

# Learning Foreground-Background Segmentation from Improved Layered GANs

Yu Yang<sup>1</sup>, Hakan Bilen<sup>2</sup>, Qiran Zou<sup>1</sup>, Wing Yin Cheung<sup>1</sup>, Xiangyang Ji<sup>1</sup>  
<sup>1</sup> Tsinghua University, BNRist <sup>2</sup> The University of Edinburgh

yang-yu16@mails.tsinghua.edu.cn, hbilen@ed.ac.uk,

{zouqr19, zhangyx20}@mails.tsinghua.edu.cn, xyji@tsinghua.edu.cn

## Abstract

Deep learning approaches heavily rely on high-quality human supervision which is nonetheless expensive, time-consuming, and error-prone, especially for image segmentation task. In this paper, we propose a method to automatically synthesize paired photo-realistic images and segmentation masks for the use of training a foreground-background segmentation network. In particular, we learn a generative adversarial network that decomposes an image into foreground and background layers, and avoid trivial decompositions by maximizing mutual information between generated images and latent variables. The improved layered GANs can synthesize higher quality datasets from which segmentation networks of higher performance can be learned. Moreover, the segmentation networks are employed to stabilize the training of layered GANs in return, which are further alternately trained with Layered GANs. Experiments on a variety of single-object datasets show that our method achieves competitive generation quality and segmentation performance compared to related methods.

## 1. Introduction

Deep learning approaches continue to dramatically push the state-of-the-art in most supervised computer vision tasks, but they demand tremendous number of human annotations which are expensive, time-consuming, and error-prone, especially for image segmentation task. Hence there is a growing interest in reducing the reliance on human supervision. Towards this target, we investigate learning to synthesize paired images and segmentation masks without segmentation annotation for the use of training segmentation network (Fig. 1).

In the recent years, generative adversarial networks (GANs) [18] repeatedly show its ability to not only synthesize highly photorealistic images [26, 27, 7] but also disentangle attributes [26, 27] and geometric factors [54]. To make GANs factorize image layers (*i.e.* foreground and background), a body of works [52, 47, 5, 10, 6], dubbed

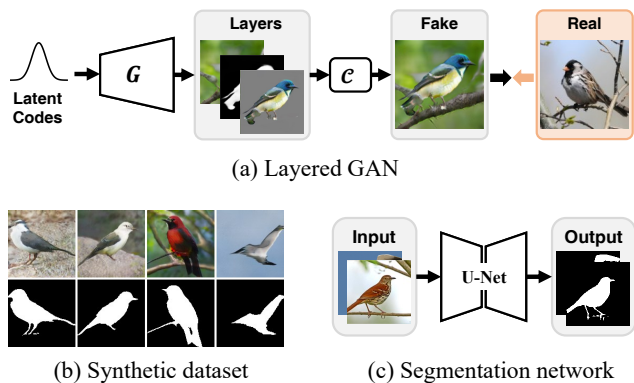


Figure 1. **Overview.** A layered GAN (a) is trained to synthesize paired image and segmentation masks (b) for training a segmentation network, *e.g.* U-Net (c).  $C$  denotes the composition module.

*layered* GANs in this paper, employ a compositional generation process which first generates background and foreground objects and then composes them to obtain the synthetic image. This compositional generation process, however, does not guarantee an expected factorization and is prone to trivial compositions such as “all as foreground” and “all as background” (Fig. 2(a) and Fig. 4) without any supervision or any regularization.

PerturbGAN [6] addresses this issue by perturbing the foreground position during composition, as such perturbation leads to inferior generated images for all-as-foreground decomposition and maintains decent generated images for expected decomposition (Fig. 2(b)). Despite so, perturbation is unable to solve the all-as-background issue, or empty foreground issue, where it can not affect the fidelity of generated image. To alleviate this problem, PerturbGAN [6] penalizes the cases where the number of foreground pixels are smaller than a threshold. However, this method is not always effective and the additionally introduced hyperparameter is sensitive to object scale, object category and dataset.

We propose a GAN that synthesizes background and foreground from a pair of latent codes. In more details, one of the latent codes, dubbed public code, is shared across

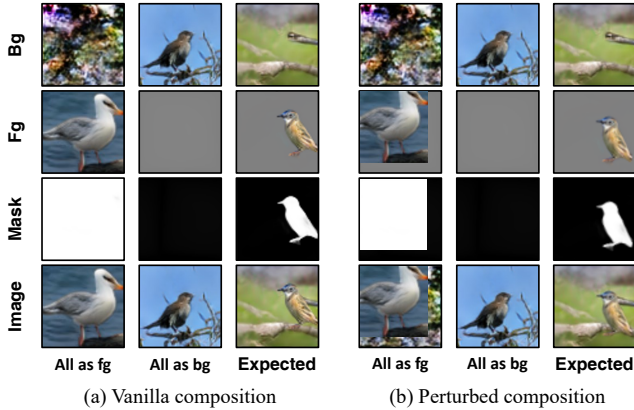


Figure 2. **Illustrations** of the learned composition and the effect of perturbation. Vanilla composition permits “all as foreground” ( $\pi = 1$ ), “all as background” ( $\pi = 0$ ), and expected composition. Perturbed composition penalizes “all as foreground” by causing unnatural border, maintains the fidelity for expected composition, but is unable to counter “all as background”.

background and foreground generation, while the other, dubbed private code, is only reserved for foreground generation. In contrast to PerturbGAN [6], we not only employ perturbation and require our generators to generate photorealistic images but also diversify the generated foreground by maximizing the mutual information between the composed image and the private code as well as the mutual information between the foreground mask and the private code. To this end, the all-as-background issue can be effectively mitigated and the layered GANs decompose foreground and background in a superior way.

Our method is reminiscent of a rich body of work [16, 8, 19, 15, 34, 14, 3] on object-centric scene generation, which is mostly based on variational autoencoders (VAEs). These models are able to decompose and generate multi-instance complex scenes. Nonetheless, these methods are mostly demonstrated on synthetic scene but seldom applied to real-world images [14]. In contrast, our method has an advantage in dealing with realistic images but lack the ability to model multi-instance generation.

