

GANSeg: Learning to Segment by Unsupervised Hierarchical Image Generation

Xingzhe He Bastian Wandt Helge Rhodin
University of British Columbia

{xingzhe, wandt, rhodin}@cs.ubc.ca

Abstract

Segmenting an image into its parts is a common pre-process for high-level vision tasks such as image editing. However, annotating masks for supervised training is expensive. Weakly-supervised and unsupervised methods exist, but they depend on the comparison of pairs of images, such as from multi-views, frames of videos, and image augmentation, which limit their applicability. To address this, we propose a GAN-based approach that generates images conditioned on latent masks, thereby alleviating full or weak annotations required by previous approaches. We show that such mask-conditioned image generation can be learned faithfully when conditioning the masks in a hierarchical manner on 2D latent points that define the position of parts explicitly. Without requiring supervision of masks or points, this strategy increases robustness of mask to viewpoint and object position changes. It also lets us generate image-mask pairs for training a segmentation network, which outperforms state-of-the-art unsupervised segmentation methods on established benchmarks. Code can be found at <https://github.com/xingzhehe/GANSeg>.

1. Introduction

This paper tackles the problem of unsupervised part segmentation. Discovering object parts in images is a fundamental problem in computer vision as parts provide an intermediate representation that is robust to object appearance and pose variation [17, 48]. Many high-level tasks benefit from part representations, such as 3D reconstruction [31, 67], pose estimation [27, 42], and image editing [14, 66]. Keypoints and part segmentation maps are among the most commonly used forms. However, their supervised training [12, 43, 56] requires pixel-level annotations for every new application domain since labels hardly generalize to other object categories and the number of parts and their granularity vary across tasks.

On the side of keypoint detection, several unsupervised detectors exist [20, 64] but segmentation methods are still in their infancy [17, 32]. Segmenting parts without pixel-

level annotation is difficult because it requires disentangling parts from other parts and the foreground from the background. Existing unsupervised¹ methods mainly follow the same strategy as applied for unsupervised keypoint detection [52]. Real images are transformed by an affine map or a thin plate spline to find those parts that are equivariant under the known deformation. For precise reconstruction they require additional information, such as saliency maps [17] or assume the objects to be consistently centered [32], which is constraining. For example, when applied to face datasets [33], the neck and shoulders are often ignored although part of almost every image.

Our goal is to improve the unsupervised part segmentation task. We propose to first train a generative adversarial network (GAN) [10] to generate images that are internally conditioned on *latent masks*. This GAN formulation alleviates the dependency on image pairs and pre-defined image transformation in existing autoencoder networks. In this way, the network learns the part distribution from the dataset instead of from pre-defined image transformations. Subsequently, we use the generator to synthesize virtually infinite mask-image pairs for training a segmentation network. Figure 1 provides an overview of our model.

The key question we address is how to design a GAN that generates images with part segmentation masks that are meaningful, i.e., group pixels into regions that typically move together and have shared appearance across images. We start from a backbone architecture that is borrowed from supervised segmentation networks [3, 43] and the GAN strategy is inspired by its recent application to unsupervised keypoint detection [14]. Our innovation is the hierarchical generation of the image via multiple abstraction levels, including the use of masks. In the first level, we use Gaussian noise to generate part appearance embeddings and a set of *2D latent points*. Unlike [14] which continues straight from points to image generation, we first group points to define the position and the scale of each part. In the second abstraction level, we use a part-relative positional encoding to

¹Most existing literature refers to *unsupervised* when training on single images without annotation, and *self-supervised* when training on auxiliary tasks using multi-views or videos. We follow this convention.

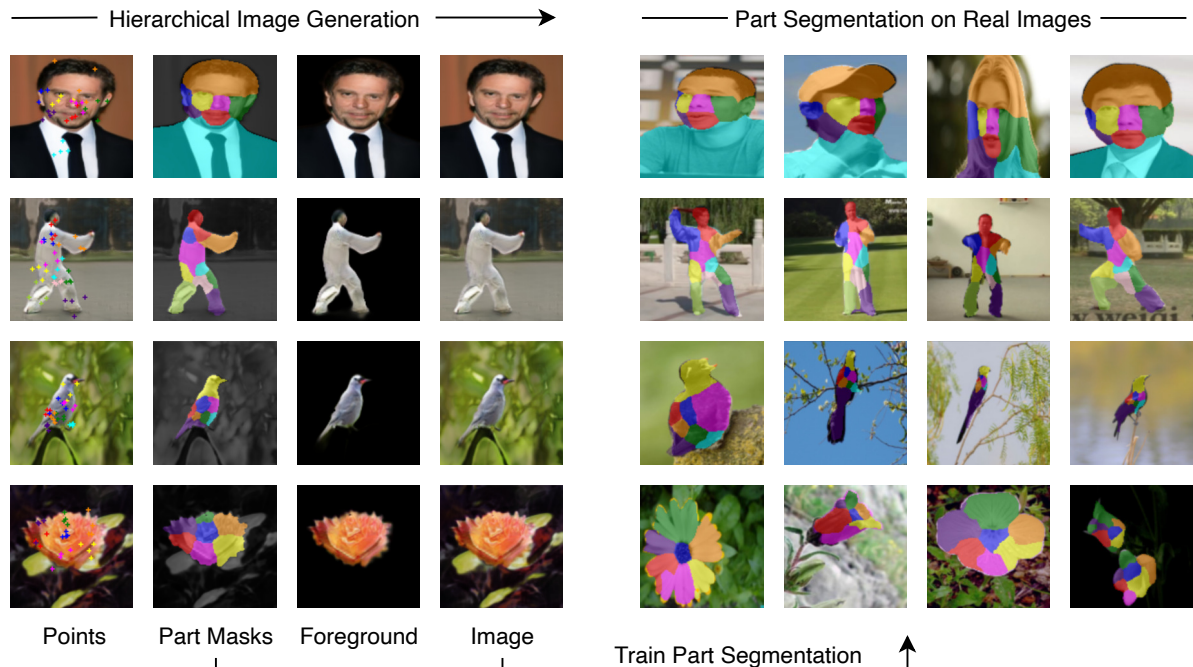


Figure 1. **GANSeg**. A segmentation network (right) is trained on mask-image pairs generated by a new hierarchical image generator (left; from points, to mask, over foreground to the image). It is unsupervised and applies to faces, persons, birds, and flowers.

generate 2D feature maps and then generate the mask with a CNN. In the third level, the foreground image is generated from a combination of feature maps with corresponding appearance embeddings. Independently, a background image is generated with a randomized position to disentangle foreground and background. Finally the foreground is blended with the background. The generated masks are used here a second time to define the blend weight.

