

# ILSGAN: Independent Layer Synthesis for Unsupervised Foreground-Background Segmentation

Qiran Zou, Yu Yang, Wing Yin Cheung, Chang Liu, Xiangyang Ji

Tsinghua University, BNRist

qiranzou@gmail.com, yang-yu16@foxmail.com, zhangyx20@mails.tsinghua.edu.cn, {liuchang2022, xyji}@tsinghua.edu.cn

## Abstract

Unsupervised foreground-background segmentation aims at extracting salient objects from cluttered backgrounds, where Generative Adversarial Network (GAN) approaches, especially layered GANs, show great promise. However, without human annotations, they are typically prone to produce foreground and background layers with non-negligible semantic and visual confusion, dubbed “information leakage”, resulting in notable degeneration of the generated segmentation mask. To alleviate this issue, we propose a simple-yet-effective explicit layer independence modeling approach, termed Independent Layer Synthesis GAN (ILSGAN), pursuing independent foreground-background layer generation by encouraging their discrepancy. Specifically, it targets minimizing the mutual information between visible and invisible regions of the foreground and background to spur inter-layer independence. Through in-depth theoretical and experimental analyses, we justify that explicit layer independence modeling is critical to suppressing information leakage and contributes to impressive segmentation performance gains. Also, our ILSGAN achieves strong state-of-the-art generation quality and segmentation performance on complex real-world data.

## Introduction

Foreground-background segmentation, also called foreground extraction, targets dividing foregrounds containing object instances from the background in an image. It is an elementary and special case of instance segmentation (He et al. 2017) that the ability to distinguish foreground-background is vital for complex scene understanding (Krishna et al. 2017; Zellers et al. 2018; Yang et al. 2018). Since it is error-prone, time-consuming, and expensive to densely label a large quantity of data, unsupervised foreground-background segmentation, including GAN-based (Melas-Kyriazi et al. 2022; Yang et al. 2017; Bielski and Favaro 2019), VAE-based (Yu et al. 2021; Engelcke, Parker Jones, and Posner 2021), and MI clustering-based (Ji, Henriques, and Vedaldi 2019; Ouali, Hudelot, and Tami 2020) methods, attracts growing attention in the community.

Among these approaches, GAN-based methods show significant representation disentanglement ability which con-

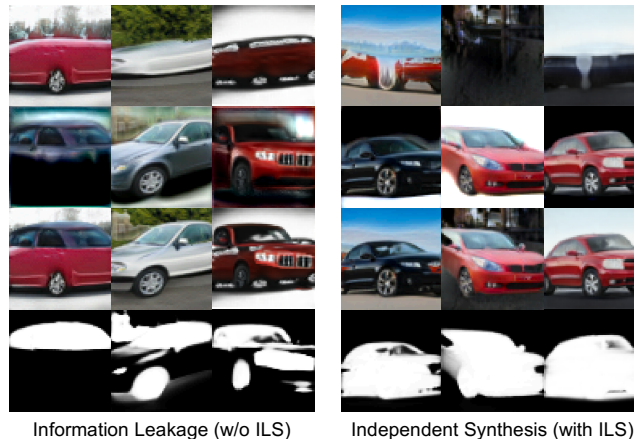


Figure 1: From top to bottom: synthetic background, foreground, composed image, and corresponding mask. Left: synthetic samples with information leakage problem when our ILS algorithm is not adopted. Right: generated samples with independent foreground and background layers and more accurate masks when ILS is employed.

tributes to accurate segmentation performance. Leveraging the intrinsic compositional nature of natural images, layered GANs have made exceptional progress for unsupervised foreground-background segmentation. Specifically, with multilayer structures of generators and corresponding inductive constraints, PerturbGAN (Bielski and Favaro 2019) considers that the foreground can be shifted to a certain extent without affecting the image fidelity. FineGAN (Singh, Ojha, and Lee 2019) constrains the background layer with an extra background discriminator.

However, during layered GANs’ training, foreground and background layers typically contain increasing semantic information about each other, resulting in their growing visual similarity. As shown in Fig. 1 (left), there exists semantic confusion that background contains a wheel and foreground contains trees, and visual similarity that foreground and background have similar colors and textures. This phenomenon, termed “information leakage”, discourages the mask layer from distinguishing the foreground and background and impairs segmentation performance severely.