Our method presents a solution to unsupervised foreground-background segmentation (Fig. 1) where layered GANs are trained to generate synthetic training data from which segmentation networks are trained. In return, the segmentation networks are used to further regularize the generation of layered GANs. These two steps are alternated which finally leads to a high-performance segmentation network. Our method learns segmentation with the analysis-by-synthesis principle, which is different from earlier work [44, 21, 25, 2, 9] that learns to discriminate foreground and background pixels using shallow classifiers on hand-crafted features and the recent trend to simultaneously learn per-pixel representations and clustering [24, 40, 29]

with deep neural networks.

We evaluate our method on a variety of single-object datasets including Caltech-UCSD Birds 200-2011 (CUB), Stanford Cars, Stanford Dogs, and Amazon Picking Challenge (APC), where it outperforms the baseline methods and achieves competitive performance compared to related methods.

Our contributions are summarized as follows:

1) We propose a mutual information based learning objective and show that it achieves significantly better performance in terms of disentangling foreground and background.

2) We show that the synthesized image and segmentation mask pairs can be successfully used to train segmentation networks which achieve competitive performance compared to related methods.

## 2. Related Work

**Layered GANs** Generative adversarial networks (GANs) [18] have been proposed and improved towards highly photorealistic image synthesis [26, 7, 27] and explored for disentangling factors of variation towards controllable generation [17, 46, 26, 41, 23, 49, 48, 22], efficient representation learning [12], and saving of human annotations [55, 54]. To gain controllable generation at object level, layered GANs [52, 47, 6, 5] leverage the compositional nature of images and synthesize images by composing image layers, *i.e.* background and foreground objects. This compositional generation process is further developed to support 3D-aware scene generation [32, 37, 39, 38] and video generation [13]. Unlike most work that is interested in generation quality, our work focus on improving and employing 2D layered GANs towards learning segmentation with minimal human supervision.

In layered GANs, it is observed that compositional generation process is prone to degeneration to single layer generation [52, 10, 6]. This issue is alleviated by composing layers at feature level [13, 39], using weak supervision (*e.g.* bounding box annotation or background images) [32, 47, 5], or specific regularization (*e.g.* perturbation strategy as in PerturbGAN [6]). We propose a method based on mutual information maximization to mitigate this issue which can be successfully employed in layered GANs under the challenging unsupervised learning setting.

**Object-Centric Scene Generation** A rich body of work [16, 8, 19, 15, 34, 14, 3], grounded on variational autoencoders (VAEs) [30], simultaneously learn scene decomposition, object-centric representations, inter-object relationships, and multi-instance scene generation. Although these works also employ a compositional generation process and interested in the automatic segmentation, they are

mostly demonstrated on synthetic scenes but rarely evaluated on real-world datasets. In contrast, our work focuses on compositional generation and segmentation of real-world datasets.

**Unsupervised Segmentation** Earlier works [44, 21, 25, 2, 9] cast unsupervised object segmentation as co-segmentation and utilise handcrafted features and shallow classifiers. With the advent of deep learning, per-pixel representation and clustering are simultaneously learned with deep neural networks [24, 40, 29] to achieve unsupervised segmentation. Apart from these works, other works learn segmentation from generative models. In particular, segmentation hint is extracted via cut and paste [42, 4, 6, 1], manipulating generation process [50, 36], erasing and re-drawing [10], or maximizing inpainting error [45]. Our methods are based on generative models and learn segmentation from compositional generation process.

### 3. Methods

#### 3.1. Revisiting Layered GANs

Let  $G : \mathcal{Z} \rightarrow \mathcal{X}$  denote a generative model that maps a latent variable  $\mathbf{z} \in \mathcal{Z}$  to an image  $\mathbf{x} \in \mathcal{X}$ . In GANs, there is also a discriminator  $D : \mathcal{X} \rightarrow \mathbb{R}$  tasked to classify real and fake images.  $G$  and  $D$  are trained by playing the following adversarial game

$$\min_G \max_D \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}(\mathbf{x})} f(D(\mathbf{x})) + \mathbb{E}_{\mathbf{z}} f(-D(G(\mathbf{z}))), \quad (1)$$

where  $f : \mathbb{R} \rightarrow \mathbb{R}$  is a concave function which differs for different kinds of GANs. In the end of adversarial learning, the images generated from  $G$  are high-fidelity and is hardly distinguished by  $D$ .

To explicitly disentangle foreground objects from background, layered GANs decompose  $G$  into multiple generators responsible for generating layers which are further composed to synthesize full images. Formally, regarding two-layer generation,  $G$  is decomposed into background generator  $G_b$  and foreground generator  $G_f$  as  $G = (G_b, G_f)$ .  $G_b$  generates background image  $\mathbf{x}_b$  and  $G_f$  generates paired foreground image  $\mathbf{x}_f$  and mask  $\pi$ . The synthetic image  $\mathbf{x}$  is completed by blending  $\mathbf{x}_f$  and  $\mathbf{x}_b$  with  $\pi$  as alpha map,

$$\mathbf{x} = \mathcal{C}(\mathbf{x}_b, \mathbf{x}_f, \pi) = (1 - \pi) \odot \mathbf{x}_b + \pi \odot \mathbf{x}_f, \quad (2)$$

where  $\mathcal{C}$  denote the composition step. However, this naive layered GAN is prone to trivial decomposition such as ‘‘all as foreground’’ and ‘‘all as background’’ (Fig. 2 and Fig. 4), which requires further regularization as follows.

**Layer perturbation** Bielski & Favaro [6] propose to perturb the position of foreground object in composition step

during training. Formally, this perturbed composition can be formulated as

$$\mathcal{C}(\mathbf{x}_b, \mathbf{x}_f, \pi; \mathcal{T}) = (\mathbf{1} - \mathcal{T}(\pi)) \odot \mathbf{x}_b + \mathcal{T}(\pi) \odot \mathcal{T}(\mathbf{x}_f), \quad (3)$$

where  $\mathcal{T}$  denotes a perturbation operator which PerturbGAN [6] instantiates as pixel-unit translation. In this paper, we employ restricted affine transformation. This method can effectively alleviate the all-as-foreground issue, because artefacts like sharp and unnatural border (Fig. 2(b)) emerges in the images under this circumstance and such inferior generation is penalized by discriminator. It should be noted that this perturbed composition does not hurt the fidelity of composed images in expected decomposition as long as the perturbation is imposed within a small range, because the relative position of foreground with respect to background can be varied up to a certain degree for most images. However, the perturbed composition is insufficient to address the all-as-background issue, which brings to our mutual information maximization as in the next section.