The key to the success of our GAN framework are several design choices that preserve translational equivariance of parts [22], which are not applicable to the traditional auto-encoder approaches, as explained in Appendix E. As a result, by not knowing the absolute location, the convolutional network is forced to condition the image purely on the spatial extent of the masks; moving the part mask will move the image part. This is a crucial inductive bias in our unsupervised learning. Our contributions are threefold:

1. An unsupervised GAN approach to generate mask-image pairs for training part segmentation;
2. A novel hierarchical image generator, which encourages the segmentation of meaningful parts;
3. Alleviating prior assumptions on saliency maps and object position.

Ethics - Risks. GANs can be abused for creating deep fakes. However, our method does not work towards editing

nor improving image quality but scene understanding. Our final output is a detector, which can not be abused to generate new images but unwanted surveillance applications are a risk. **Benefits.** Since our method is entirely unsupervised, it could be applied on objects, animals, or situations that have not yet been labeled.

2. Related Work

Unsupervised Landmark Detection methods discover keypoints in the images without any supervision signals. Most existing works discover keypoints by comparing pairs of images of the same object category. The common idea is that the keypoints change as the image changes. The change can be inferred from videos [7, 21, 28, 30, 36, 47] and multi-views [40, 41, 50] of the same object category. While videos and multi-views naturally contain pairs of images, those relying on unsupervised learning from image collections require pairs created by pre-defined random transformations [20, 34, 52, 64] and learned transformation [57, 61] that are tuned for the dataset. Their underlying idea is similar. Keypoints must follow the transformation that is applied to the original images—equivariance. Recently, [14] introduced an alternative. They use a GAN to generate images along with corresponding latent keypoints and use them to train a detector. By contrast to sparse keypoints, our goal is to generate masks at the pixel level, which we attain by introducing a hierarchical generator.

Unsupervised Foreground Segmentation aims to segment foreground objects from the background in an unsupervised manner. Gupta et al. [11] learn a patch level mask of foreground to benefit representation learning, which, however, is only coarse. Singh et al. [49] use a multi-stage GAN to disentangle the foreground shape and texture but they only focus on editing. Bielski et al. [2] propose to reposition the generated foreground to disentangle it from the background. Katircioglu et al, [24, 25] detect that region as foreground that cannot be inpainted from the surrounding, a strategy previously used on optical flow [63], which however requires a similar background in all training examples or an optical flow estimator. Chen et al. [4] achieved unsupervised foreground segmentation by resampling the foreground appearance to disentangle foreground and background. Voynov et al. [54] and Yang et al. [62] introduce comparison-based segmentation methods using pre-trained GANs, achieving better results than generation-based methods. As pointed out by [2, 4], such unsupervised methods can easily get into trivial solutions, where the background includes the whole foreground. To counteract, we introduce two losses that overcome trivial solutions. Besides, we found that preserving translation equivariance in the network architecture can naturally mitigate trivial solutions. This is another important reason why we chose to use a GAN, as we will explain in Section 3 and supplementary. In comparison to the above methods, this yields comparable results even if they are specialized for foreground separation while we provide more fine-grained part segmentation.

Unsupervised Part Segmentation aims at pixel-level masks for multiple parts of an object, including separating foreground from background without mask annotation. Collins et al. [6] use matrix factorization to find similar parts in images, but it requires test images at training time, which makes it computational prohibitive. Hung et al. [17] draws lessons from unsupervised keypoint detection and extends them to predict the part segmentation masks of an object using various loss functions that preserve the geometry and semantic consistency of the masks. However, it needs off-the-shelf saliency maps or ground truth background masks. Liu et al. [32] alleviate the need for the background mask, but they use a central prior to constrain the object mask to the center of the image, which can be a constraining bias. For example, in portrait images, the hair is often not masked. Temporal information can be used [9, 48] to achieve better segmentation results. In comparison to all of these approaches, our model uses less information (single images without video or saliency map) yet outperform these on half of the most established metrics and datasets, as we evaluate in our experiments, Section 4. The recently proposed DatasetGAN [65] also utilize the a GAN to segment parts, but they still need some ground truth masks as parts are not explicitly disentangled within the network.

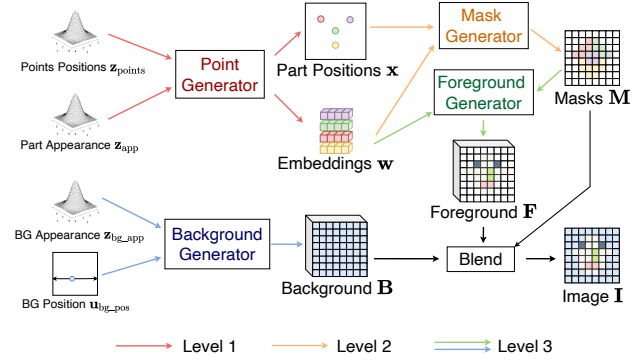


Figure 2. **Hierarchical generation overview.** **Level 1 (red):** The Point Generator converts two Gaussian noise vectors to part position and appearance embedding. **Level 2 (orange):** The Mask Generator turns part positions and embeddings into masks defining the part support. **Level 3 (green and blue):** The Foreground Generator uses the quantities from the previous level to generate a foreground image that is finally blended with the independently generated background.

3. Method

We train a Generative Adversarial Network (GAN) [10] to generate points, part masks, foreground, background, and the image in a hierarchy. Figure 2 gives an overview of our method. In a second stage, we generate mask-image pairs to train the Deeplab V3 [3] segmentation network, thereby enabling unsupervised part segmentation. Our core network architecture design principle is to build a hierarchy that preserves the translation equivariance of its part representations at each of the following three stages.