Though, it is still largely under-studied by now.

Intuitively, the semantic and visual confusion can be characterized as the excessive correlation between foreground and background layers, indicating a lack of independence. Therefore, a feasible way to solve information leakage problem is to enhance the independence between layers. However, such independence has not been valued in previous works, which restricts the ability of layered GANs to distinguish foreground from background.

In this work, we propose **Independent Layer Synthesis GAN (ILSGAN)**, based on layered GAN and fueled by information theory, which generates images composed of independent layers and corresponding high-quality masks for unsupervised foreground-background segmentation. Specifically, ILSGAN minimizes the mutual information (MI) between foreground and background layers to improve the interlayer independence. To circumvent the difficulty of deriving the sample density in MI, it employs a variational upper bound and reduce MI by minimizing the upper bound. Furthermore, ILSGAN can specifically optimize the mask layer to learn partitions with less MI, which effectively enhances segmentation performance.

We perform experiments on various datasets including Cars (Krause et al. 2013), CUB (Wah et al. 2011), and Dogs (Khosla et al. 2011). Our ILSGAN successfully suppresses information leakage and efficiently generates complete and accurate segmentation masks. Compared with existing unsupervised segmentation methods, ILSGAN shows a significant advantage in segmentation performance.

We summarize our contributions as follows:

- We find that interlayer semantic and visual confusion, termed information leakage, exists in layered GAN and undermines the segmentation performance.
- We propose ILSGAN, improving interlayer independence by minimizing mutual information, which suppresses information leakage while powerfully enhancing segmentation performance.
- We achieve solid state-of-the-art performances on commonly used unsupervised foreground-background segmentation datasets, with IoU results of 81.4, 72.1, and 86.3 on Cars, CUB and Dogs, beating the closet competitors (70.8, 69.7, 72.3) with significant margins.

## Related Work

### Generative Adversarial Networks

GANs (Goodfellow et al. 2014) provide an unsupervised way for learning the representation of data and generating high-fidelity synthetic images. They have been successfully used for image editing (Ling et al. 2021; Kim et al. 2021), style transfer (Karras, Laine, and Aila 2019; Karras et al. 2020b), controllable image generation (Chen et al. 2016; Singh, Ojha, and Lee 2019; Benny and Wolf 2020), representation learning (Donahue and Simonyan 2019), image classification (Chen et al. 2016), super-resolution (Ledig et al. 2017), and generating annotations (Zhang et al. 2021; Yang et al. 2021; Abdal et al. 2021).

Unlike traditional GANs, a series of works (Chen et al. 2016; Singh, Ojha, and Lee 2019; Benny and Wolf 2020;

Yang et al. 2022) disentangle GANs by decomposing the input latent into a main subset and a controlling subset for disentanglement. The controlling subset can be used to classify the synthetic images and disentangle the semantic meaning like rotation and width.

On the other hand, layered GANs disentangle the representation explicitly by introducing a multi-layer structure that corresponds to the elements in nature scenes. They first generate multiple layers and then blend these layers into a composed image. It gives layered GANs the ability to decompose visual scenes and makes it possible for layered GANs to achieve fine-grained controllable generation (Singh, Ojha, and Lee 2019; Benny and Wolf 2020; Yang et al. 2022), video generation (Ehrhardt et al. 2020), and unsupervised segmentation (Yang et al. 2017; Bielski and Favaro 2019).

### Unsupervised Foreground-Background Segmentation

A typical line of recent unsupervised methods partition images into foreground and background with layered GANs. Some works spur foreground extraction with foreground perturbation (Yang et al. 2017; Bielski and Favaro 2019; Yang et al. 2022) or shape prior (Kim and Hong 2021), distinguish foreground and background with an additional foreground discriminator (Singh, Ojha, and Lee 2019; Benny and Wolf 2020), or decompose images by redrawing real images (Chen, Artières, and Denoyer 2019). These methods show impressive gains in segmentation performance. However, the independence between foreground and background is overlooked which leads to information leakage and degradation of segmentation performance.