#### 3.2. Improving Layered GANs with Mutual Information Maximization

Fig. 3 presents the framework of our model. Generally, we improve layered GANs with in terms of disentangling foreground object and background from a special design of generative model and mutual information maximization, which is described as follows.

Our generative model takes as input two kinds of latent code, denoted as  $\mathbf{z}$  and  $\mathbf{c}$ , respectively. In particular,  $\mathbf{z}$  is shared across  $G_b$  and  $G_f$ , whereas  $\mathbf{c}$  is private to  $G_f$ . We therefore refer to  $\mathbf{z}$  and  $\mathbf{c}$  as public code and private code, respectively (Fig. 3). Based on this design, we introduce an additional learning objective aiming at maximizing mutual information between the composed image  $\mathbf{x}$  and the private code  $\mathbf{c}$  and mutual information between the generated mask  $\pi$  and the private code  $\mathbf{c}$ ,

$$\max I(\mathbf{x}, \mathbf{c}) + I(\pi, \mathbf{c}). \quad (4)$$

This learning objective can penalize  $G_f$  from always generating zero mask, which is shown as following.

For the first term in Equ. (4), it is the composed image  $\mathbf{x}$  that is measured the mutual information with  $\mathbf{c}$ . If  $G_f$  produces all-zero mask, *i.e.*  $\pi = \mathbf{0}$ , the composed image would only contain background, *i.e.*  $\mathbf{x} = \mathbf{x}_b$ , and  $I(\mathbf{x}, \mathbf{c}) = I(\mathbf{x}_b, \mathbf{c}) = 0$  since  $\mathbf{c}$  is private to foreground generation process. Therefore, non-zero mask is essential for maximizing  $I(\mathbf{x}, \mathbf{c})$ . For the second term, the mutual information can be rewritten as difference between two entropies,  $I(\pi, \mathbf{c}) = H(\pi) - H(\pi|\mathbf{c})$ . Note that our generative model is deterministic, meaning there is no stochastic in  $\pi$  given  $\mathbf{c}$ , *i.e.*  $H(\pi|\mathbf{c}) = 0$ . Hence, maximizing  $I(\pi, \mathbf{c})$  is equivalent to maximizing  $H(\pi)$ . If  $G_f$  constantly generates all-zero masks,  $H(\pi)$ , which depicts the diversity of

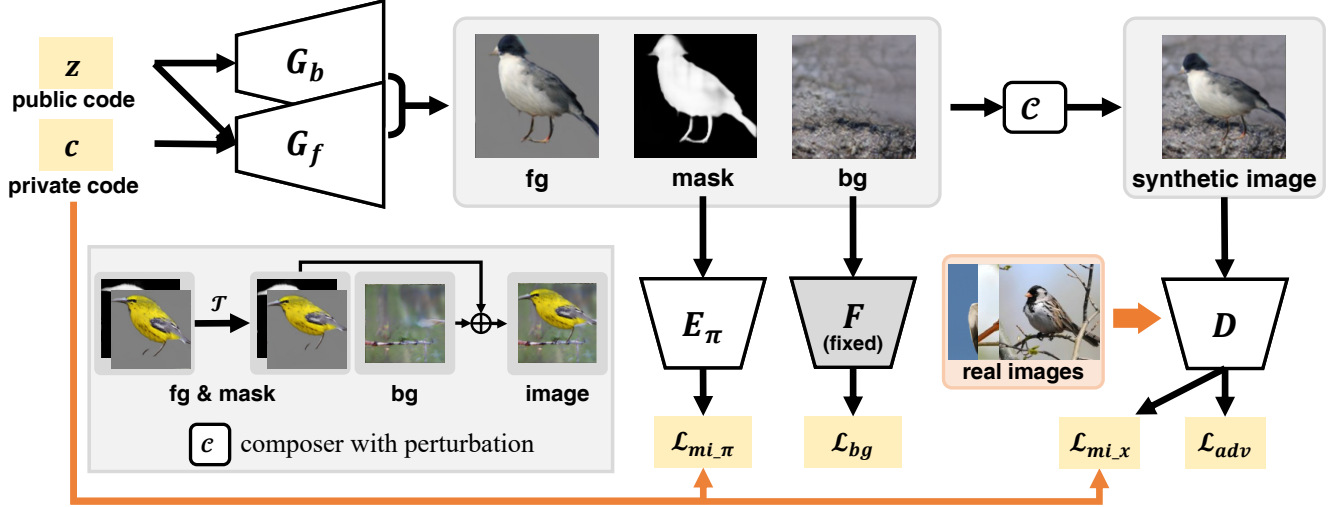


Figure 3. **Framework.** Our generative model comprises background generator  $G_b$ , foreground generator  $G_f$ , and composer with perturbation  $C$ . The public code  $\mathbf{z}$  is shared across foreground and background generation, while the private code  $\mathbf{c}$  is only used for foreground generation. Our generative model is mainly trained via adversarial learning against a discriminator  $D$ , denoted as  $\mathcal{L}_{adv}$  and regularized by mutual information loss, denoted as  $\mathcal{L}_{mi,\pi}$  and  $\mathcal{L}_{mi,x}$ , whose estimation is assisted by  $E_\pi$  and  $E_x$ .  $E_x$ , which is omitted in the figure, shares backbone parameters with  $D$  and bifurcates at the top layers. In the alternate training, background generation is regularized by segmentation network  $F$ , denoted as  $\mathcal{L}_{bg}$ .

masks, would be very low. Therefore, maximizing  $I(\pi, \mathbf{c})$  also advocates non-zero masks.

In practice, computing the mutual information,  $I(\mathbf{x}, \mathbf{c})$  and  $I(\pi, \mathbf{c})$ , is intractable. Following InfoGAN [11], we instead maximize the variational lower bound of the mutual information, which is realized as following. First, neural networks  $E_\pi$  and  $E_x$  are introduced to approximate the posterior distribution  $p(\mathbf{c}|\pi)$  and  $p(\mathbf{c}|\mathbf{x})$ , respectively. Second,  $G$  and  $E_\pi$  and  $E_x$  are learned by minimizing the following loss

$$\mathcal{L}_{mi} = \mathbb{E}_{\mathbf{z}, \mathbf{c}} \log E_x(\mathbf{c}|\mathbf{x}) + \log E_\pi(\mathbf{c}|\pi) \quad (5)$$

where  $E_x(\mathbf{c}|\mathbf{x})$  and  $E_\pi(\mathbf{c}|\pi)$  denote the likelihood of a sample  $\mathbf{c}$  under the approximate posterior distribution of  $p(\mathbf{c}|\pi)$  and  $p(\mathbf{c}|\mathbf{x})$ , respectively. At the implementation level,  $E_x$  shares backbone network parameters with discriminator  $D$  and bifurcates at the last two layers, as commonly done in [11, 33, 47, 5]. More implementation details can be found in the appendices.