3.1. Level 1: Point Generation and Part Scale

In the first level, we utilize independent noise vectors to generate the locations and appearances of K parts. We found training to be most stable by first predicting $n_{\text{per}} \times K$ points separated into K groups with n_{per} points. The part location and scale are computed from the mean and standard deviation of the corresponding n_{per} points, which regularizes training. Figure 3 gives an overview of the underlying **Point Generator** module. It takes two noise vectors as input $\mathbf{z}_{\text{point}}, \mathbf{z}_{\text{app}} \sim \mathcal{N}(\mathbf{0}^{D_{\text{noise}}}, \mathbf{I}^{D_{\text{noise}} \times D_{\text{noise}}})$, where D_{noise} is the noise dimension. We use a 3-layer multi-layer perceptron (MLP) to map $\mathbf{z}_{\text{point}}$ to $n_{\text{per}} \times K$ points $\{\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_{\text{per}}}\}_{k=1}^K$. Then we calculate the part locations $\{\mathbf{x}_1, \dots, \mathbf{x}_K\}$ and part scales $\{\sigma_1, \dots, \sigma_K\}$,

$$\mathbf{x}_k = \frac{1}{n_{\text{per}}} \sum_{i=1}^{n_{\text{per}}} \mathbf{x}_k^i, \quad \sigma_k = \frac{\sqrt{\sum_{i=1}^{n_{\text{per}}} \|\mathbf{x}_k^i - \mathbf{x}_k\|^2}}{n_{\text{per}} - 1},$$

with $\{\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_{\text{per}}}\}_{k=1}^K = \text{MLP}_{\text{point}}(\mathbf{z}_{\text{point}})$, (1)

where $k = 1, \dots, K$.

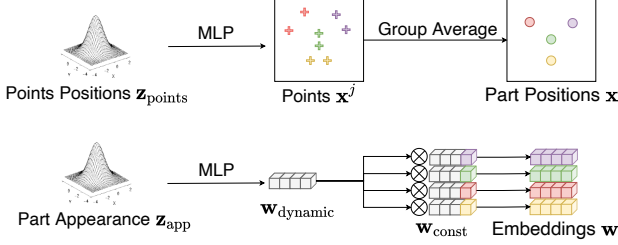


Figure 3. **Level 1.** The Point Generator uses two Gaussian noise vectors to generate part location, and part appearance embeddings, respectively.

We use another 3-layer MLP to map \mathbf{z}_{app} to a part appearance vector $\mathbf{w}^{\text{dynamic}} \in \mathbb{R}^{D_{\text{emb}}}$. Following [14], we define a constant embedding vector $\mathbf{w}_k^{\text{const}} \in \mathbb{R}^{D_{\text{emb}}}$ for each part. We then perform an elementwise multiplication between $\mathbf{w}^{\text{dynamic}}$ and $\mathbf{w}_k^{\text{const}}$ to obtain the final part embedding $\mathbf{w}_k \in \mathbb{R}^{D_{\text{emb}}}$. That is,

$$\mathbf{w}^{\text{dynamic}} = \text{MLP}_{\text{app}}(\mathbf{z}_{\text{app}}) \quad (2)$$

$$\mathbf{w}_k = \mathbf{w}^{\text{dynamic}} \otimes \mathbf{w}_k^{\text{const}} \quad (3)$$

where \otimes is the elementwise product. It is important that the noise source for appearance and position is independent to prevent the appearance from interfering with the location information.

3.2. Level 2: From Points to Masks

In the second level, the mask generation, we use Gaussian heatmaps to model the local independence and positional encoding [39] to generate masks relative to the predicted part location. We encode relative instead of absolute position between the points and the image pixels to keep long-distance relations and to prevent leaking the absolute coordinate information, which would violate the translation equivariance. To further preserve the translation equivariance, we initialize the positional encoding in a larger grid than the real image range and crop to a fixed margin size after each 2x upsampling [22], which prevents convolutional layers from passing on boundary information (see Section 3.3 for additional details).

These operations are implemented with the **Mask Generator** illustrated in Figure 4. It takes the $n_{\text{per}} \times K$ points $\{\mathbf{x}_k^1, \dots, \mathbf{x}_k^{n_{\text{per}}}\}_{k=1}^K$, part locations \mathbf{x}_k and part scales σ_k , and part embeddings \mathbf{w}_k as input. We generate a Gaussian heatmap for each part using the mean and standard deviation of each part defined in Equation 1. The embedding \mathbf{w}_k is then multiplied with every pixel of the corresponding heatmap, generating a spatially localized embedding map. We assume the additivity of feature maps (see supplementary for more details). All K part-specific embeddings are summed to form a single feature map $\mathbf{W}_{\text{mask}} \in$

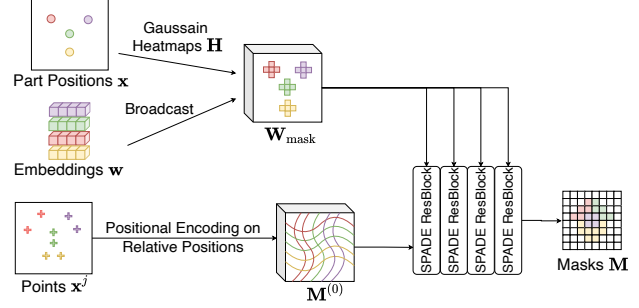


Figure 4. **Level 2.** The Mask Generator uses points, part locations, part scales, and part appearance embeddings to generate masks.

$\mathbb{R}^{D_{\text{emb}} \times H \times W}$. Formally we write,

$$\mathbf{H}_k(\mathbf{p}) = \exp(-\|\mathbf{p} - \mathbf{x}_k\|_2^2 / \sigma_k^2), \quad (4)$$

$$\mathbf{W}_{\text{mask}}(\mathbf{p}) = \sum_{k=1}^K \mathbf{H}_k(\mathbf{p}) \mathbf{w}_k.$$

Note that we use σ_k^2 instead of $2\sigma_k^2$ to obtain sharper heatmaps, making it easier for the network to generate sharp masks.

The generated embedding map \mathbf{W}_{mask} will subsequently be used to generate masks, together with the mask starting tensor $\mathbf{M}^{(0)} \in \mathbb{R}^{D_{\text{emb}} \times H \times W}$. To avoid leaking absolute position information, we do not use a constant tensor [23] or a linearly mapped noise [39]. Instead, we use low frequency positional encoding [51] of the difference between the pixel position and the $n_{\text{per}} \times K$ points. That is,

$$\mathbf{M}^{(0)}(\mathbf{p}) = [\sin(\pi \text{FC}([\mathbf{p} - \mathbf{x}_1^1, \dots, \mathbf{p} - \mathbf{x}_K^{n_{\text{per}}}])), \cos(\pi \text{FC}([\mathbf{p} - \mathbf{x}_1^1, \dots, \mathbf{p} - \mathbf{x}_K^{n_{\text{per}}}]))] \quad (5)$$

where FC stands for a fully connected layer without activation function followed (a linear projection).