On another front, methods based on clustering attain unsupervised segmentation through self-supervised learning with mutual information maximization (Ji, Henriques, and Vedaldi 2019; Ouali, Hudelot, and Tami 2020) or contrastive learning (Cho et al. 2021; Van Gansbeke et al. 2021; Choudhury et al. 2021; Hwang et al. 2019). However, when these methods are applied to binary segmentation of foreground and background, they show disadvantages (Benny and Wolf 2020) in segmentation performance compared to methods based on layered GANs.

Apart from these works, other works extract foreground by modeling with VAE (Yu et al. 2021; Burgess et al. 2019; Engelcke et al. 2020; Engelcke, Parker Jones, and Posner 2021; Locatello et al. 2020), searching directions in latent space (Melas-Kyriazi et al. 2022; Voynov, Morozov, and Babenko 2021), or maximizing inpainting error (Savarese et al. 2021). Our methods leverage the high-fidelity and foreground extraction ability of layered GANs and at the same time improve interlayer independence which contributes to more accurate and clear segmentations.

## Method

### Layer Synthesis

Our ILSGAN synthesizes independent foreground, background, and mask layers with a layered structure, which are

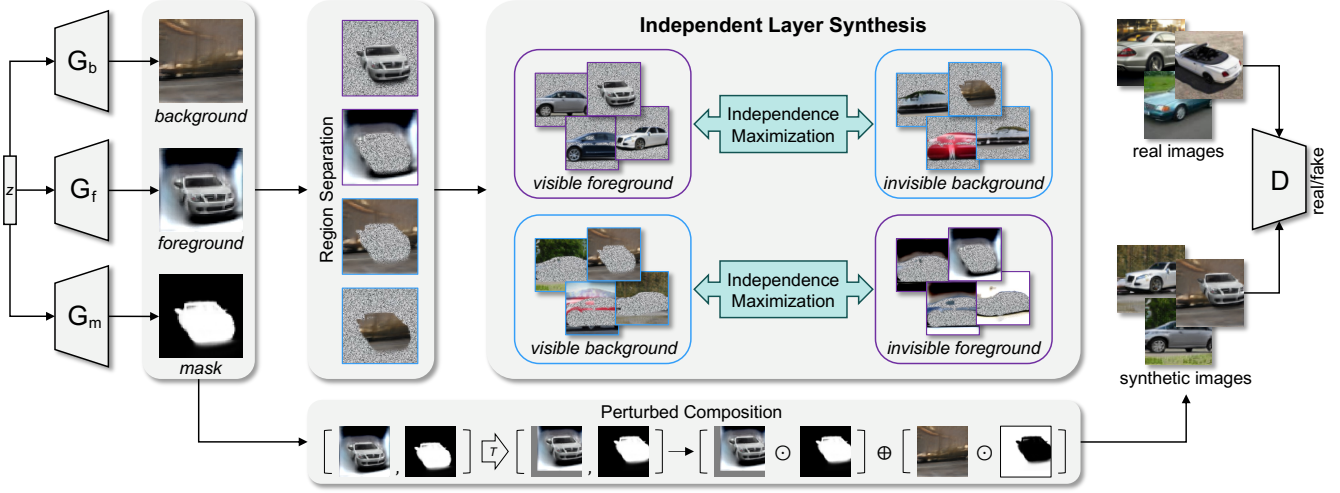


Figure 2: Framework of ILSGAN.  $G_b$ ,  $G_f$  and  $G_m$  are the generators sharing latent code  $z$  and producing background, foreground, and mask layers. The generated layers are fed into two branches. One branch is the perturbed composition and adversarial learning against a discriminator  $D$ . The other includes separating regions according to visibility in the final image and our Independent Layer Synthesis algorithm which maximizes the independence between visible and invisible regions of foreground and background layers.

