

Unsupervised Part Segmentation through Disentangling Appearance and Shape

Shilong Liu¹, Lei Zhang², Xiao Yang¹, Hang Su¹, Jun Zhu^{1*}

¹ Dept. of Comp. Sci. and Tech., BNRist Center, Institute for AI, Tsinghua-Bosch Joint ML Center

¹ Tsinghua University, Beijing, 100084, China ² Microsoft Corporation

{liusl20, yangxiao19}@mails.tsinghua.edu.cn, leizhang@microsoft.com, {suhangss, dcszj}@mail.tsinghua.edu.cn

Abstract

We study the problem of unsupervised discovery and segmentation of object parts, which, as an intermediate local representation, are capable of finding intrinsic object structure and providing more explainable recognition results. Recent unsupervised methods have greatly relaxed the dependency on annotated data which are costly to obtain, but still rely on additional information such as object segmentation mask or saliency map. To remove such a dependency and further improve the part segmentation performance, we develop a novel approach by disentangling the appearance and shape representations of object parts followed with reconstruction losses without using additional object mask information. To avoid degenerated solutions, a bottleneck block is designed to squeeze and expand the appearance representation, leading to a more effective disentanglement between geometry and appearance. Combined with a self-supervised part classification loss and an improved geometry concentration constraint, we can segment more consistent parts with semantic meanings. Comprehensive experiments on a wide variety of objects such as face, bird, and PASCAL VOC objects demonstrate the effectiveness of the proposed method.

1. Introduction

Object parts and landmarks are two widely used intermediate representations of object local structures, which have received increasing attention recently for their robustness to the variations of viewpoint and appearance [12]. While there have been many related works on object landmark detection by either supervised learning [41, 25] or unsupervised discovery [13, 40], semantic part segmentation is relatively less studied [12] due to the costly efforts required for data annotation.

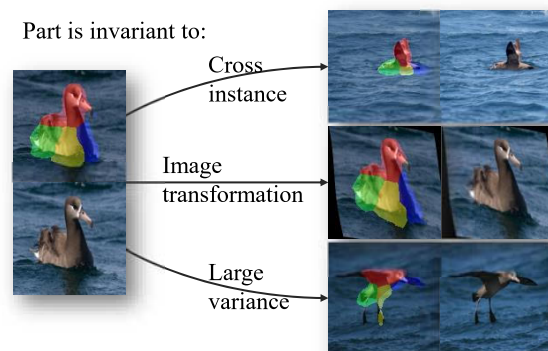


Figure 1: Semantic consistency of object parts. Segmentation regions of the same part should be semantically consistent across object instances and robust to appearance and shape changes. Best view in colors.

This paper primarily focuses on unsupervised semantic part segmentation, as object parts provide richer information about intrinsic object structures which are complementary to landmark points. To alleviate the efforts in annotations, we investigate the unsupervised learning strategy for automatic semantic part segmentation [12]. It has also shown an encouraging impact on single-view 3D object reconstruction by using self-supervised learned semantic parts which enforces semantic consistency between the reconstructed 3D mesh and original 2D images [18].

A good object part segmentation should satisfy two essential constraints on *semantics* and *geometry*, which correspond to object appearance and shape, respectively. The *semantic constraint* means that segmentation regions of the same part should be semantically consistent across object instances and robust to appearance and shape changes. The *geometric constraint* means a part segmentation region should be locally connected and the union of all parts should entirely cover the corresponding object. These two constraints can be seen in Fig. 1, in which the parts with the same semantic information (e.g., heads of bird) are covered by the same segmentation color (e.g., red) in different

*Corresponding author

images, and the union of all parts provides a good coverage over an object. Hung *et al.* [12] also discuss similar constraints including geometric concentration, robustness to variations, semantic consistency, and objects as union of parts, and accordingly develop an unsupervised solution for semantic part segmentation. However, for the objects-as-union-of-parts constraint, their solution lacks a principled way to enforce such a constraint and has to take a dependency on saliency map, which is not always reliable and has to be replaced with object segmentation mask in some cases¹. This dependency makes it hard to scale up with more data or to other object categories.

To relax the dependency on saliency map (or object segmentation mask) and further improve the performance of part segmentation, we develop a novel unsupervised learning approach by disentangling the appearance and shape representations of semantic parts, followed with reconstruction losses without using any additional object segmentation mask or saliency map². The idea of disentanglement and synthesis is inspired by [13], despite that it was proposed for automatic landmark point discovery whereas this work is for semantic part segmentation. By disentangling the appearance and shape representations of a target object category, we can effectively add semantic and geometric constraints on the appearance part and the shape part, respectively.

More specifically, our framework is an encoder-decoder based deep neural network. The network firstly takes a pair of images which only differ by geometric transformations, and then uses two encoders to extract the appearance and shape representations of the two images. After exchanging the appearance representations between two images, a decoder is used to reconstruct the input images to close the loop for unsupervised learning. To avoid degenerated solutions, we design a bottleneck block to squeeze and expand the appearance representation, leading to a more complete disentanglement between appearance and shape, which in return yields a better shape representation for part segmentation. On the squeezed appearance representation, we apply a classification loss to learn distinctive part features while enforcing cross-instance semantic consistency. On the disentangled shape representation, we use an improved geometric concentration loss with a special treatment for background pixels. This concentration loss, together with the reconstruction loss, essentially ensures the geometric constraint. To validate the effectiveness of the proposed

¹For example, for CUB-Birds, the open-source implementation of SCOPS uses ground truth silhouettes rather than saliency maps and crops birds w.r.t bounding boxes rather than using the original images. See <https://github.com/NVlabs/SCOPS#scops-on-caltech-ucsd-birds>

²Note that we still assume the availability of object bounding boxes, which are used in all unsupervised object landmark/part learning works to bound the learning problem to single object images.

framework, we perform comprehensive experiments on a wide variety of objects such as face, bird, and PASCAL VOC objects, and obtain superior part segmentation results compared to previous works.