In addition to penalizing trivial masks, introducing private code and maximizing  $I(\mathbf{x}, \mathbf{c})$  is also favorable for *foreground-oriented clustering*. In particular,  $\mathbf{c}$  can be designed as discrete variable and its prior distribution can be a uniform categorical distribution. In this way,  $\mathbf{c}$  is interpreted as clustering category and  $E_x$  is regarded as a clustering network trained from synthetic data. Note that the unique design of our generative model, where  $\mathbf{c}$  is only used for foreground generation, makes the clustering based only on the variety of foreground objects and not interfered by backgrounds. This mechanism conforms to the logics of category annotation for major datasets, where it is foreground

objects that are categorized rather than background. Experimental results in Section 4.3 (Table 4) justifies the merit of our foreground-oriented clustering.

### 3.3. Alternate Training of Layered GAN and Segmentation Network

**Training a layered GAN** The improved Layered GAN is optimized by updating generator and discriminator alternately,

$$\begin{aligned} \min_G \mathcal{L}_{adv}^{(G)} + \gamma_{mi} \mathcal{L}_{mi}^{(G)} + \gamma_b \mathcal{L}_b^{(G)}, \\ \min_{D, E_x, E_\pi} \mathcal{L}_{adv}^{(D)} + \gamma_{mi} \mathcal{L}_{mi}^{(E_x, E_\pi)}, \end{aligned} \quad (6)$$

where  $\mathcal{L}_{mi}$  denotes the mutual information loss as defined in Equ. (5),  $\mathcal{L}_{adv}$  denotes the adversarial loss for which hinge loss is employed, *i.e.*  $\mathcal{L}_{adv}^{(G)} = -\mathbb{E}_{\mathbf{z}, \mathbf{c}} D(G(\mathbf{z}, \mathbf{c}))$ ,  $\mathcal{L}_{adv}^{(D)} = \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} \max(0, 1 - D(\mathbf{x})) + \mathbb{E}_{\mathbf{z}, \mathbf{c}} \max(0, 1 + D(G(\mathbf{z}, \mathbf{c})))$ ,  $\mathcal{L}_b^{(G)}$  denotes the binarization loss  $\mathcal{L}_b^{(G)} = \mathbb{E}_{\mathbf{z}, \mathbf{c}} \min(\pi, \mathbf{1} - \pi)$  to encourage the binarization of masks as in [6], and  $\gamma_{mi}$ , and  $\gamma_b$  denote the corresponding loss weights.

**Training a Segmentation Network** Improved layered GAN can synthesize paired images and segmentation masks by sampling latent variables from prior distribution and forwarding them to the generator. This procedure is repeated for  $N$  times to generate a synthetic dataset  $\mathcal{D}_{syn} = \{(x_i, \pi_i)\}_{i=1}^N$  which is further used for training a foreground-background segmentation network  $F : \mathcal{X} \rightarrow$

**Algorithm 1: Alternate training**


---

```

1 for  $k = 1, \dots, K$  do
2   Set  $\hat{\gamma}_{bg} = 0$  if  $k = 1$  else  $\hat{\gamma}_{bg} = \gamma_{bg}$ 
3   for  $n = 1, \dots, N_{gan}$  do
4     Update  $G_b$  and  $G_f$  with loss
4      $\mathcal{L}^{(G)} = \mathcal{L}_{adv}^{(G)} + \gamma_{mi} \mathcal{L}_{mi}^{(G)} + \gamma_b \mathcal{L}_b^{(G)} + \hat{\gamma}_{bg} \mathcal{L}_{bg}^{(G)}$ 
5     Update  $D, E_x$  and  $E_\pi$  with loss
5      $\mathcal{L}^{(D)} = \mathcal{L}_{adv}^{(D)} + \gamma_{mi} \mathcal{L}_{mi}^{(E_x, E_\pi)}$ 
6   for  $n = 1, \dots, N_{seg}$  do
7     Generate a minibatch of synthetic dataset  $\mathcal{B}$ .
8     Update  $F$  on  $\mathcal{B}$  with loss
8      $\mathcal{L}^{(F)} = \mathbb{E}_{(x, \pi) \sim \mathcal{B}} [(1 - \pi) \log(1 - F(x)) + \pi \log(F(x))]$ 

```

---

$[0, 1]^{H \times W}$  as follows,

$$\min_F \mathbb{E}_{(x, \pi) \sim \mathcal{D}_{syn}} [(1 - \pi) \log(1 - F(x)) + \pi \log(F(x))]. \quad (7)$$

Since the layered GAN approximates the real data distribution and produces reasonable masks, the optimized  $F$  could fairly generalize to real images.

**Alternate Training** As presented in Algorithm 1, an alternate training schedule is employed to exploit the synergy between layered GAN and segmentation network. In particular, the training is conducted for  $K$  rounds. In the initial round, a layered GAN is trained without the help of segmentation network. The initially learned layered GAN is sufficient to generate synthetic dataset for training a segmentation network. In the rest rounds, the segmentation network is used to regularize the background generation with an additional loss term

$$\mathcal{L}_{bg}^{(G)} = \mathbb{E}_{\mathbf{z}, \mathbf{c}} \log(1 - F(\mathbf{x}_b)), \quad (8)$$

where the parameters of segmentation network are fixed. This loss encourages the generated background pixels,  $\mathbf{x}_b$ , to be classified as background and helps to stabilize the learned disentanglement of foreground and background. The training of layered GAN is then followed by another time of updating segmentation network. In this way, layered GAN and segmentation network are alternately trained. More training details are available in the appendices.