With the mask starting tensor $\mathbf{M}^{(0)}$ and mask embedding map \mathbf{W}_{mask} defined, we generate masks $\mathbf{M} = [\mathbf{M}_{\text{bg}}, \mathbf{M}_1, \dots, \mathbf{M}_K] \in \mathbb{R}^{(K+1) \times H \times W}$ with SPADE ResBlocks [39],

$$\mathbf{M}^{(i)} = \text{SPADE ResBlock}(\mathbf{M}^{(i-1)}, \mathbf{W}_{\text{mask}}) \quad (6)$$

$$\mathbf{M} = \text{softmax}(\mathbf{M}^{(T_{\text{mask}})})$$

where $i = 1, \dots, T_{\text{mask}}$, T_{mask} is the number of blocks, and an additional channel is reserved for the background. For more details of SPADE ResBlock, we refer the reader to our supplementary document and the original paper [39]. In theory, the Batch Normalization [18] may leak absolute position information and break the translation equivariance. However, in practice, experiments have shown SPADE has strong local disentanglement [14, 39, 66].

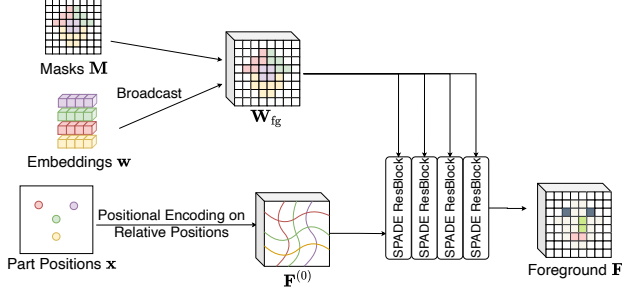


Figure 5. **Level 3 - Part I.** The foreground Generator uses part locations, part appearance embeddings, and masks.

3.3. Level 3: Mask-conditioned Image Generation

In the third level, we generate the foreground and the background separately and blend them linearly by reusing the masks from the previous level. The **Foreground Generator** is illustrated in Figure 5. It takes the $K + 1$ masks \mathbf{M} , K part locations \mathbf{x}_k , and K part appearance embedding \mathbf{w}_k as input. Similar to the procedure that generates masks, we first broadcast the embedding with the corresponding mask to generate the foreground embedding map $\mathbf{W}_{\text{fg}} \in \mathbb{R}^{D_{\text{emb}} \times H \times W}$,

$$\mathbf{W}_{\text{fg}}(\mathbf{p}) = \sum_{k=1}^K \mathbf{M}_k(\mathbf{p}) \mathbf{w}_k. \quad (7)$$

We then use K part locations to generate the foreground starting tensor $\mathbf{F}^{(0)}$ with low frequency positional encoding similar to Equation 5. Finally we use SPADE ResBlocks to generate the Foreground feature map $\mathbf{F} \in \mathbb{R}^{D_{\text{emb}} \times H \times W}$,

$$\begin{aligned} \mathbf{F}^{(i)} &= \text{SPADE ResBlock}(\mathbf{F}^{(i-1)}, \mathbf{W}_{\text{fg}}) \\ \mathbf{F} &= \mathbf{F}^{(T_{\text{fg}})}, \end{aligned} \quad (8)$$

where $i = 1, \dots, T_{\text{fg}}$ and T_{fg} is the number of SPADE ResBlocks. Independent of this, the **Background Generator** takes two noise vectors as input $\mathbf{z}_{\text{bg_app}} \sim \mathcal{N}(\mathbf{0}^{D_{\text{noise}}}, \mathbf{I}^{D_{\text{noise}} \times D_{\text{noise}}})$, $\mathbf{u}_{\text{bg_pos}} \sim \mathcal{U}([-1, 1]^2)$. We first use a 3-layer MLP to map $\mathbf{z}_{\text{bg_app}}$ to a background appearance vector $\mathbf{w}_{\text{bg}} \in \mathbb{R}^{D_{\text{emb}}}$,

$$\mathbf{w}_{\text{bg}} = \text{MLP}_{\text{bg_app}}(\mathbf{z}_{\text{bg_app}}). \quad (9)$$

Positional encoding is used on the difference between the background center $\mathbf{u}_{\text{bg_pos}}$ and the pixel position, to generate the background starting tensor $\mathbf{B}^{(0)}$, similar to Equation 5.

Finally we use AdaIN ConvBlocks [15] to generate the background feature map $\mathbf{B} \in \mathbb{R}^{D_{\text{emb}} \times H \times W}$,

$$\begin{aligned} \mathbf{B}^{(i)} &= \text{AdaIN ConvBlock}(\mathbf{B}^{(i-1)}, \mathbf{W}_{\text{bg}}) \\ \mathbf{B} &= \mathbf{B}^{(T_{\text{bg}})} \end{aligned} \quad (10)$$

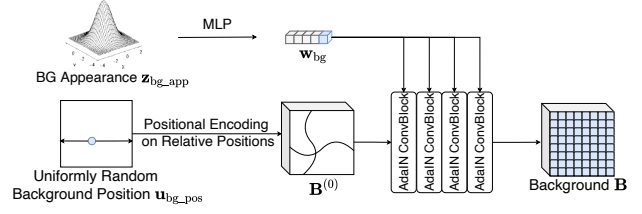


Figure 6. **Level 4 - Part II.** The Background Generator uses a Gaussian noise vector, and a random position to generate translation-invariant background.

where $i = 1, \dots, T_{\text{bg}}$ and T_{bg} is the number of AdaIN ConvBlocks. For more detail on AdaIN ConvBlocks, we refer readers to our supplementary, or the original paper [15].

Combining Foreground and Background. Recall that we generate the background mask $\mathbf{M}_{\text{bg}} \in \mathbb{R}^{H \times W}$ along with the part masks. This is used to combine our Foreground and Background together. The final image is generated by feeding the feature map into a two-layer CNN. That is,

$$\mathbf{I} = \text{Conv}((1 - \mathbf{M}_{\text{bg}}) \otimes \mathbf{F} + \mathbf{M}_{\text{bg}} \otimes \mathbf{B}) \quad (11)$$

where \otimes is the pixel-wise product.

Upsampling and Cropping. For simplicity, we used size $H \times W$ for all above feature maps. For mask, foreground and background generation, the starting tensor has a 10-pixel-wide margin at each border, same as [22], so that the central feature map will not be interfering with the boundary. For example, instead of generating a $H_0 \times W_0$ grid in range $[-1, 1]^2$, we generate a $(H_0 + 20) \times (W_0 + 20)$ grid in range

$$[-1 - 20/H_0, 1 + 20/H_0] \times [-1 - 20/W_0, 1 + 20/W_0]. \quad (12)$$

We use this grid to calculate the starting tensors for masks, foreground, and background. After each SPADE ResBlock and each AdaIN ConvBlock, we use 2x upsampling on the feature maps. The margin becomes 20-pixel wide. To bound the otherwise increasing boundary width, we subsequently crop the feature map to keep the 10-pixel margin. The Gaussian heatmaps are calculated on a grid with a 10-pixel-wide-margin, separately for each resolution.