Our contribution can be summarized as follows:

- To our best knowledge, this is the first attempt that can unsupervised learn semantic part segmentation without using any additional object segmentation mask (or saliency map).
- By utilizing a squeeze-and-expand bottleneck block, the proposed method can effectively disentangle the appearance and shape representations of objects and remove the dependency of using additional segmentation masks.

2. Related Work

2.1. Unsupervised Landmark Discovery

There have been many efforts to learn landmarks in a supervised manner [43, 42, 39, 37, 30, 3, 25, 35]. However, supervised landmark detection is largely limited by the availability of annotated data and mainly works for popular categories, such as face [21, 23], bird [32], *etc.* To avoid relying massive annotated data, unsupervised landmark discovery has attracted much more attention in recent years. Thewlis *et al.* [29] proposed to address landmark discovery by using the equivariance constraint and then expanding the formulation to a dense situation [28]. Jakab *et al.* [13] proposed to learn landmarks by first disentangling pose and appearance of an image and then recombining them for synthesizing a new image. We term this approach as a *disentanglement and synthesis paradigm*. This work shows that combining disentanglement with synthesis can effectively lead to the desired equivariance property, which has inspired many follow-up works, including this paper. Many subsequent works suggest improvement upon [13]. For example, Thewlis *et al.* [27] proposed to add a vector exchanging procedure for better semantic consistency across instances. Jakab *et al.* [14] then tried to learn human-interpretable landmarks from unpaired labels and Xu *et al.* [38] designed a cycle framework to learn from unpaired data. Nevertheless, such works are limited to landmark discovery, whereas our work is to obtain part segmentation from unlabeled data, which is a more challenging problem.

2.2. Part Segmentation

Our work is closely related to weakly supervised object co-segmentation. Co-segmentation aims at identifying pixels of interested foreground objects from background. There have been many related works on this subject [11, 16, 1, 26, 2, 24, 10], but they mainly deal with semantic information at the object level rather than at the part

level. Besides, most works need to take a whole dataset as input and perform co-segmentation, while we are more interested in learning a model that can predict part level semantic information, shown as a part segmentation map, for each image.

Several recent works have attempted to learn part-level information in an unsupervised manner. Hung *et al.* [12] proposed to learn part features that are semantically consistent across images and have achieved good results in their paper. However, it takes a dependency on object saliency map, which is not always reliable and has to be replaced with object segmentation mask in some cases, to distinguish a foreground object from background. In this work, we remove this dependency and make semantic part segmentation applicable to more object categories. Lorenz *et al.* [22] proposed to disentangle object shape and appearance to obtain a part modeling result, which has a similar goal to this work. However they make an elliptical assumption for each part and do not assign a part label for each pixel. Thus this approach does not work well for images whose main object occupies a large area, like human faces, as it cannot generate a complete segmentation map.

3. Methodology

As discussed in Sec. 1, we hope the learned part segmentation results could meet both the *semantic constraint* and *geometry constraint* without additional saliency map or segmentation mask. To apply the two constraints on the appearance and shape features separately, we leverage an encoder-decoder neural network with a squeeze-and-expand bottleneck block in the middle to disentangle the appearance and shape representations, which will be described in Sec. 3.1. To ensure the semantic constraint, we use a reconstruction loss and a classification loss, which will be detailed in Sec. 3.2 and Sec. 3.3 respectively. For the geometry constraint, we improve the concentration loss and describe the details in Sec. 3.4.

3.1. Overall Framework

The overall framework is an encoder-decoder based deep neural network, which includes three modules: a part segmentation encoder E_s , an appearance encoder E_a , and a decoder D for image generation. Between the two encoders and the decoder, a bottleneck block is designed to squeeze and expand the appearance representation, for a more effective disentanglement between appearance and shape. The whole framework is shown in Fig. 2.

To learn such a neural network, we assume an image collection of the interested object category is available as training data. Given an image from the training dataset, we firstly construct a pair of images which only differ by geometric transformations. That is, the transformed objects included in the two images have the same appearance but

different shape representations. As suggested in previous works [13, 22], we use the thin-plate-spline transformation (TPS) [7, 33] to obtain two transformed images $x^{(1)}, x^{(2)}$. Using $x^{(1)}$ and $x^{(2)}$ as inputs, we can extract an appearance feature map $\mathbf{A} \in \mathbb{R}^{L \times H \times W}$ and a part segmentation map $\mathbf{S} \in \mathbb{R}^{(K+1) \times H \times W}$ (corresponding to shape) by E_a and E_s for each image as

$$\mathbf{A}^{(i)} = E_a(x^{(i)}), \quad \mathbf{S}^{(i)} = E_s(x^{(i)}), \quad i = 1, 2 \quad (1)$$

where K denotes the number of object parts, L denotes the dimension of appearance features, and $H \times W$ is the resolution of two output tensors, including both the appearance feature map and the part segmentation map. The 0-th channel is used for representing the background region and the 1-st to K -th channels are used for representing K object parts. For notational convenience, we denote $K' = K + 1$. Moreover, each part segmentation map $\mathbf{S}^{(i)}, i = 1, 2$ is normalized along the channel dimension by softmax.