## 4. Experiments

### 4.1. Settings

**Datasets** Our method is evaluated on a variety of datasets including *Caltech-UCSD Birds 200-2011* (CUB) [51], *Stanford Cars* [31], *Stanford Dogs* [28], and *Amazon Picking*

	$\mathcal{L}_{size}$	$\gamma_{mi}$				
		0.1	0.5	1.0	5.0	10.0
FID ↓	19.4	19.1	17.6	<b>11.6</b>	15.5	17.8
bg-FID ↑	20.9	18.6	17.4	99.5	107.1	<b>120.2</b>
IoU ↑	6.3	0.	0.1	<b>68.6</b>	66.5	66.3

Table 1. Ablation study with respect to mutual information maximization on CUB.  $\gamma_{mi}$  denotes the loss weight of mutual information maximization.  $\mathcal{L}_{size}$  denotes the mask size loss.

*Challenge* (APC) [53]. CUB, Stanford Cars and Stanford Dogs are fine-grained category datasets for bird, car, and dog, respectively. Most images in these three datasets only contain single object. APC is a image dataset collected from robot context which contains images of single object either on a shelf or in a tray. Unlike the above four datasets, objects in APC belong to a variety of categories such as ball, scissors, book, *etc.* Ground truth segmentation masks of CUB and APC are available from their public release, whereas pre-trained Mask R-CNN are employed to approximate the ground truth segmentation of Stanford Cars, and Stanford Dogs, as in previous work [47, 5]. More details about the datasets are available in the appendices.

**Evaluation metrics** Our methods are quantitatively evaluated with respect to generation quality and segmentation performance. To measure the quality of generation, Fréchet Inception Distance (FID) [20] between 10k synthetic images and all the images are computed. The segmentation performance is evaluated with per-pixel accuracy (ACC), foreground intersection over union (IoU) between predicted masks and ground truth ones, and mean IoU (mIoU) across foreground and background. Moreover, adjusted random index (ARI) [19] and mean segmentation covering (MSC) [15] are computed to compare the segmentation performance level of our method to that of VAE-based methods. In addition to above metrics, normalized mutual information (NMI) is computed to measure the category clustering performance of our method which is compared to other layered GANs.

### 4.2. Ablation Studies

We conduct multiple ablation studies to investigate (i) the effect of mutual information maximization, (ii) the influence of different choices for private code  $\mathbf{c}$  distribution, and (iii) the effect of alternate training by controlled experiments which are mainly conducted in CUB dataset at  $64 \times 64$  resolution.

**The effect of mutual information maximization** Table 1 and Fig. 4 presents the results of ablation study with respect to mutual information maximization, where  $\gamma_{mi}$  is varied or mask size loss  $\mathcal{L}_{size}$  is used to replace mutual information

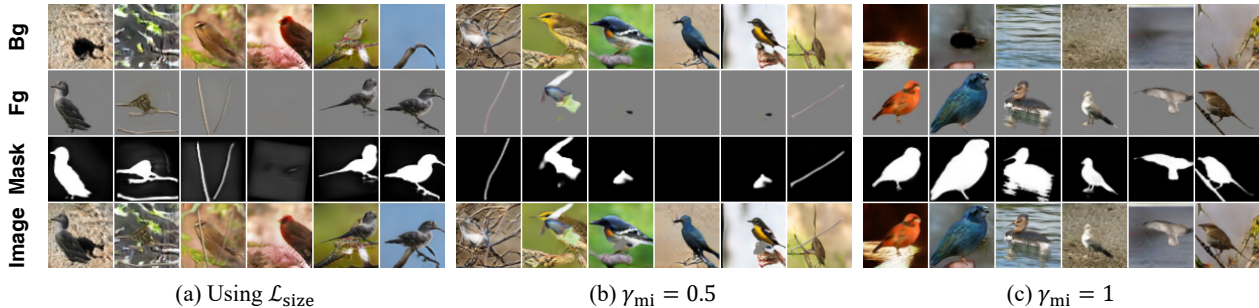


Figure 4. **Ablation study on mutual information maximization** on CUB.  $\gamma_{mi}$  denotes the loss weight of mutual information maximization.  $\mathcal{L}_{size}$  denotes that mutual information maximization is replaced with with mask size loss.

	Normal				Categorical				
	$d_c$	4	6	8	10	50	100	200	400
$H(c)$	5.6	8.5	11.3	14.1	3.9	4.6	5.2	5.9	
FID	12.3	12.1	11.8	11.1	11.3	11.6	11.6	11.9	
IoU	72.3	72.5	66.7	65.4	71.5	69.1	68.6	68.6	

Table 2. **Comparison of different prior distribution for  $c$**  in CUB.  $d_c$  denotes the dimension of  $c$  for normal distribution and the number of categories for categorical distribution.  $H(c)$  denotes the entropy of  $c$  drawn from the prior distribution.

maximization. In addition to FID and IoU, we also compute bg-FID which is the FID between synthetic background image and realistic images. To some extent, a bg-FID that is close to FID can manifest that “all as background” occurs. Therefore, we would expect there is sufficient gap between bg-FID and FID. It can be seen from Fig. 4 that “all as background” still occurs when using mask size loss. This is also evidenced by the poor segmentation performance and the negligible gap between FID and bg-FID (Table 1). In contrast, the utilization of mutual information loss with appropriate weight (*e.g.*  $\gamma = 1, 5, 10$ ) can successfully suppress the occurrence of “all as background”. It is noteworthy that the benefits from mutual information loss saturates and the FID score, which reflects the fidelity and diversity of synthetic images, drops as  $\gamma_{mi}$  increases beyond 1.0. We attribute this to adversarial learning being overwhelmed by mutual information maximization. To trade off between the quality of synthetic fidelity and suppression of degeneration, a proper  $\gamma_{mi}$  needs to be chosen, but the secure range of  $\gamma_{mi}$  is quite wide as the performances of  $\gamma_{mi}$  from 1 to 10 are quite close. Additional results on other datasets are available in the supplementary material.

**Prior distribution of  $c$**  We investigate the influence of different choices for prior distribution of  $c$ . In particular, normal distribution and categorical distribution are compared and the dimension (or number of categories) of  $c$  is varied. Results in Table 2 show that our method achieves similar performance with different settings, suggesting that

	CUB		Stanford Dogs		Stanford Cars	
	FID ↓	IoU ↑	FID ↓	IoU ↑	FID ↓	IoU ↑
w/o A.T.	<b>11.6</b>	68.6	67.5	53.2	13.0	58.3
w/ A.T.	12.9	<b>70.1</b>	<b>59.3</b>	<b>61.3</b>	<b>12.8</b>	<b>62.0</b>

Table 3. **Ablation study with respect to alternate training** on CUB, Stanford Dogs, and Stanford Cars. A.T. denotes alternate training.

the proposed mutual information maximization is robust to different choices for prior distribution of  $c$ . Although normal distribution leads to marginally higher performance than categorical distribution, the categorical distribution advocates the foreground-oriented clustering as discussed in 3.2 and makes our method comparable to other layered GANs [47, 5].