finally blended into photo-realistic images and corresponding binary masks. Formally, let  $G_f$ ,  $G_b$ , and  $G_m$  denote a foreground generator, a background generator, and a mask generator, respectively. Given a latent variable  $\mathbf{z} \in \mathcal{Z}$  drawn from a prior distribution,  $\mathbf{z} \sim p_z$ , the generation process can be written as

$$\begin{aligned} \mathbf{f} &= G_f(\mathbf{z}), \quad \mathbf{b} = G_b(\mathbf{z}), \quad \mathbf{m} = G_m(\mathbf{z}), \\ \mathbf{x} &= \mathbf{m} \odot \mathbf{f} + (\mathbf{1} - \mathbf{m}) \odot \mathbf{b}, \end{aligned} \quad (1)$$

where  $\mathbf{f}$ ,  $\mathbf{b}$ , and  $\mathbf{m}$  denote foreground, background, and mask, respectively, which constitutes an image  $\mathbf{x}$  by a composition step (Eq. 1).

The full synthesis process is wrapped as  $G : \mathcal{Z} \rightarrow \mathcal{X}$ , which can be learned via adversarial learning against a discriminator  $D : \mathcal{X} \rightarrow [0, 1]$  to tell real or fake images as

$$\begin{aligned} \min_G \max_D V(G, D) &= \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} [\log D(\mathbf{x})] + \\ &\quad \mathbb{E}_{\mathbf{z} \sim p_z} [\log (1 - D(G(\mathbf{z})))] , \end{aligned} \quad (2)$$

where  $\mathcal{Z} \subset \mathbb{R}^d$  and  $\mathcal{X} \subset \mathbb{R}^{3 \times H \times W}$  denote the latent variable space and the image space.

Since vanilla adversarial learning (Eq. 2) lacks incentive for non-trivial layerwise generation, inductive biases such as network structure, hand-crafted regularization, *etc.* are essential. Hence, we briefly review the inductive biases inherited from prior work as follows. First, the generators are parameterized with deep convolutional neural networks. Concretely,  $G_f$  and  $G_b$  are different networks;  $G_m$  shares backbone networks with  $G_f$  and has a separate head to output masks. More details about networks structure are available in Sec. . Second, to counter that  $G_m$  constantly outputs  $\mathbf{m} = \mathbf{1}$ , we replace composition step in Eq. 1 with perturbed composition (Bielski and Favaro 2019),

$$\mathbf{x} = T(\mathbf{m}) \odot T(\mathbf{f}) + (\mathbf{1} - T(\mathbf{m})) \odot \mathbf{b}, \quad (3)$$

where  $T$  is an operation that randomly shifts masks and foreground pixels with a small amount. Third,  $G_m$  is penalized if it outputs masks that have smaller area than a pre-defined threshold. Finally, binary masks are encouraged with an additional loss function. The loss functions are defined as follows,

$$\begin{aligned} \mathcal{L}_m &= \mathbb{E}_{\mathbf{z} \sim p_z, \mathbf{m} = G_m(\mathbf{z})} \max\{0, \eta - \frac{1}{HW} \|\mathbf{m}\|_1\}, \\ \mathcal{L}_b &= \mathbb{E}_{\mathbf{z} \sim p_z, \mathbf{m} = G_m(\mathbf{z})} \frac{1}{HW} \sum_{u \in \Omega} \min\{m_u, 1 - m_u\}, \end{aligned} \quad (4)$$

where  $\mathcal{L}_m$  and  $\mathcal{L}_b$  respectively denote the aforementioned loss functions that encourage non-zero and binary masks.  $\eta$  represents the tolerable minimal mask area,  $m_u$  denotes the element in  $\mathbf{m}$  indexed by  $u$  and  $\Omega = \{1, \dots, H\} \times \{1, \dots, W\}$  denotes the index set.