Appearance Transfer Bottleneck. To disentangle the appearance and shape representations, reconstruction is normally performed to provide a self-supervised loss³ after exchanging the shape (or appearance) representations between two input samples. However, directly performing reconstruction will lead to a degenerated solution. For example, the part segmentation map \mathbf{S} may contain some appearance information as well as shape information, making itself capable of reconstructing the original image, which is an undesired trivial solution. Inspired by [13] and [22], we design a simple appearance transfer bottleneck to *squeeze* and *expand* the part appearance features, aiming at generating a rendered appearance feature map that contains appearance feature from only one image and shape information from another image.

In the *squeeze* step (Eq. (2)), each channel in the part segmentation map can be regarded as an attention map, which is used to weighted average the learned appearance features over all pixels. This step reduces the dimension of the appearance feature and squeezes out the shape information of $x^{(1)}$. In the *expand* step (Eq. (3)), the squeezed part appearance feature of $x^{(1)}$ is expanded according to the part segmentation map of $x^{(2)}$, treating each response channel as an attention map to generate a rendered appearance feature map. This appearance transfer bottleneck, as well as the reconstruction, effectively helps avoid degenerated disentanglement learning.

For the sake of simplicity, we only write out the process of transferring the appearance of $x^{(1)}$ to combine with the shape of $x^{(2)}$, and it is straightforward to generalize to the inverse direction.

³In the paper, we mainly use “self-supervised” to refer to a specific loss, but use “unsupervised” to refer to the overall part segmentation problem which does not require supervised annotations.

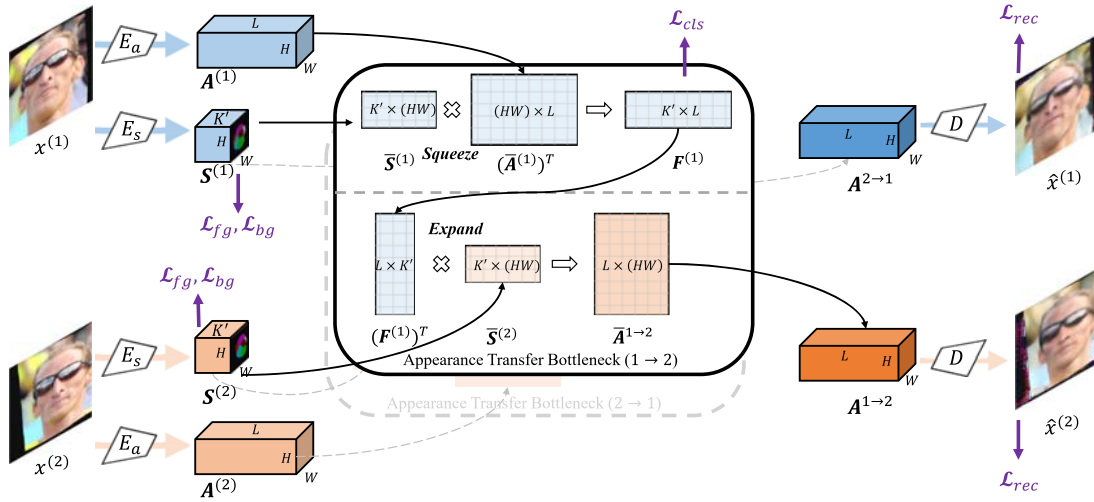


Figure 2: The overall framework of our model. The network firstly takes a pair of images which only differ by geometric transformations and then uses two encoders to extract the appearance $\mathbf{A}^{(i)}$, $i = 1, 2$ and shape representation $\mathbf{S}^{(i)}$, $i = 1, 2$ of the two images. In the middle are two bottleneck blocks, each for squeezing the appearance representation and combining it with the shape representation of another image by an expansion operation. After exchanging the appearance representations between two images, a decoder is used to reconstruct the input images to close the loop for unsupervised learning.

First, we reshape the part segmentation map $\mathbf{S}^{(1)}$ to $\bar{\mathbf{S}}^{(1)} \in \mathbb{R}^{K' \times (HW)}$ and the appearance feature map $\mathbf{A}^{(1)}$ to $\bar{\mathbf{A}}^{(1)} \in \mathbb{R}^{L \times (HW)}$ for notational convenience. Then, we squeeze the part appearance feature by using the part segmentation map to obtain $\mathbf{F}^{(1)} \in \mathbb{R}^{K' \times L}$ as follows to push out any shape information:

$$\begin{aligned} \mathbf{D} &= \text{diag}(\bar{\mathbf{S}}^{(1)} \mathbf{1}_{HW}), \\ \mathbf{F}^{(1)} &= \mathbf{D}^{-1} \bar{\mathbf{S}}^{(1)} (\bar{\mathbf{A}}^{(1)})^T, \end{aligned} \quad (2)$$

where $\mathbf{1}_{HW}$ is a vector of ones with dimension HW and \mathbf{D}^{-1} is to perform row normalization on $\bar{\mathbf{S}}^{(1)}$.