3.4. Training Objective

Our hierarchical GAN is trained end-to-end on image collections using the following loss functions.

Adversarial Loss. We denote \mathcal{G} as the generator and \mathcal{D} as the discriminator. We use the non-saturating loss [10],

$$\mathcal{L}_{\text{GAN}}(\mathcal{G}) = \mathbb{E}_{\mathbf{z} \sim \mathcal{N}} \log(\exp(-\mathcal{D}(\mathcal{G}(\mathbf{z}))) + 1) \quad (13)$$

for the generator, and logistic loss,

$$\begin{aligned} \mathcal{L}_{\text{GAN}}(\mathcal{D}) &= \mathbb{E}_{\mathbf{z} \sim \mathcal{N}} \log(\exp(\mathcal{D}(\mathcal{G}(\mathbf{z}))) + 1) + \\ &\quad \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \log(\exp(-\mathcal{D}(\mathbf{x})) + 1) \end{aligned} \quad (14)$$

for the discriminator, with gradient penalty [35] applied only on real data,

$$\mathcal{L}_{\text{gp}}(\mathcal{D}) = \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \nabla \mathcal{D}(\mathbf{x}). \quad (15)$$

Geometric Concentration Loss Pixels from the same segment are usually connected and concentrated around its center as assumed by [17]. We enforce the mask to be in an area around its center with the geometric concentration loss

$$\mathcal{L}_{\text{con}}(\mathcal{G}) = \sum_{k=1}^K \sum_{\mathbf{p}} \frac{\mathbf{M}_k(\mathbf{p})}{\sum_{\mathbf{p}'} \mathbf{M}_k(\mathbf{p}')} \|\mathbf{p} - \mathbf{x}_k\|_2^2. \quad (16)$$

Note that the background mask is not constrained by $\mathcal{L}(\mathcal{G})_{\text{con}}$. We found that this loss alone can cause parts to collapse since it encourages a small mask area. To mitigate this problem, we introduce the Area Loss below.

Area Loss. We force the mask area to be larger than the area of the Gaussian heatmap, which is not predefined but predicted from the generated points in level 1. If the part is visible, this loss encourages the visibility of the mask. Otherwise, the area loss encourages a smaller (close to a zero area) Gaussian heatmap.

$$\mathcal{L}_{\text{area}}(\mathcal{G}) = \sum_{k=1}^K \max \left(0, \sum_{\mathbf{p}} \mathbf{H}_k(\mathbf{p}) - \sum_{\mathbf{p}} \mathbf{M}_k(\mathbf{p}) \right). \quad (17)$$

We will empirically show that this loss makes the masks more consistent.

The final loss for the discriminator is

$$\mathcal{L}(\mathcal{D}) = \mathcal{L}_{\text{GAN}}(\mathcal{D}) + \lambda_{\text{gp}} \mathcal{L}_{\text{gp}}(\mathcal{D}), \quad (18)$$

and the final loss for the generator is

$$\mathcal{L}(\mathcal{G}) = \mathcal{L}_{\text{GAN}}(\mathcal{G}) + \lambda_{\text{con}} \mathcal{L}_{\text{con}}(\mathcal{G}) + \lambda_{\text{area}} \mathcal{L}_{\text{area}}(\mathcal{G}). \quad (19)$$

4. Experiments

Following related work [17, 32, 48], we analyze the improvement of our method in terms of the positioning of parts and the mask coverage on the established benchmarks. We provide results for a wide variety of images containing faces, animals, flowers, and humans. The supplemental provides additional examples.

4.1. Datasets and Metrics

Our evaluation follows the dataset-specific protocols and metrics as established in prior work:

CelebA-in-the-wild [33] shows celebrity faces in unconstrained conditions and is used to estimate the consistency of part positions and the mask center of mass. We follow [17] removing images with a face-covering fewer than 30% of the pixel area from the MAFL train and test sets. As

in [17, 32], we train a linear regression model without bias from the part centers to the ground truth keypoints. The error metric is the landmark regression error in terms of mean L_2 distance normalized by inter-ocular distance. The splits are 45609 images for GAN training, 5379 with keypoint labels for regression, and 283 for testing.

CelebA-aligned [33] contains 200k faces, each centered such that eyes align. Following [52], we use three subsets: CelebA training set without MAFL (160k images), MAFL training set (19k images), MAFL test set (1k images). The error metric is the same as on CelebA-in-the-wild.

CUB-2011 [55] consists of 11,788 images of birds. We follow [4] to use 10,000 images for training, 1,000 for testing, and the remaining 778 for validation. We use this dataset to analyze segmentation coverage accuracy. We aggregate part segments to form the foreground mask and calculate the Intersection Over Union (IoU) between the predicted foreground masks and the ground truth foreground masks [17, 32]. To this end, we calculate the foreground mask as the sum of our part masks.

Flowers [38] consists of 8,189 images of flowers. The ground truth masks are obtained by an automated method built specifically for flowers [38]. We follow [4] to use 6,149 images for training, 1,020 for testing, and 1020 for validation. The metric is IoU of the foreground mask.

Taichi [33] contains 3049 training videos and 285 test videos of people performing Tai-Chi. We train our GAN on the training images (*not* videos) using 5000 images for fitting the regression model and 300 other images for testing. For a fair comparison, we use the same 5300 images as [48]. The metric, mean average error (MAE), is calculated as the sum of the L_2 distances between 18 regressed keypoints and their ground truth. We also calculate IoU based on the provided foreground masks.

4.2. Baselines

We compare to the following unsupervised methods, most requiring stronger assumptions: **DEF** [6] uses test images at training time, which has a high computational cost. **SCOPS** [17] uses saliency maps. **Liu et al.** [32] requires a strong prior on the mask center being close to the image center. **Siarohin et al.** [48] trained on videos exploiting temporal information.