**The effect of alternate training** We compare the performance of our methods with and without the alternate training strategy on CUB, Stanford Dogs, and Stanford Cars at  $64 \times 64$  resolution. Table 3 shows that the alternate training helps to improve the segmentation performance without compromising the generation quality.

### 4.3. Results in Fine-Grained Category Datasets

Results of our method in three fine-grained category datasets, CUB, Stanford dogs, and Stanford Cars, are presented and compared to that of other layered GANs, FineGAN [47] and OneGAN [5]. All of the experiments are conducted at  $128 \times 128$  resolution. It should be noted that FineGAN and OneGAN requires all of or a portion of training images annotated with bounding boxes, which should be regarded as weak supervision. On the contrary, our models only learn from images, which is should be thought of as unsupervised learning. Table 4 compares our method to FineGAN and OneGAN in the aspect of generation quality, segmentation performance, and clustering performance with FID, IoU, and NMI, respectively, and Fig. 5 presents the qualitative generation results of our method.

Despite in the more challenging unsupervised learning

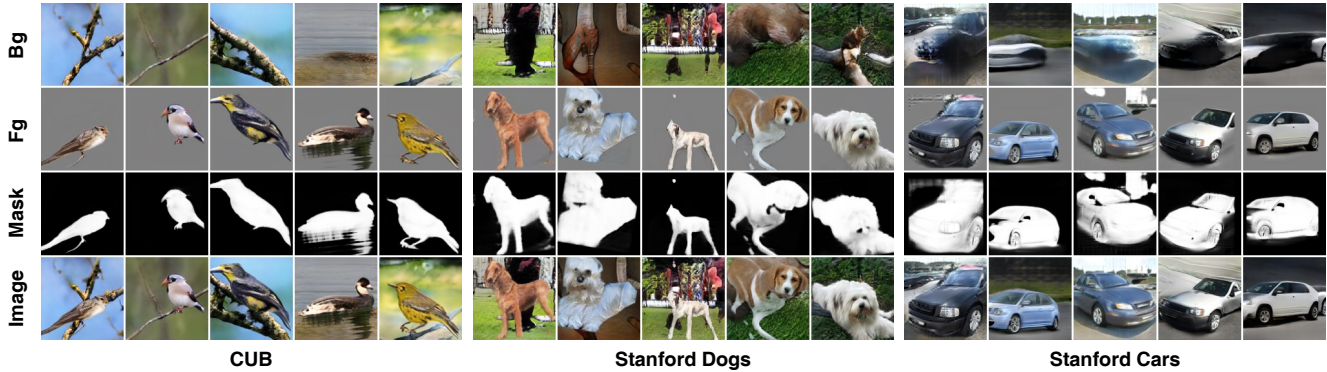


Figure 5. **Qualitative generation results** in fine-grained category datasets: CUB, Stanford Dogs, Stanford Cars. From top to bottom, the background (Bg), foreground multiplied with mask (Fg), foreground mask (Mask), and composed image (Image) are presented.

Methods	Sup.	CUB			Stanford Dogs			Stanford Cars		
		FID ↓	IoU ↑	NMI ↑	FID ↓	IoU ↑	NMI ↑	FID ↓	IoU ↑	NMI ↑
FineGAN [47]	Weak	23.0	44.5	.349	54.9	48.7	.194	24.8	53.2	.233
OneGAN [5]	Weak	20.5	55.5	.391	<b>48.7</b>	<b>69.7</b>	.211	24.2	<b>71.2</b>	.272
Ours	Unsup.	<b>12.9</b>	<b>69.7</b>	<b>.585</b>	59.3	61.3	<b>.257</b>	<b>19.0</b>	63.3	<b>.301</b>

Table 4. **Comparison to other layered GANs** in fine-grained category datasets: CUB, Stanford Dogs, and Stanford Cars. Generation quality, segmentation performance, and clustering performance are evaluated with FID, IoU, and NMI, respectively. Note that FineGAN and OneGAN is trained under weakly supervised setting, while our model is trained under unsupervised setting.

setting, our method is still able to learn disentangled foreground and background and synthesize quality paired images and segmentation masks (Fig. 5). Our method even achieves higher generation quality than FineGAN and OneGAN in CUB and Stanford Cars, with 7.6 and 5.2 higher FID, respectively. The segmentation performance of our method is still competitive to that of FineGAN and OneGAN. It is notable that our method even outperforms OneGAN in CUB with large margin (14.2 IoU) and achieves close performance to OneGAN on Stanford Dogs and Stanford Cars. We also take the trained  $E_x$  as a clustering network and evaluate its category clustering performance (Table 4). It shows that our method constantly obtains higher NMI score than FineGAN and OneGAN. This result justifies the merit of our foreground-oriented clustering as the ground truth category label is annotated based on foreground rather than background as discussed in Section 3.2.

#### 4.4. Object Segmentation Results

Our methods are compared to other unsupervised segmentation methods on CUB at  $128 \times 128$  resolution. Quantitative and qualitative results are presented in Table 5 and Fig. 6, respectively. It is noteworthy that the method proposed by Voynov *et al.* [50] requires manual inspection which should be regarded as external supervision. Our method outperform PerturbGAN [6] with a large margin, showing the merit of our proposed mutual information maximization as discussed in the ablation study. It is notable

Methods	CUB		
	ACC	IoU	mIoU
Supervised U-Net	98.0	88.8	93.2
Voynov <i>et al.</i> [50]	94.0	71.0	-
GrabCut [43] by [45]	72.3	36.0	52.3
ReDO [10]	84.5	42.6	-
PerturbGAN [6]	-	-	38.0
IEM + SegNet [45]	89.3	55.1	71.4
Melas-Kyriazi <i>et al.</i> [36]	92.1	66.4	-
Ours	<b>94.3</b>	<b>69.7</b>	<b>81.7</b>

Table 5. **Comparison to other unsupervised segmentation methods** in CUB. “-” indicates the unreported in original paper.

that performance level of our method is comparable to the state-of-the-art method [45, 36]. In particular, our method gain significantly higher performance than Melas-Kyriazi *et al.* [36] with 3.3 IoU on CUB.