For the next step, we expand the squeezed part appearance feature $\mathbf{F}^{(1)}$ of $x^{(1)}$ according to the shape (part segmentation map) of the other image $x^{(2)}$. After reshaping $\mathbf{S}^{(2)}$ to $\bar{\mathbf{S}}^{(2)} \in \mathbb{R}^{K' \times (HW)}$, we can obtain the rendered appearance feature map $\bar{\mathbf{A}}^{1 \rightarrow 2} \in \mathbb{R}^{L \times (HW)}$ as

$$\bar{\mathbf{A}}^{1 \rightarrow 2} = (\mathbf{F}^{(1)})^T \bar{\mathbf{S}}^{(2)}. \quad (3)$$

By reshaping $\bar{\mathbf{A}}^{1 \rightarrow 2}$ back to $\mathbf{A}^{1 \rightarrow 2} \in \mathbb{R}^{L \times H \times W}$, we get the rendered appearance feature map $\mathbf{A}^{1 \rightarrow 2}$ which contains the appearance information of image $x^{(1)}$ only and the shape information of $x^{(2)}$ alone, as we expected. The inverse direction for $\mathbf{A}^{2 \rightarrow 1}$ can be derived in a similar way.

Image Generation. Using the rendered appearance feature maps $\mathbf{A}^{2 \rightarrow 1}$ and $\mathbf{A}^{1 \rightarrow 2}$ as inputs, we can obtain two reconstructed images using the decoder D , as shown in the right part of Fig. 2.

$$\hat{x}^{(1)} = D(\mathbf{A}^{2 \rightarrow 1}), \quad \hat{x}^{(2)} = D(\mathbf{A}^{1 \rightarrow 2}). \quad (4)$$

3.2. Reconstruction Loss

To ensure the *semantic constraint* being applied effectively on the disentangled appearance and shape representations, we use reconstruction as self-supervised signals. Some previous works [13, 22] have shown that the disentangling and synthesis paradigm could help learn features with the desired semantic consistency property, which is also termed as equivariance. Our reconstruction loss is defined as

$$\mathcal{L}_{rec} = \mathcal{P}(x^{(1)}, \hat{x}^{(1)}) + \mathcal{P}(x^{(2)}, \hat{x}^{(2)}), \quad (5)$$

where \mathcal{P} represents the perceptual loss [15, 4] as in previous works [13, 38].

3.3. Part Classification Loss

Similar to semantic segmentation, part segmentation can also be considered as a multi-class classification problem to assign a class label to each pixel. Although we do not have part labels for each pixel, we can perform classification on the squeezed part appearance feature $\mathbf{F} \in \mathbb{R}^{K' \times L}$. That is, the k -th part feature $\mathbf{F}_k \in \mathbb{R}^L$ should be classified correctly to the k -th class.

Such a self-supervised classification loss not only helps learn distinctive features for different parts, but also implies *semantic consistency* across object instances, which leads to the desired feature invariance property. Inspired by recent progress in face recognition feature learning [36, 19, 34, 9, 6], to further reduce the intra-class variance, we adopt the

ArcFace loss [6], which can be formulated as:

$$\mathcal{L}_{cls} = \frac{1}{K} \sum_{i=1}^K \log \frac{e^{s(\cos(\theta_{i,i+m}))}}{e^{s(\cos(\theta_{i,i+m}))} + \sum_{j=1, j \neq i}^K e^{s \cos \theta_{j,i}}}. \quad (6)$$

The background channel is excluded from the loss for its inconsistency across instances. To implement this loss function, we add a linear classification layer in the model, which has parameters $W \in \mathbb{R}^{K \times L}$ representing K feature basis of dimension L . In the loss function, $\theta_{j,i}$ is the angle between part feature $F_i \in \mathbb{R}^L$ and basis $W_j \in \mathbb{R}^L$, and s and m are two constants representing radius scaling and angular margin penalty.

Unlike previous image-level constraints, we apply self-supervised classification constraints directly on object parts. To our best knowledge, this is the first time such a constraint is used in unsupervised keypoint discovery or part segmentation. The use of this constraint effectively leads to distinctive part features while keeping cross-instance semantic consistency, which further helps improve the part segmentation performance as demonstrated in Sec. 4.2.

3.4. Foreground/Background Concentration Loss

As mentioned in Sec. 1, the learned segmentation map should meet the *geometry constraint* as well. Our method requires each pixel to be assigned with a part label. Hence the requirement of an object being entirely covered by the union of all parts is easy to implement.

For the concentration constraint which requires each object part to be locally connected, we follow the concentration loss as proposed in [12] for foreground object pixels but make a special modification for background pixels as they are scattered around an object and do not follow the concentration constraint. For foreground pixels:

$$\mathcal{L}_{fg} = \frac{1}{K} \sum_{k=1}^K \sum_{u,v} \left\| \begin{bmatrix} u \\ v \end{bmatrix} - \begin{bmatrix} c_u^k \\ c_v^k \end{bmatrix} \right\|_2^2 \cdot \mathbf{S}_{k,u,v}/z_k, \quad (7)$$

where c_u^k and c_v^k are the coordinates of the k -th part center along axis u, v , and z_k is a normalization term. c_u^k, c_v^k , and z_k can be calculated as follows:

$$c_u^k = \sum_{u,v} u \cdot \mathbf{S}_{k,u,v}/z_k, \quad c_v^k = \sum_{u,v} v \cdot \mathbf{S}_{k,u,v}/z_k, \quad (8)$$

$$z_k = \min \left\{ \sum_{u,v} \mathbf{S}_{k,u,v}, 1 \right\}. \quad (9)$$

Note that Eq. (9) is slightly different from the definition of z_k in [12], as we found it works better for avoiding degenerated solutions in our framework.