4.3. Qualitative analysis

We show detection examples in Figure 7, and qualitatively compare our predicted masks with the baselines in Figure 8. Our mask has better coverage with the foreground object, less fragmented parts, and better consistency, with the same part masks mapping to the same body parts across different images. For example, the shoulder exists in almost all images in wild CelebA [33], but SCOPS [17] cannot consistently discover this apparent part, likely because

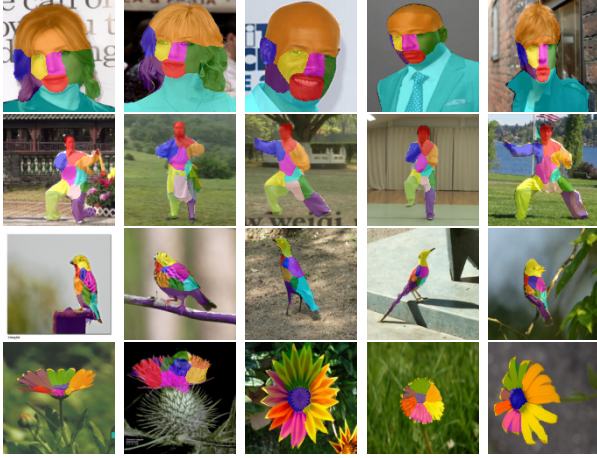


Figure 7. **Detection Examples** showing faithful part segmentation on the tested datasets, with varying object size and complexity.

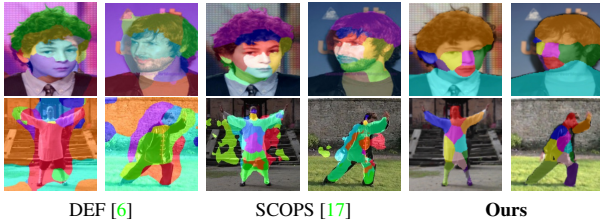


Figure 8. **Mask quality comparison.** We qualitatively compare the mask on wild CelebA (first row) and Taichi (second row). Our masks have better quality than other methods due to our hierarchical generator.

Method	Type	Aligned (K=10)	Wild (K=4)	Wild (K=8)
Thewlis et al. [52]	Landmark	7.95%	-	31.30% *
Zhang et al. [64]	Landmark	3.46%	-	40.82% *
LatentKeypointGAN [14]	Landmark	5.85%	25.81%	21.90%
Lorenz et al. [34]	Landmark	3.24%	15.49% †	11.41% †
IMM [20]	Landmark	3.19%	19.42% †	8.74% †
DEF [6] by [17]	Part	-	-	31.30% *
SCOPS [17] (w/o saliency)	Part	-	46.62%	22.11%
SCOPS [17] (w/ saliency)	Part	-	21.76%	15.01%
Liu et al. [32]	Part	-	15.39%	12.26%
Huang et al. [16] (w/ detailed label)	Part	-	-	8.40%
Ours	Part	3.98%	12.26%	6.18%

Table 1. **Landmark detection on CelebA.** The metric is the landmark regression (without bias) error in terms of mean L_2 distance normalized by inter-ocular distance (lower is better). Although the pair-based methods work better on aligned CelebA, they do not generalize well. Our generation-based methods are more robust on Wild CelebA. The sign * means reported by [17] and † by [32].

their employed saliency map focuses on the face. For more qualitative comparisons, refer to the supplementary.

4.4. Part Center Consistency

Table 1 shows the results for keypoint detection. Our method outperforms other part segmentation methods on the part consistency metric on the challenging CelebA-in-the-wild. The keypoint-based methods that perform better

Method	Type	CUB	Flowers
GrabCut [44]	Foreground Segmentation	0.360	0.692
PerturbGAN [2]	Foreground Segmentation	0.380	-
ReDO [4]	Foreground Segmentation	0.426	0.764
IEM + SegNet [45]	Foreground Segmentation	0.551	0.789
SCOPS [17] w/ GT BG by us	Part Segmentation	0.329 †	0.544 †
Ours	Part Segmentation	0.629	0.739

Table 2. **Foreground-Background Segmentation.** The metric on CUB and Flowers is the IoU of foreground (higher is better). We use $K = 8$ for both ours and SCOPS [17]. The sign † means being trained by us using their official implementation.

Method	Type	MAE ↓	IoU ↑
DEF [6]	no train/test split	494.48 †	-
SCOPS [17]	images with saliency maps	411.38 †	0.5485 †
Strohwin et al. [48]	trained on videos	389.78	0.7686
Ours	trained on single images	417.17	0.8538

Table 3. **Part Segmentation on Taichi.** The sign † means being reported by [48]. All the results are shown in the case of $K = 10$.

on aligned face images do not generalize to real-world examples. The results on Taichi in Table 3 confirm that our model is comparable in consistency (MAE metric) even though others are trained with additional information. In addition, we evaluate our model on CUB using the protocol of [34]. Our estimation error 3.23% (L_2 error normalized by edge length of the image) is lower (better) than Zhang et al. [64] (5.36%) and Lorenz et al. [34] (3.91%).

4.5. Mask Coverage

Since the datasets used in [17, 32] are too small for training a GAN (see limitations) we only use the sufficiently large CUB-2011 [33] and Flowers [38] dataset. To nevertheless compare with SCOPS [17] we retrain their model. We use the ground truth masks and the required saliency masks on our train/val/test split with their default parameters. Furthermore, we compare our results with the unsupervised foreground-background segmentation methods that already report on these datasets. Note that these methods focus on foreground-background segmentation instead of part segmentation. Their approaches are tailored for this task while it is rather a byproduct for us.

As shown in Table 2, our model outperforms the unsupervised part segmentation methods with better mask coverage, and is comparable with the dedicated foreground segmentation methods. On the most challenging Taichi dataset, we outperform all methods in Table 3 in terms of mask coverage (foreground IOU) even though they use additional information, such as saliency maps and videos.

4.6. Ablation Tests

Figure 9 shows a qualitative comparison when removing different parts from our full model. Instead of showing the segmentation on real images, we show the generated masks



Figure 9. **Ablation Tests.** From top to bottom: models disabling 1) points, 2) separated background, 3) points and separated background, 4) geometric concentration loss 5) area loss, and 6) nothing (full model). For visualization purposes we show the background in grey-scale. All contributions improve across datasets.

from our GAN. These masks explain our model more directly. For ablation tests on the number of parts, we refer readers to the supplementary document.

Disabling Points. We remove the point generation from the network and in turn also the Gaussian heatmaps and SPADE blocks. Instead, we use AdaIN [15], which is a spatially invariant version of SPADE. We use the mean of the part embeddings as the embedding for the feature map of AdaIN and replace the starting tensor with a learned constant tensor as in StyleGAN [23].