### Independent Layer Synthesis

We propose independent layer synthesis (ILS) loss,  $\mathcal{L}_{\text{ILS}}$ , to improve the independence between foreground and background layers and thereby suppress information leakage. ILS loss, as well as aforementioned losses, are integrated into GAN training, resulting in our ILSGAN framework which is illustrated in Fig. 2. The optimization of generator and discriminator are alternated (Goodfellow et al. 2014) as follows,

$$\begin{aligned} \min_D & -V(\bar{G}, D) + \gamma R_1(D), \\ \min_G & V(G, \bar{D}) + \lambda_m \mathcal{L}_m + \lambda_b \mathcal{L}_b + \lambda_{\text{ils}} \mathcal{L}_{\text{ILS}}, \end{aligned} \quad (5)$$

where  $V(\bar{G}, D)$  and  $V(G, \bar{D})$  respectively denote the variants of  $V(G, D)$  when  $G$  or  $D$  is fixed,  $R_1(D)$  denotes the gradient penalty (Mescheder, Geiger, and Nowozin 2018)

that stabilizes the GAN training with weight  $\gamma$ . And  $\lambda_m$ ,  $\lambda_b$ , and  $\lambda_{vis}$  denote the corresponding loss weights. The construction of ILS loss is described below.

**Measuring layer independence with MI** As observed in practice, information leakage issue (Fig. 1) frequently occurs in the unsupervised learning of layer synthesis. This problem can be formally stated as the unexpected low interlayer independence. Therefore, to suppress information leakage, we are motivated to increase the independence between foreground and background layers. By information theory, this independence can be measured with mutual information (MI),

$$\begin{aligned} I(\mathbf{b}; \mathbf{f}) &= \int p(\mathbf{b}, \mathbf{f}) \log \frac{p(\mathbf{b}, \mathbf{f})}{p(\mathbf{b})p(\mathbf{f})} d\mathbf{b}d\mathbf{f} \\ &= H(\mathbf{b}) + \mathbb{E}_{p(\mathbf{b}, \mathbf{f})} \log p(\mathbf{b}|\mathbf{f}), \end{aligned} \quad (6)$$

where  $\mathbf{b}, \mathbf{f}$  are drawn from the joint probability  $p(\mathbf{b}, \mathbf{f})$  induced by our layer synthesis model. To this end, one straightforward way to increase interlayer independence is to directly minimize  $I(\mathbf{b}; \mathbf{f})$ . Nonetheless, estimation of  $I(\mathbf{b}; \mathbf{f})$  is challenging since the implicit model (*i.e.* GAN) on which our layer synthesis model is built only supports the sampling process and is unable to yield density estimation of samples. Therefore, we instead appeal to minimizing the upper bound of  $I(\mathbf{b}; \mathbf{f})$  as follows.

**Minimizing MI upper bound** To circumvent the difficulty of estimating sample density, *i.e.*  $p(\mathbf{b}, \mathbf{f})$ ,  $p(\mathbf{b})$ , and  $p(\mathbf{f})$  in Eq. 6, we consider the variational CLUB (Cheng et al. 2020) upper bound of MI,

$$\begin{aligned} I(\mathbf{b}; \mathbf{f}) \leq I_{VC}(\mathbf{b}, \mathbf{f}) &= \mathbb{E}_{p(\mathbf{b}, \mathbf{f})} [\log q(\mathbf{b}|\mathbf{f})] \\ &\quad - \mathbb{E}_{p(\mathbf{b})} \mathbb{E}_{p(\mathbf{f})} [\log q(\mathbf{b}|\mathbf{f})], \end{aligned} \quad (7)$$

where  $q(\mathbf{b}|\mathbf{f})$  is a variational approximation of conditional probability  $p(\mathbf{b}|\mathbf{f})$ . This inequality Eq. 7 holds under mild conditions (Cheng et al. 2020). It is noteworthy that computing  $I_{VC}(\mathbf{b}, \mathbf{f})$  only requires samples from generative models and estimating density under the variational distribution, significantly easing the computation. We further choose a Laplace distribution with identity covariance (Savarese et al. 2021) as the variational distribution, *i.e.*  $q(\mathbf{b}|\mathbf{f}) = \mathcal{L}(\mathbf{b}; \mathbf{f}, \mathbf{I}) \propto \exp(-\|\mathbf{b} - \mathbf{f}\|_1)$ . We have

$$\begin{aligned} I_{VC}(\mathbf{b}, \mathbf{f}) &= \mathbb{E}_{p(\mathbf{b}, \mathbf{f})} [-\|\mathbf{b} - \mathbf{f}\|_1] \\ &\quad - \mathbb{E}_{p(\mathbf{b})} \mathbb{E}_{p(\mathbf{f})} [-\|\mathbf{b} - \mathbf{f}\|_1] + \text{Const}. \end{aligned} \quad (8)$$