The foreground concentration loss works well for foreground pixels, but it is inappropriate for background pixels which are scattered around an object and lack semantic

consistency. [12] is not suffered from this issue because it has an extra input of saliency map to bound the region of interested object. However, as we want to remove the dependency on an extra saliency map input, we propose a background concentration loss as follows to separate out background pixels. As the channel 0 of \mathbf{S} is regarded as background, we can write the loss as:

$$\mathcal{L}_{bg} = \sum_{u,v} \|h(\langle u, v \rangle)\|^2 \cdot \mathbf{S}_{0,u,v}/z_0, \quad (10)$$

where $h(\langle u, v \rangle)$ is the distance between the pixel $\langle u, v \rangle$ and its nearest boundary which can be calculated as

$$h(\langle u, v \rangle) = \min\{u, W - u, v, H - v\}. \quad (11)$$

Compared with the foreground concentration loss as in Eq. (7), this constraint encourages background pixels to be close to image boundaries. This constraint is also based on the assumption that the target object has been cropped and roughly centered with the help of an object bounding box, as mentioned in Sec. 3.1, which is a commonly used assumption in unsupervised object part learning. Since it is hard to apply a semantic consistency constraint for diverse background regions, this geometric constraint helps the proposed algorithm focus on the main object in the center of an image.

3.5. Final Loss Function

The overall training loss is to minimize the weighted combination of all the losses in Eq. (5), (6), (7), and (10):

$$\mathcal{L}_{sum} = \lambda_{rec} \mathcal{L}_{rec} + \lambda_{cls} \mathcal{L}_{cls} + \lambda_{fg} \mathcal{L}_{fg} + \lambda_{bg} \mathcal{L}_{bg}, \quad (12)$$

where $\lambda_{rec}, \lambda_{cls}, \lambda_{fg}$ and λ_{bg} are hyper parameters for balancing the final loss.

4. Experiments

To evaluate the proposed approach, we conduct experiments on several datasets, including CelebA [21] and AFLW [23] for human faces, CUB [8] for birds, and PASCAL VOC [31] for other common object categories. As there is no ground truth for unsupervised part segmentation, we mainly follow recent works [29, 40, 12] and use landmark regression as a proxy metric to evaluate the performance of part segmentation for datasets with landmark annotations, e.g. CelebA, AFLW, and CUB. For PASCAL VOC which does not have landmark annotations, we evaluate the aggregated part segmentations with the foreground segmentation using the IOU (intersection over union) metric.

The experiments on human faces (including CelebA and AFLW) and common objects (including CUB and VOC) are detailed in Sec. 4.2 and Sec. 4.3, respectively. We also present qualitative results of appearance and shape disentanglement and transfer on DeepFashion [20] in Sec. 4.4.

4.1. Implementation Details

As mentioned in Sec 3, our framework consists of three modules: a part segmentation encoder E_p , an appearance encoder E_a , and a decoder D , all of which are implemented as deep convolutional networks using PyTorch.

We use all layers before layer4 (as denoted in the PyTorch implementation) of an unpretrained ResNet34 as the appearance encoder, except that we set the stride of layer3 to 1. For the part segmentation encoder, we replace the backbone model with ResNet18 (using all layers before layer4 and setting the stride of layer3 to 1 as well) and add an extra 1×1 convolution layer at the end of the backbone. We resize all the images to 128×128 as input resolution, and the size of the output tensor is $H \times W = 32 \times 32$. For different tasks, we will use different K . Our generator is composed of seven convolution blocks, each containing a convolution or transposed convolution layer, a batch normalization layer and a ReLU layer, and an extra 3×3 convolution layer at the end of the generator. Of these seven blocks, the second and fourth contain 4×4 transposed convolution layers for upsampling, and the rest contains 3×3 convolution layers.

We initialize all model weights with random Gaussian noise ($\sigma = 0.01$), and use Adam [17] as our optimizer, with $(\beta_1, \beta_2) = (0.9, 0.999)$, a weight decay of 5×10^{-4} , and a learning rate of 1×10^{-4} . For the ArcFace loss, we set the constants s to 20 and m to 0.5. More implementation details are available in the supplementary materials.

4.2. Evaluation on Human Faces

We first test our method on two commonly used human face dataset: CelebA [21] and AFLW [23]. CelebA is a large-scale face dataset containing more than 200K images of celebrities collected from Internet. Each image in CelebA has 40 binary annotations, a bounding box and five landmark coordinates for each face. Our setting is the same as that in [12]. We use the *wild image* for training and testing, and filter out the images with face covering less than 30% of the pixel area. Since unsupervised results cannot be measured directly, most previous works [28, 13, 22] use proxy tasks for evaluation, such as landmark regression. Following [12], we first convert part segmentations into landmarks by taking part centers as in Eq. (8) and then learn a linear regressor to map the converted landmarks to ground truth landmarks and evaluate the regression error on test data. The result can be found in Table 1.

The result shows that our method, without using any additional information, outperforms all other methods, including SCOPS which uses an additional saliency map. Although this is a proxy evaluation metric, it in some ways illustrates the good semantic consistency of our approach. To visualize the part segmentation result, we show some resulting images from CelebA in Fig. 3.