Disabling Separated Background Generation. We integrate the background generator into the foreground generator and treat the background as one of the parts.

Disabling Losses. We remove losses and show examples from each dataset in Figure 9. Without the Geometric Concentration Loss, masks become fragmented. Without the Area Loss some parts vanish. Surprisingly, even without points and separate background generation, the model can

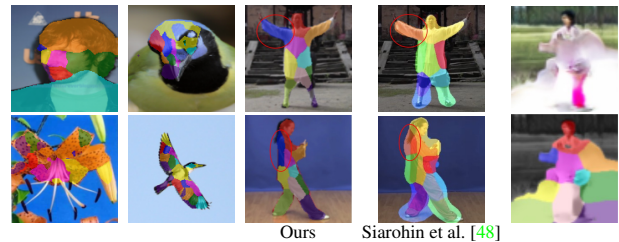


Figure 10. **Limitations.** The method fails for poses underrepresented in the training set (1st and 2nd col) and under occlusions (left eye of the face is marked as blue, which is usually on the ear), which we have in common with other methods such as [48] (4th col). In some rare cases the generator fails to generate good body shape on Taichi (5th col).

discover some commonly shared parts in the images except for human bodies, as shown in the third row in Figure 9. We hypothesize it is due to our spatially variant network design with translation equivariance encouraged at all levels. Overall, it is evident that our hierarchical architecture is crucial for mask generation.

5. Limitations and Future Work

Our GAN training requires a larger dataset (>5000 images), which does not apply to some existing benchmarks but is acceptable for domains with large unlabelled image collections. Figure 10 shows representative failure cases that are typical also for related work. In some rare cases and only for the Taichi dataset, the generator fails to generate good shape. Nevertheless, the detection remains accurate as there is no such unusual shape in the real data. When a part is occluded in the image the associated mask will still cover a region nearby, such as the left arm moving to the back instead of being occluded. Sometimes, front/back for bird and left/right for humans are flipped. The comparison to [48] shows that this is common for 2D methods and demands for extensions to 3D and occlusion handling.

6. Conclusion

We presented a GAN-based approach for learning part segmentation from unlabelled image collections. Crucial is our hierarchical generator design that synthesizes images in a coarse-to-fine manner, with independence and invariance built in. It forms a viable alternative to existing autoencoder techniques and opens up a path for learning part-based 3D models from 2D images in the future.

Acknowledgement

This work was supported by the UBC Advanced Research Computing (ARC) GPU cluster, the Compute Canada GPU servers, and a Huawei-UBC Joint Lab project.

References

- [1] Bilal Alsallakh, Narine Kokhlikyan, Vivek Miglani, Jun Yuan, and Orion Reblitz-Richardson. Mind the pad – {cnn}s can develop blind spots. In *International Conference on Learning Representations*, 2021. 14
- [2] Adam Bielski and Paolo Favaro. Emergence of object segmentation in perturbed generative models. *Advances in Neural Information Processing Systems*, 32:7256–7266, 2019. 3, 7
- [3] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 1, 3
- [4] Mickaël Chen, Thierry Artières, and Ludovic Denoyer. Unsupervised object segmentation by redrawing. In *Advances in Neural Information Processing Systems 32 (NIPS 2019)*, pages 12705–12716, 2019. 3, 6, 7
- [5] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. 14
- [6] 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. 3, 6, 7
- [7] Xuanyi Dong, Shou-I Yu, Xinshuo Weng, Shih-En Wei, Yi Yang, and Yaser Sheikh. Supervision-by-registration: An unsupervised approach to improve the precision of facial landmark detectors. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 360–368, 2018. 2
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 14
- [9] Qingzhe Gao, Bin Wang, Libin Liu, and Baoquan Chen. Unsupervised co-part segmentation through assembly. *arXiv preprint arXiv:2106.05897*, 2021. 3
- [10] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1, 3, 5
- [11] Kamal Gupta, Saurabh Singh, and Abhinav Shrivastava. Patchvae: Learning local latent codes for recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4746–4755, 2020. 3
- [12] 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. 1
- [13] 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. 14
- [14] Xingzhe He, Bastian Wandt, and Helge Rhodin. Latentkeypointgan: Controlling gans via latent keypoints. *arXiv preprint arXiv:2103.15812*, 2021. 1, 2, 4, 7
- [15] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1501–1510, 2017. 5, 8, 12
- [16] Zixuan Huang and Yin Li. Interpretable and accurate fine-grained recognition via region grouping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8662–8672, 2020. 7
- [17] Wei-Chih Hung, Varun Jampani, Sifei Liu, Pavlo Molchanov, Ming-Hsuan Yang, and Jan Kautz. Scops: Self-supervised co-part segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 869–878, 2019. 1, 3, 6, 7
- [18] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015. 4, 12
- [19] Md Amirul Islam, Sen Jia, and Neil DB Bruce. How much position information do convolutional neural networks encode? *arXiv preprint arXiv:2001.08248*, 2020. 14
- [20] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks through conditional image generation. In *Advances in neural information processing systems*, pages 4016–4027, 2018. 1, 2, 7
- [21] Tomas Jakab, Ankush Gupta, Hakan Bilen, and Andrea Vedaldi. Self-supervised learning of interpretable keypoints from unlabelled videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8787–8797, 2020. 2
- [22] Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *arXiv preprint arXiv:2106.12423*, 2021. 2, 4, 5, 14
- [23] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4401–4410, 2019. 4, 8, 12
- [24] Isinsu Katircioglu, Helge Rhodin, Victor Constantin, Jorg Sporri, Mathieu Salzmann, and Pascal Fua. Self-supervised human detection and segmentation via background inpainting. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 01:1–1, 2021. 3
- [25] Isinsu Katircioglu, Helge Rhodin, Jorg Sporri, Mathieu Salzmann, and Pascal Fua. Human detection and segmentation via multi-view consensus. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2855–2864, 2021. 3
- [26] Osman Semih Kayhan and Jan C van Gemert. On translation invariance in cnns: Convolutional layers can exploit absolute spatial location. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14274–14285, 2020. 14
- [27] Martin Kiefel and Peter Vincent Gehler. Human pose estimation with fields of parts. In David Fleet, Tomas Pajdla,