Given a mini-batch of samples  $\{(\mathbf{b}^{(i)}, \mathbf{f}^{(i)})\}_{i=1}^N$ , neglecting the constant term, the unbiased  $N$ -sample estimation of  $I_{VC}(\mathbf{b}, \mathbf{f})$  writes

$$\widehat{I}_{VC}(\mathbf{b}, \mathbf{f}) = \frac{1}{N} \sum_{i=1}^N -\|\mathbf{b}^{(i)} - \mathbf{f}^{(i)}\|_1 + \|\mathbf{b}^{(k_i)} - \mathbf{f}^{(i)}\|_1, \quad (9)$$

where  $N$  denotes the number of samples,  $\mathbf{b}^{(i)} = G_b(\mathbf{z}^{(i)})$  and  $\mathbf{f}^{(i)} = G_f(\mathbf{z}^{(i)})$  are generated layers by forwarding randomly sampled latent variables  $\mathbf{z}^{(i)} \sim p_z$  to our generative models, and  $k_i$  denotes a uniformly sampled index in  $\{1, 2, \dots, N\}$ . To this end, we minimize  $\widehat{I}_{VC}(\mathbf{b}, \mathbf{f})$  to indirectly reduce the  $I(\mathbf{b}; \mathbf{f})$ .

**ILS Loss** Moreover, in order to optimize the mask to find the segmentation with less mutual information, we separate foreground and background layers respectively according to their visibility in synthetic image  $\mathbf{x}$ , Fig. 2,

$$\begin{aligned} \mathbf{f}_{vis} &= \mathbf{f} \odot \mathbf{m}, & \mathbf{f}_{inv} &= \mathbf{f} \odot (\mathbf{1} - \mathbf{m}), \\ \mathbf{b}_{vis} &= \mathbf{b} \odot (\mathbf{1} - \mathbf{m}), & \mathbf{b}_{inv} &= \mathbf{b} \odot \mathbf{m}, \end{aligned} \quad (10)$$

where  $\mathbf{f}_{vis}$ ,  $\mathbf{f}_{inv}$ ,  $\mathbf{b}_{vis}$ , and  $\mathbf{b}_{inv}$  denote the visible and invisible regions of foreground and background layers. Thus, our ILS loss minimizes MI of two pairs of visible and invisible regions,

$$\mathcal{L}_{ILS} = I(\mathbf{b}_{inv}; \mathbf{f}_{vis}) + I(\mathbf{b}_{vis}; \mathbf{f}_{inv}). \quad (11)$$

Estimating MI with  $\widehat{I}_{VC}$ , our ILS-MI loss is as follows,

$$\begin{aligned} \mathcal{L}_{ILS}^{MI} &= \widehat{I}_{VC}(\mathbf{b}_{inv}; \mathbf{f}_{vis}) + \widehat{I}_{VC}(\mathbf{b}_{vis}; \mathbf{f}_{inv}) \\ &= \frac{1}{N} \sum_{i=1}^N \left[ -\|\mathbf{b}_{inv}^{(i)} - \mathbf{f}_{vis}^{(i)}\|_1 + \|\mathbf{b}_{inv}^{(k_i)} - \mathbf{f}_{vis}^{(i)}\|_1 \right] \\ &\quad + \frac{1}{N} \sum_{j=1}^N \left[ -\|\mathbf{b}_{vis}^{(j)} - \mathbf{f}_{inv}^{(j)}\|_1 + \|\mathbf{b}_{vis}^{(k_j)} - \mathbf{f}_{inv}^{(j)}\|_1 \right]. \end{aligned} \quad (12)$$

## Discussion

As the information leakage problem can be interpreted as there existing unexpected interlayer visual similarity, *e.g.* car body textures or wheels present in the background layer (Fig. 1), one intuitive way to suppress this issue is to directly minimize the visual similarity between foreground and background layers. Based on this motivation, by measuring the visual similarity using  $L1$  distance, the ILS- $L1$  loss can be constructed,

$$\mathcal{L}_{ILS}^{L1} = -\|\mathbf{b}_{inv} - \mathbf{f}_{vis}\|_1 - \|\mathbf{b}_{vis} - \mathbf{f}_{inv}\|_1. \quad (13)$$

In practice, we observe that ILS- $L1$  can help to mitigate the information leakage issue and improve the segmentation performance. Moreover, ILS- $L1$  is also able to contribute a bit to the decrease of MI between foreground and background layers. Despite so, we believe the effect of this intuitive method is limited and our ILS-MI loss is a more fundamental approach that achieves more impressive results. See Sec. for detailed experiments.