Method	K=4	K=8
ULD [29]	-	31.30
Zhang <i>et al.</i> [40]	-	40.82
IMM [13]	19.42	8.74
Lorenz <i>et al.</i> [22]	15.49	11.41
SCOPS [12](w/o saliency)	46.62	22.11
SCOPS [12](with saliency)	21.76	15.01
Ours	15.39	12.26

Table 1: Landmark regression results on wild CelebA. We report the landmark regression error in terms of mean L2 distance normalized by inter-ocular distance. Note that the first three methods are specially designed for landmarks discovery and Lorenz *et al.* is not for precise part segmentation for its elliptical shape assumption. We list them for reference.



Figure 3: Visualization and ablation studies on wild CelebA. It shows that all the \mathcal{L}_{rec} , \mathcal{L}_{fg} and \mathcal{L}_{bg} are essential for avoiding a degenerated solution and the classification loss further improves the distinctiveness of parts.

In addition to CelebA, we also test our method on another human face dataset: AFLW [23]. AFLW contains 25K faces, each of which is annotated with 21 landmarks. AFLW is a more challenging dataset because the variation across pictures is even larger. Some other works [27] provide a subset of AFLW and select five face annotations for each face. We use their *unaligned* dataset for testing. We compare our method with relevant works in Table 2. In

Method	Mean Error
IMM [13]	13.31
Lorenz <i>et al.</i> [22]	13.60
SCOPS [12](w/o saliency)	56.08
SCOPS [12](with saliency)	16.05
Ours	13.13

Table 2: Landmark regression results on unaligned AFLW. All methods are in the setting of $K = 8$.

the setting of $K = 8$, we achieve a mean error of 13.13, whereas SCOPS achieves 16.05 even with a saliency prior. The comparisons on two different datasets demonstrate the superiority of our proposed method in terms of semantic consistency of object parts.

Ablation Study. To validate the effectiveness of each loss function, we conduct ablation studies on CelebA. Table 3 shows the contribution of each loss function by excluding one loss function each time. For a more intuitive understanding, the visual results of each ablation study are shown in the Fig. 3. The results show that the combination of the reconstruction loss, the foreground loss, and the background loss can effectively help avoid a degenerated result. When any constraint is missing, the regression mean error increases dramatically. The visualization in Fig. 3 can show the impact of each loss more intuitively. The reconstruction loss \mathcal{L}_{rec} is the key to close the loop for unsupervised learning and ensure semantic consistency. Without \mathcal{L}_{rec} , the segmentation results are no better than uniform squares. Without the foreground loss \mathcal{L}_{fg} , the parts learned are scattered all over an image. If we remove the background loss \mathcal{L}_{bg} , the result leads to another trivial solution, with one part covering the whole object and others covering the background. The classification loss \mathcal{L}_{cls} helps further improve the distinctiveness of the learned part features and enforce cross-instance semantic consistency. After adding \mathcal{L}_{cls} , the mean error is further reduced from 14.38 to 12.26 as in Table 3. Without \mathcal{L}_{cls} , the part learned may cover some semantically unrelated regions. For example, the facial area contains some pixels from neck and clothes.

As we discussed in Sec. 3, the whole framework is based on an encoder-decoder structure, with the reconstruction loss to setup unsupervised learning tasks. The reconstruction loss helps ensure both the *semantic constraint* and the *geometry constraint*. The squeeze-and-expand bottleneck block makes it possible to further add the classification loss for improving the semantic consistency and add the foreground and background losses for satisfying the geometry constraint. The ablation studies demonstrate that all the four losses are essential and effective for good segmentation results.

4.3. Evaluation on Common Objects

To test the generalization ability of our method, we also evaluate it on other two datasets: CUB-200-2011 [8],

Method	K=4	K=8
w/o \mathcal{L}_{rec}	49.32	46.01
w/o \mathcal{L}_{fg}	45.77	41.42
w/o \mathcal{L}_{bg}	31.83	29.93
w/o \mathcal{L}_{cls}	17.20	14.38
Full Model	15.39	12.26

Table 3: Ablation study on CelebA. The results shows that all the four losses are essential for semantically consistent part segmentation.

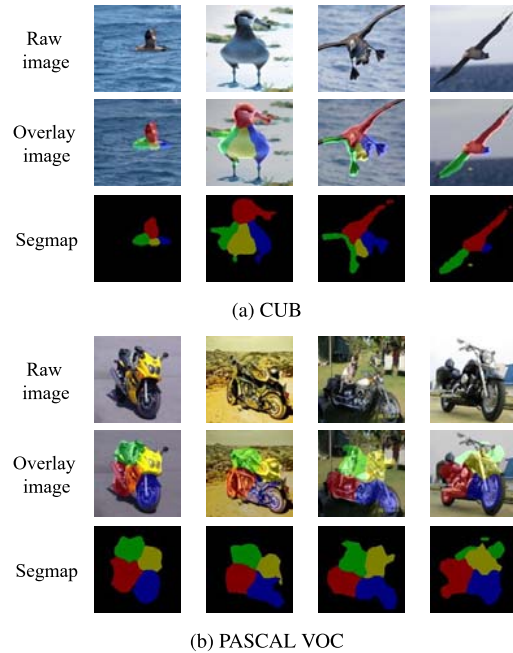


Figure 4: Visualization on CUB and PASCAL VOC. Each instance (shown in one column) has three images: the input image in the first row, the segmentation map in the last row, and the overlay of segmentation result on the input image in the second row.

termed as CUB for short, and PASCAL VOC [31], termed as VOC in the following paragraphs.