- Bernt Schiele, and Tinne Tuytelaars, editors, *Computer Vision – ECCV 2014*, pages 331–346, Cham, 2014. Springer International Publishing. 1
- [28] Yunji Kim, Seonghyeon Nam, In Cho, and Seon Joo Kim. Unsupervised keypoint learning for guiding class-conditional video prediction. In *Advances in Neural Information Processing Systems*, pages 3814–3824, 2019. 2
- [29] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015. 12
- [30] Tejas D Kulkarni, Ankush Gupta, Catalin Ionescu, Sebastian Borgeaud, Malcolm Reynolds, Andrew Zisserman, and Volodymyr Mnih. Unsupervised learning of object keypoints for perception and control. In *Advances in neural information processing systems*, pages 10724–10734, 2019. 2
- [31] 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. In *European Conference on Computer Vision*, pages 677–693. Springer, 2020. 1
- [32] Shilong Liu, Lei Zhang, Xiao Yang, Hang Su, and Jun Zhu. Unsupervised part segmentation through disentangling appearance and shape. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8355–8364, 2021. 1, 3, 6, 7
- [33] 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. 1, 6, 7
- [34] Dominik Lorenz, Leonard Bereska, Timo Milbich, and Bjorn Ommer. Unsupervised part-based disentangling of object shape and appearance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10955–10964, 2019. 2, 7
- [35] Lars Mescheder, Andreas Geiger, and Sebastian Nowozin. Which training methods for gans do actually converge? In *International conference on machine learning*, pages 3481–3490. PMLR, 2018. 6
- [36] Matthias Minderer, Chen Sun, Ruben Villegas, Forrester Cole, Kevin P Murphy, and Honglak Lee. Unsupervised learning of object structure and dynamics from videos. In *Advances in Neural Information Processing Systems*, pages 92–102, 2019. 2
- [37] Rito Murase, Masanori Suganuma, and Takayuki Okatani. How can cnns use image position for segmentation? *arXiv preprint arXiv:2005.03463*, 2020. 14
- [38] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008. 6, 7
- [39] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019. 4, 12
- [40] Helge Rhodin, Victor Constantin, Isinsu Katircioglu, Mathieu Salzmann, and Pascal Fua. Neural scene decomposition for multi-person motion capture. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [41] Helge Rhodin, Mathieu Salzmann, and Pascal Fua. Unsupervised geometry-aware representation for 3d human pose estimation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 750–767, 2018. 2
- [42] Iasonas Kokkinos Riza Alp Güler, Natalia Neverova. Densepose: Dense human pose estimation in the wild. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 1
- [43] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 1, 14
- [44] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ”grabcut” interactive foreground extraction using iterated graph cuts. *ACM transactions on graphics (TOG)*, 23(3):309–314, 2004. 7
- [45] Pedro Savarese, Sunnie SY Kim, Michael Maire, Greg Shakhnarovich, and David McAllester. Information-theoretic segmentation by inpainting error maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4029–4039, 2021. 7
- [46] Maximilian Seitzer. pytorch-fid: FID Score for PyTorch. <https://github.com/mseitzer/pytorch-fid>, August 2020. Version 0.1.1. 12
- [47] Aliaksandr Siarohin, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Animating arbitrary objects via deep motion transfer. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2377–2386, 2019. 2
- [48] Aliaksandr Siarohin, Subhankar Roy, Stéphane Lathuilière, Sergey Tulyakov, Elisa Ricci, and Nicu Sebe. Motion-supervised co-part segmentation. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 9650–9657. IEEE, 2021. 1, 3, 6, 7, 8
- [49] Krishna Kumar Singh, Utkarsh Ojha, and Yong Jae Lee. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6490–6499, 2019. 3
- [50] Supasorn Suwajanakorn, Noah Snavely, Jonathan J Tompson, and Mohammad Norouzi. Discovery of latent 3d keypoints via end-to-end geometric reasoning. In *Advances in neural information processing systems*, pages 2059–2070, 2018. 2
- [51] Matthew Tancik, Pratul P Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *arXiv preprint arXiv:2006.10739*, 2020. 4
- [52] James Thewlis, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of object landmarks by factorized spatial embeddings. In *Proceedings of the IEEE international conference on computer vision*, pages 5916–5925, 2017. 1, 2, 6, 7

- [53] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016. [12](#)
- [54] Andrey Voynov, Stanislav Morozov, and Artem Babenko. Object segmentation without labels with large-scale generative models. In *International Conference on Machine Learning*, pages 10596–10606. PMLR, 2021. [3](#)
- [55] 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. [6](#)
- [56] Shih-En Wei, Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Convolutional pose machines. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 4724–4732, 2016. [1](#)
- [57] Wayne Wu, Kaidi Cao, Cheng Li, Chen Qian, and Chen Change Loy. Transgaga: Geometry-aware unsupervised image-to-image translation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8012–8021, 2019. [2](#)
- [58] Shiyong Xiong and Yue Yang. The boundary-constraint method for constructing vortex-surface fields. *Journal of Computational Physics*, 339:31–45, 2017. [14](#)
- [59] Shiyong Xiong and Yue Yang. Identifying the tangle of vortex tubes in homogeneous isotropic turbulence. *Journal of Fluid Mechanics*, 874:952–978, 2019. [14](#)
- [60] Rui Xu, Xintao Wang, Kai Chen, Bolei Zhou, and Chen Change Loy. Positional encoding as spatial inductive bias in gans. *arXiv preprint arXiv:2012.05217*, 2020. [14](#)
- [61] Yinghao Xu, Ceyuan Yang, Ziwei Liu, Bo Dai, and Bolei Zhou. Unsupervised landmark learning from unpaired data. *arXiv preprint arXiv:2007.01053*, 2020. [2](#)
- [62] Yu Yang, Hakan Bilen, Qiran Zou, Wing Yin Cheung, and Xiangyang Ji. Unsupervised foreground-background segmentation with equivariant layered gans. *arXiv preprint arXiv:2104.00483*, 2021. [3](#)
- [63] Yanchao Yang, Antonio Loquercio, Davide Scaramuzza, and Stefano Soatto. Unsupervised moving object detection via contextual information separation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 879–888, 2019. [3](#)
- [64] Yuting Zhang, Yijie Guo, Yixin Jin, Yijun Luo, Zhiyuan He, and Honglak Lee. Unsupervised discovery of object landmarks as structural representations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2694–2703, 2018. [1](#), [2](#), [7](#)
- [65] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10145–10155, 2021. [3](#)
- [66] Peihao Zhu, Rameen Abdal, Yipeng Qin, and Peter Wonka. Sean: Image synthesis with semantic region-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5104–5113, 2020. [1](#), [4](#)
- [67] Silvia Zuffi and Michael J. Black. The stitched puppet: A graphical model of 3d human shape and pose. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3537–3546, 2015. [1](#)