#### 4.5. Comparison to VAE-based methods

Our methods are evaluated and compared to VAE-based models including MONet [8], GENESIS [15], SLOT ATTENTION [35], and GENESIS-V2 [14] in APC dataset. The quantitative and qualitative results are presented in Table 6, Fig. 7, and Fig. 8. The majority of VAE-based models perform poorly in realistic dataset as suggested in Table 6. GENESIS-V2 is the latest work that is successfully demonstrated in real-world dataset like APC. Therefore, our

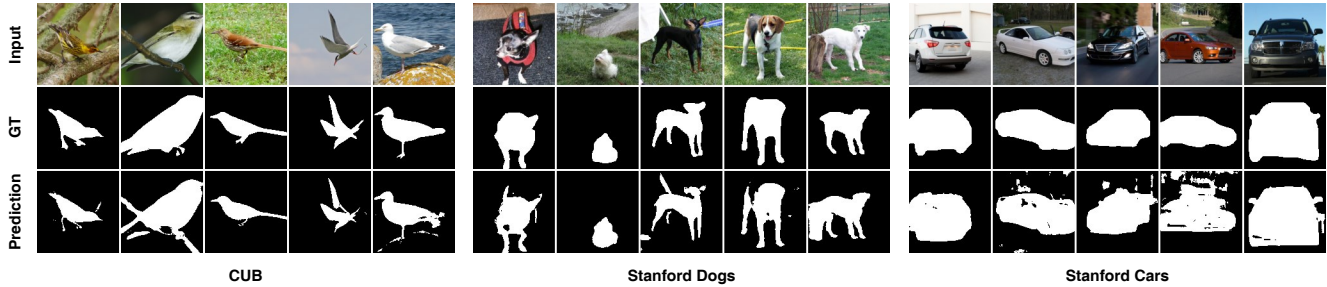


Figure 6. **Qualitative segmentation results** in CUB. The test images (Input), ground truth masks (GT), and predicted masks by segmentation network (Prediction), and are presented.

Methods		FID ↓	ARI ↑	MSC ↑
VAE-based	MONET-G [8]	269.3	0.11	0.48
	GENESIS [15]	183.2	0.04	0.29
	SLOT-ATT. [35]	-	0.03	0.25
	GENESIS-V2 [14]	245.6	0.55	0.67
GAN-based	Ours	<b>103.9</b>	<b>0.59</b>	<b>0.71</b>

Table 6. **Comparison to VAE-based methods** in APC.




Figure 7. **Qualitative generation results** in APC. Our method can generate higher quality samples than GENESIS-V2, a VAE-based model.




Figure 8. **Qualitative segmentation results** in APC. The test images (Input), ground truth masks (GT), predicted masks by GENESIS-V2 and our methods are presented.

methods are mainly compared to GENESIS-V2. To obtain the results of GENESIS-V2, we use its released codes and models<sup>1</sup>.

It can be seen from Fig. 7 and Fig. 8 that our meth-

<sup>1</sup><https://github.com/applied-ai-lab/genesis>

ods are successfully applied to APC dataset which contains multi-category objects. These results show that our methods are not limited to single-category dataset. It is not surprising that our methods, as based on GAN, generate synthetic images of significant higher quality than GENESIS-V2 does, as suggested by the more visual pleasing images in Fig. 7 and much lower FID score in Table 6. Furthermore, our methods achieve higher segmentation performance than GENESIS-V2 and predict finer segmentation results than GENESIS-V2 does, as shown in Fig. 8.

#### 4.6. Limitations

Our generative model is limited to synthesize single object and single instance images. Our method also inherits the characteristic problems of GAN training including instability and requiring certain number of unlabelled training images. In the future, we plan to extend our model into multi-instance image generation by modelling inter-object relationship and also improve the stability of training process with powerful architectures and techniques.

### 5. Conclusion

In this paper, we propose a method to improve the capability of layered GANs to disentangle foreground and background. This improved layered GAN can successfully generate synthetic datasets to train segmentation networks which can be, in return, employed to further stabilize the training of layered GANs. Our method achieves competitive generation quality and segmentation performance compared to related methods on a variety of single-object image datasets. We plan to extend our method to handle multiple object instances and the noisy data in future work.

**Acknowledgement** This work was supported by the National Key R&D Program of China under Grant 2018AAA0102801, National Natural Science Foundation of China under Grant 61620106005, and EPSRC Visual AI grant EP/T028572/1.