CUB [8] is a dataset with 11,788 photos of 200 bird species in total. This dataset is challenging due to the large variation in bird attitudes. In this experiment, we follow [12] and select the photos from the first 3 categories for segmentation. Similar to the previous experiments for human faces, to quantify the performance of our result, we use the expected coordinates of each part as a key point position estimation for landmark regression, where we use the annotated 15 part locations as the ground truth landmarks. We show the mean error of landmark regression in Table 4, where CUB-1, CUB-2 and CUB-3 represents the results of three categories of birds. Similarly, our results are superior to other methods, which indicates that our method keeps a good semantic consistency for objects with large variations. Visualization of CUB is shown in Fig. 4a.

Method	CUB-1	CUB-2	CUB-3
zhang <i>et al.</i> [40]	30.12	29.36	28.19
DFF [5]	22.42	21.62	21.98
SCOPS [12]	18.50	18.82	21.07
Ours	18.15	17.54	19.40

Table 4: Landmark regression results on CUB. We report the mean error of landmark regression, where the horizontal and vertical coordinates are normalized by the width and height of ground truth bounding box, respectively.

category	sheep	bus	motor	horse	cow	areo	car
DFF [5]	51.03	58.63	54.80	49.51	56.39	48.38	56.48
SCOPS [12]	56.95	73.82	58.53	55.76	60.79	69.02	65.18
Ours	64.87	78.17	59.98	53.65	65.60	52.82	70.49

Table 5: Evaluation results on PASCAL VOC. We report IOU of part-aggregated segmentation compared with the ground truth mask.

Next, we test our method on PASCAL VOC [31] for more object categories. VOC is a large-scale dataset which has been widely used for object detection and semantic segmentation tasks. It contains more than 10,000 images of 20 categories in total, and about 7,000 object instances annotated with segmentation maps. We follow [12] and select seven categories to test our method. As VOC does not have landmark annotations, we use the ground truth segmentation masks of object instances and report in Table 5 the IOU of our part-aggregated segmentation result compared with the ground truth mask. The part-aggregated segmentation is simply taken as the union of all the learned foreground parts. It is worth noting that, as SCOPS [12] explicitly uses an additional saliency map (or silhouette), this is not a fair comparison to some extent. Nevertheless, our results still outperforms previous methods on most categories, and are better than DFF [5] on all categories. Visualization of VOC is shown in Fig. 4b. Note that while our method performs well in terms of image transformation and large appearance variance, we find it still struggles with distinguishing between orientations of near-symmetric objects. Distinguishing flipped object parts remains a challenging task for self-supervised learning and deserves dedicated studies in our future work.

4.4. Disentanglement on DeepFashion

To illustrate the effect of disentanglement, which is a key component in our framework, we evaluate our method on DeepFashion [20] dataset, which contains over 800,000 diverse images. In our experiment, we select full-body images of their “in-shop clothes retrieval benchmark” for learning.

We visualize the part segmentation results and appearance transfer results in Fig. 5. Every generated image is synthesized by the appearance taking from one image in the leftmost column and the shape taking from one image in the topmost row. The visualization is mainly to show if the ap-




Figure 5: Appearance and shape disentanglement on DeepFashion. Images in the same column have the same posture, but their appearances come from the leftmost images.

pearance and shape representations are well disentangled, rather than showing high-quality synthesis images.

5. Conclusion

In this paper, we have presented an unsupervised method that can learn part segmentation without supervised annotations by disentangling appearance and shape. We considered two essential constraints: *semantic constraint* and *geometry constraint* for a good part segmentation result, and accordingly developed an encoder-decoder neural network with a squeeze-and-expand bottleneck in the middle for disentangling the appearance and shape appearances. This design effectively removes the dependency on a saliency map or object segmentation mask, yet leads to an improved result. Comprehensive experiments on a wide variety of objects such as human faces, birds, and PASCAL VOC objects verified the effectiveness of the proposed method in terms of both segmentation quality and disentanglement result.

6. Acknowledgment

This work was supported by the National Key Research and Development Program of China (No. 2017YFA0700904, No.2020AAA0104304), NSFC Projects (Nos. 61620106010, 62076147, U19A2081, U19B2034, U1811461), Beijing Academy of Artificial Intelligence (BAAI), Tsinghua-Huawei Joint Research Program, a grant from Tsinghua Institute for Guo Qiang, Tiangong Institute for Intelligent Computing, and the NVIDIA NVAAIL Program with GPU/DGX Acceleration.

