- [9] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 2, 3, 6, 8
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 5
- [11] Jian Dong, Qiang Chen, Xiaohui Shen, Jianchao Yang, and Shuicheng Yan. Towards unified human parsing and pose estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 843–850, 2014. 1
- [12] Xianzhi Du, Barret Zoph, Wei-Chih Hung, and Tsung-Yi Lin. Simple training strategies and model scaling for object detection. *arXiv preprint arXiv:2107.00057*, 2021. 6
- [13] S Eslami and Christopher Williams. A generative model for parts-based object segmentation. *Advances in Neural Information Processing Systems*, 25, 2012. 1
- [14] 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. 5
- [15] Li Fei-Fei, Rob Fergus, and Pietro Perona. One-shot learning of object categories. *IEEE transactions on pattern analysis and machine intelligence*, 28(4):594–611, 2006. 2
- [16] Pedro F Felzenszwalb and Daniel P Huttenlocher. Pictorial structures for object recognition. *International journal of computer vision*, 61(1):55–79, 2005. 2
- [17] Martin A Fischler and Robert A Elschlager. The representation and matching of pictorial structures. *IEEE Transactions on computers*, 100(1):67–92, 1973. 2
- [18] Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2918–2928, 2021. 6
- [19] Ross Girshick, Pedro Felzenszwalb, and David McAllester. Object detection with grammar models. *Advances in neural information processing systems*, 24, 2011. 2
- [20] Ju He, Shuo Yang, Shaokang Yang, Adam Kortylewski, Xiaoding Yuan, Jie-Neng Chen, Shuai Liu, Cheng Yang, and Alan Yuille. Partimagenet: A large, high-quality dataset of parts. *arXiv preprint arXiv:2112.00933*, 2021. 1, 4, 5, 6
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 2
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 2, 5, 6
- [23] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. pages 869–878, 2019. 2
- [24] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019. 2, 5
- [25] Adam Kortylewski, Ju He, Qing Liu, and Alan L Yuille. Compositional convolutional neural networks: A deep architecture with innate robustness to partial occlusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8940–8949, 2020. 1, 3
- [26] Harold W Kuhn. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97, 1955. 5
- [27] Brenden M Lake, Ruslan Salakhutdinov, and Joshua B Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338, 2015. 1
- [28] Xia Li, Zhisheng Zhong, Jianlong Wu, Yibo Yang, Zhouchen Lin, and Hong Liu. Expectation-maximization attention networks for semantic segmentation. pages 9167–9176, 2019. 3

- [29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3
- [30] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. Unsupervised part segmentation through disentangling appearance and shape. pages 8355–8364, 2021. 2
- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10012–10022, 2021. 5, 6
- [32] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33:11525–11538, 2020. 3
- [33] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 3
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [35] Wufei Ma, Angtian Wang, Alan Yuille, and Adam Kortylewski. Robust category-level 6d pose estimation with coarse-to-fine rendering of neural features. In *European Conference on Computer Vision*, pages 492–508. Springer, 2022. 3
- [36] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *2016 fourth international conference on 3D vision (3DV)*, pages 565–571. IEEE, 2016. 5
- [37] Keval Morabia, Jatin Arora, and Tara Vijaykumar. Attention-based joint detection of object and semantic part. *arXiv preprint arXiv:2007.02419*, 2020. 3
- [38] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7262–7272, 2021. 3
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 3
- [40] Angtian Wang, Yihong Sun, Adam Kortylewski, and Alan L Yuille. Robust object detection under occlusion with context-aware compositionalnets. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12645–12654, 2020. 1, 2, 3, 8
- [41] Huiyu Wang, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Max-deeplab: End-to-end panoptic segmentation with mask transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5463–5474, 2021. 2, 3, 5
- [42] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *European Conference on Computer Vision*, pages 108–126. Springer, 2020. 3
- [43] Peng Wang, Xiaohui Shen, Zhe Lin, Scott Cohen, Brian Price, and Alan L Yuille. Joint object and part segmentation using deep learned potentials. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1573–1581, 2015. 1, 3
- [44] Markus Weber, Max Welling, and Pietro Perona. Unsupervised learning of models for recognition. In *European conference on computer vision*, pages 18–32. Springer, 2000. 2
- [45] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34, 2021. 6
- [46] Yi Yang and Deva Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *CVPR 2011*, pages 1385–1392. IEEE, 2011. 1
- [47] Qihang Yu, Huiyu Wang, Dahun Kim, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Cmt-deeplab: Clustering mask transformers for panoptic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2560–2570, 2022. 3
- [48] Qihang Yu, Huiyu Wang, Siyuan Qiao, Maxwell Collins, Yukun Zhu, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. k-means mask transformer. In *European Conference on Computer Vision*, pages 288–307. Springer, 2022. 3
- [49] Ning Zhang, Jeff Donahue, Ross Girshick, and Trevor Darrell. Part-based r-cnns for fine-grained category detection. In *European conference on computer vision*, pages 834–849. Springer, 2014. 1
- [50] Wenwei Zhang, Jiangmiao Pang, Kai Chen, and Chen Change Loy. K-net: Towards unified image segmentation. *Advances in Neural Information Processing Systems*, 34, 2021. 3
- [51] Song-Chun Zhu and David Mumford. *A stochastic grammar of images*. Now Publishers Inc, 2007. 2
- [52] Adrian Ziegler and Yuki M Asano. Self-supervised learning of object parts for semantic segmentation. *arXiv preprint arXiv:2204.13101*, 2022. 3