## References

- [1] Rameen Abdal, Peihao Zhu, Niloy Mitra, and Peter Wonka. Labels4free: Unsupervised segmentation using stylegan. *arXiv preprint arXiv:2103.14968*, 2021.
- [2] Bogdan Alexe, Thomas Deselaers, and Vittorio Ferrari. Classcut for unsupervised class segmentation. In *ECCV*, 2010.
- [3] Titas Anciukevicius, Christoph H Lampert, and Paul Henderson. Object-centric image generation with factored depths, locations, and appearances. *arXiv preprint arXiv:2004.00642*, 2020.
- [4] Relja Arandjelović and Andrew Zisserman. Object discovery with a copy-pasting gan. *arXiv preprint arXiv:1905.11369*, 2019.
- [5] Yaniv Benny and Lior Wolf. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. In *ECCV*, 2020.
- [6] Adam Bielski and Paolo Favaro. Emergence of object segmentation in perturbed generative models. In *NeurIPS*, 2019.
- [7] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. In *ICLR*, 2019.
- [8] Christopher P Burgess, Loic Matthey, Nicholas Watters, Rishabh Kabra, Irina Higgins, Matt Botvinick, and Alexander Lerchner. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*, 2019.
- [9] Yuning Chai, Victor Lempitsky, and Andrew Zisserman. Bicos: A bi-level co-segmentation method for image classification. In *ICCV*, 2011.
- [10] Mickaël Chen, Thierry Artières, and Ludovic Denoyer. Unsupervised object segmentation by redrawing. In *NeurIPS*, 2019.
- [11] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *NIPS*, 2016.
- [12] Jeff Donahue and Karen Simonyan. Large scale adversarial representation learning. In *NeurIPS*, 2019.
- [13] Sébastien Ehrhardt, Oliver Groth, Aron Monszpart, Martin Engelcke, Ingmar Posner, Niloy J. Mitra, and Andrea Vedaldi. Relate: Physically plausible multi-object scene synthesis using structured latent spaces. In *NeurIPS*, 2020.
- [14] Martin Engelcke, Oiwi Parker Jones, and Ingmar Posner. Genesis-v2: Inferring unordered object representations without iterative refinement. *arXiv preprint arXiv:2104.09958*, 2021.
- [15] Martin Engelcke, Adam R Kosiorek, Oiwi Parker Jones, and Ingmar Posner. Genesis: Generative scene inference and sampling with object-centric latent representations. In *ICLR*, 2020.
- [16] S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, David Szepesvari, koray kavukcuoglu, and Geoffrey E Hinton. Attend, infer, repeat: Fast scene understanding with generative models. In *NIPS*, 2016.
- [17] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. Ganalyze: Toward visual definitions of cognitive image properties. In *ICCV*, 2019.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [19] Klaus Greff, Raphaël Lopez Kaufmann, Rishabh Kabra, Nick Watters, Christopher Burgess, Daniel Zoran, Loic Matthey, Matthew Botvinick, and Alexander Lerchner. Multi-object representation learning with iterative variational inference. In *ICML*, 2019.
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017.
- [21] Dorit S Hochbaum and Vikas Singh. An efficient algorithm for co-segmentation. In *ICCV*, 2009.
- [22] Erik Härkönen, Aaron Hertzmann, Jaakko Lehtinen, and Sylvain Paris. Ganspace: Discovering interpretable gan controls. In *NeurIPS*, 2020.
- [23] Ali Jahanian, Lucy Chai, and Phillip Isola. On the” steerability” of generative adversarial networks. In *ICLR*, 2020.
- [24] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, 2019.
- [25] Armand Joulin, Francis Bach, and Jean Ponce. Discriminative clustering for image co-segmentation. In *CVPR*, 2010.
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019.
- [27] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *CVPR*, 2020.
- [28] Aditya Khosla, Nityananda Jayadevaprakash, Bangpeng Yao, and Fei-Fei Li. Novel dataset for fine-grained image categorization: Stanford dogs. In *CVPR Workshop*, 2011.
- [29] Wonjik Kim, Asako Kanezaki, and Masayuki Tanaka. Unsupervised learning of image segmentation based on differentiable feature clustering. *TIP*, 2020.
- [30] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *ICLR*, 2014.
- [31] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCV Workshop*, 2013.
- [32] Yiyi Liao, Katja Schwarz, Lars Mescheder, and Andreas Geiger. Towards unsupervised learning of generative models for 3d controllable image synthesis. In *CVPR*, 2020.
- [33] Zinan Lin, Kiran Thekumparampil, Giulia Fanti, and Seungwon Oh. Infogan-cr and modelcentrality: Self-supervised model training and selection for disentangling gans. In *ICML*, 2020.
- [34] Zhixuan Lin, Yi-Fu Wu, Skand Vishwanath Peri, Weihao Sun, Gautam Singh, Fei Deng, Jindong Jiang, and Sungjin Ahn. Space: Unsupervised object-oriented scene representation via spatial attention and decomposition. In *ICLR*, 2020.
- [35] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In *NeurIPS*, 2020.

- [36] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. Finding an unsupervised image segmenter in each of your deep generative models. *arXiv preprint arXiv:2105.08127*, 2021.
- [37] Thu Nguyen-Phuoc, Christian Richardt, Long Mai, Yong-Liang Yang, and Niloy Mitra. Blockgan: Learning 3d object-aware scene representations from unlabelled images. In *NeurIPS*, 2020.
- [38] Michael Niemeyer and Andreas Geiger. Campari: Camera-aware decomposed generative neural radiance fields. *arXiv preprint arXiv:2103.17269*, 2021.
- [39] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *CVPR*, 2021.
- [40] Yassine Ouali, Céline Hudelot, and Myriam Tami. Autoregressive unsupervised image segmentation. In *ECCV*, 2020.
- [41] Antoine Plumerault, Hervé Le Borgne, and Céline Hudelot. Controlling generative models with continuous factors of variations. In *ICLR*, 2020.
- [42] Tal Remez, Jonathan Huang, and Matthew Brown. Learning to segment via cut-and-paste. In *ECCV*, 2018.
- [43] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. "grabcut" interactive foreground extraction using iterated graph cuts. *TOG*, 2004.
- [44] Carsten Rother, Tom Minka, Andrew Blake, and Vladimir Kolmogorov. Cosegmentation of image pairs by histogram matching-incorporating a global constraint into mrfs. In *CVPR*, 2006.
- [45] Pedro Savarese, Sunnie SY Kim, Michael Maire, Greg Shakhnarovich, and David McAllester. Information-theoretic segmentation by inpainting error maximization. *arXiv preprint arXiv:2012.07287*, 2020.
- [46] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. Interpreting the latent space of gans for semantic face editing. In *CVPR*, 2020.
- [47] Krishna Kumar Singh, Utkarsh Ojha, and Yong Jae Lee. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In *CVPR*, 2019.
- [48] Nurit Spingarn-Eliezer, Ron Banner, and Tomer Michaeli. Gan steerability without optimization. *ICLR*, 2021.
- [49] Andrey Voynov and Artem Babenko. Unsupervised discovery of interpretable directions in the gan latent space. In *ICML*, 2020.
- [50] Andrey Voynov, Stanislav Morozov, and Artem Babenko. Big gans are watching you: Towards unsupervised object segmentation with off-the-shelf generative models. *arXiv preprint arXiv:2006.04988*, 2020.
- [51] 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.
- [52] Jianwei Yang, Anitha Kannan, Dhruv Batra, and Devi Parikh. Lr-gan: Layered recursive generative adversarial networks for image generation. *arXiv preprint arXiv:1703.01560*, 2017.
- [53] Andy Zeng, Kuan-Ting Yu, Shuran Song, Daniel Suo, Ed Walker Jr, Alberto Rodriguez, and Jianxiong Xiao. Multi-view self-supervised deep learning for 6d pose estimation in the amazon picking challenge. In *ICRA*, 2017.
- [54] Yuxuan Zhang, Wenzheng Chen, Huan Ling, Jun Gao, Yinan Zhang, Antonio Torralba, and Sanja Fidler. Image gans meet differentiable rendering for inverse graphics and interpretable 3d neural rendering. In *ICLR*, 2020.
- [55] 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 *CVPR*, 2021.