References

- [1] Dhruv Batra, Adarsh Kowdle, Devi Parikh, Jiebo Luo, and Tsuhan Chen. icoseg: Interactive co-segmentation with intelligent scribble guidance. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3169–3176, 2010.
- [2] Yuning Chai, Victor Lempitsky, and Andrew Zisserman. Bicos: A bi-level co-segmentation method for image classification. In *2011 International Conference on Computer Vision*, pages 2579–2586, 2011.
- [3] Lisha Chen, Hui Su, and Qiang Ji. Deep structured prediction for facial landmark detection. In *Advances in Neural Information Processing Systems*, pages 2450–2460, 2019.
- [4] Qifeng Chen and Vladlen Koltun. Photographic image synthesis with cascaded refinement networks. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1520–1529, 2017.
- [5] Edo Collins, Radhakrishna Achanta, and Sabine Susstrunk. Deep feature factorization for concept discovery. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 336–352, 2018.
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4690–4699, 2019.
- [7] Jean Duchon. Splines minimizing rotation-invariant seminorms in sobolev spaces. In *Constructive theory of functions of several variables*, pages 85–100. Springer, 1977.
- [8] Mark Everingham, Luc Gool, Christopher K. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [9] Yandong Guo and Lei Zhang. One-shot face recognition by promoting underrepresented classes, 2018.
- [10] Junwei Han, Rong Quan, Dingwen Zhang, and Feiping Nie. Robust object co-segmentation using background prior. *IEEE Transactions on Image Processing*, 27(4):1639–1651, 2018.
- [11] Dorit S. Hochbaum and Vikas Singh. An efficient algorithm for co-segmentation. In *2009 IEEE 12th International Conference on Computer Vision*, pages 269–276, 2009.
- [12] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [13] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *NIPS 2018: The 32nd Annual Conference on Neural Information Processing Systems*, pages 4016–4027, 2018.
- [14] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8787–8797, 2020.
- [15] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision (ECCV)*, pages 694–711, 2016.
- [16] Armand Joulin, Francis Bach, and Jean Ponce. Discriminative clustering for image co-segmentation. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 1943–1950, 2010.
- [17] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [18] Xueting Li, Sifei Liu, Kihwan Kim, Shalini De Mello, Varun Jampani, Ming-Hsuan Yang, and Jan Kautz. Self-supervised single-view 3d reconstruction via semantic consistency. *arXiv preprint arXiv:2003.06473*, 2020.
- [19] Weiyang Liu, Yandong Wen, Zhiding Yu, Ming Li, Bhiksha Raj, and Le Song. Sphereface: Deep hypersphere embedding for face recognition. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6738–6746, 2017.
- [20] Ziwei Liu, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. Deepfashion: Powering robust clothes recognition and retrieval with rich annotations. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2016.
- [21] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of International Conference on Computer Vision (ICCV)*, December 2015.
- [22] Dominik Lorenz, Leonard Bereska, Timo Milbich, and Bjorn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10955–10964, 2019.
- [23] Peter M. Roth, Martin Koestinger, Paul Wohlhart and Horst Bischof. Annotated Facial Landmarks in the Wild: A Large-scale, Real-world Database for Facial Landmark Localization. In *Proc. First IEEE International Workshop on Benchmarking Facial Image Analysis Technologies*, 2011.
- [24] Fanman Meng, Hongliang Li, Guanghui Liu, and King Ngi Ngan. Object co-segmentation based on shortest path algorithm and saliency model. *IEEE Transactions on Multimedia*, 14(5):1429–1441, 2012.
- [25] Rajeev Ranjan, Vishal M. Patel, and Rama Chellappa. Hyperface: A deep multi-task learning framework for face detection, landmark localization, pose estimation, and gender recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(1):121–135, 2019.
- [26] Oana Sidi, Oliver van Kaick, Yanir Kleiman, Hao Zhang, and Daniel Cohen-Or. Unsupervised co-segmentation of a set of shapes via descriptor-space spectral clustering. In *Proceedings of the 2011 SIGGRAPH Asia Conference on*, volume 30, page 126, 2011.
- [27] J. Thewlis, S. Albanie, H. Bilen, and A. Vedaldi. Unsupervised learning of landmarks by exchanging descriptor vectors. In *International Conference on Computer Vision*.

- [28] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object frames by dense equivariant image labelling. In *NIPS'17 Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 844–855, 2017.
- [29] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 3229–3238, 2017.
- [30] Nikolai Ufer and Bjorn Ommer. Deep semantic feature matching. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5929–5938, 2017.
- [31] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-ucsd birds-200-2011 dataset. 2011.
- [32] 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.
- [33] Grace Wahba. *Spline models for observational data*. SIAM, 1990.
- [34] Hao Wang, Yitong Wang, Zheng Zhou, Xing Ji, Dihong Gong, Jingchao Zhou, Zhifeng Li, and Wei Liu. Cosface: Large margin cosine loss for deep face recognition. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5265–5274, 2018.
- [35] Rongsheng Wang, Jinzhao Tian, and Changlong Jin. Joint face detection and alignment using focal loss-based multi-task convolutional neural networks. In *Chinese Conference on Biometric Recognition*, pages 266–273, 2019.
- [36] Yandong Wen, Kaipeng Zhang, Zhifeng Li, and Yu Qiao. A discriminative feature learning approach for deep face recognition. In *European Conference on Computer Vision*, pages 499–515, 2016.
- [37] Shengtao Xiao, Jiashi Feng, Luoqi Liu, Xuecheng Nie, Wei Wang, Shuicheng Yan, and Ashraf Kassim. Recurrent 3d-2d dual learning for large-pose facial landmark detection. In *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 1642–1651, 2017.
- [38] Yinghao Xu, Ceyuan Yang, Ziwei Liu, Bo Dai, and Bolei Zhou. Unsupervised landmark learning from unpaired data. *arXiv preprint arXiv:2007.01053*, 2020.
- [39] Xiang Yu, Feng Zhou, and Manmohan Chandraker. Deep deformation network for object landmark localization, 2016.
- [40] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2694–2703, 2018.
- [41] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Facial landmark detection by deep multi-task learning. In *European Conference on Computer Vision*, pages 94–108, 2014.
- [42] Zhanpeng Zhang, Ping Luo, Chen Change Loy, and Xiaoou Tang. Learning deep representation for face alignment with auxiliary attributes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(5):918–930, 2016.
- [43] Shizhan Zhu, Cheng Li, Chen Change Loy, and Xiaoou Tang. Face alignment by coarse-to-fine shape searching. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4998–5006, 2015.