## Experiments

### Settings

**Datasets** We evaluate our method on Stanford Cars (Cars) (Krause et al. 2013), Stanford Dogs (Dogs) (Khosla et al. 2011), and Caltech-UCSD Birds 200-2011 (CUB) (Wah et al. 2011). We follow the train-test split from (Yu et al. 2021) to split Cars and Dogs dataset, leading to 6218 training images and 6104 testing images for Cars as well as 3286 training and 1738 testing images for Dogs. The CUB dataset is split into 10k images for training and 1k images for testing following (Chen, Artières, and Denoyer 2019). We use the approximated ground-truth by (Yu et al. 2021) on Cars and Dogs, and manually labeled ground-truth masks on CUB to evaluate the performance on test set.

Methods	Sup.	Cars			CUB			Dogs		
		FID↓	IoU↑	DICE↑	FID↓	IoU↑	DICE↑	FID↓	IoU↑	DICE↑
FineGAN (Singh, Ojha, and Lee 2019) <sup>†</sup>	Weak	24.8	53.2	60.3	23.0	44.5	56.9	54.9	48.7	59.3
OneGAN (Benny and Wolf 2020)	Weak	24.2	71.2	82.6	20.5	55.5	69.2	<b>48.7</b>	71.0	81.7
IEM (Savarese et al. 2021)	Uns.	-	-	-	-	55.1	68.7	-	-	-
ReDO (Chen, Artières, and Denoyer 2019) <sup>‡</sup>	Uns.	-	52.5	68.6	-	47.1	61.7	-	52.8	67.9
Impr.LGAN (Yang et al. 2022)	Uns.	19.0	64.6	-	12.9	69.7	-	59.3	61.3	-
DRC (Yu et al. 2021)	Uns.	-	70.8	82.5	-	54.6	69.4	-	72.3	83.6
ILSGAN (Ours)	Uns.	<b>9.0</b>	<b>81.2</b>	<b>89.3</b>	<b>8.7</b>	<b>72.1</b>	<b>82.8</b>	49.5	<b>86.3</b>	<b>92.5</b>

Table 1: Quantitative comparison to related methods with respect to generation quality and segmentation. <sup>†</sup>: Results are reported by OneGAN (Benny and Wolf 2020). <sup>‡</sup>: Results are reported by DRC(Yu et al. 2021).



Figure 3: Qualitative generation results (128×128). Left: Cars, middle: CUB, right: Dogs. From top to bottom: background, foreground, composed image, and foreground mask.

**Evaluation** As our ILSGAN is only able to generate images with segmentation masks, we train a post hoc UNet (Ronneberger, Fischer, and Brox 2015) on these synthetic data at 64×64 resolution to evaluate its segmentation performance. For segmentation evaluation at 128×128 resolution, we upsample the predicted masks to 128×128 for computing metrics. The trained UNet is tested on real test set and the DICE coefficient and intersection over union (IoU) metrics are reported. To more accurately measure the independence between synthesized foreground and background layers during test phase, we employ a heavier neural estimator, MINE (Belghazi et al. 2018), which is a lower bound of mutual information and can serve as a reasonable MI estimator. Similar to Eq. 11, we respectively compute the MI of regions divided by segmentation masks and sum them up, which is reported as “MI” in the following sections. Besides, we also compute the Fréchet Inception Distance (FID) (Heusel et al. 2017) of 10k synthetic images against the whole training images to quantitatively evaluate the generation quantity.

**Implementation Details** We employ StyleGAN2 (Karras et al. 2020b) with adaptive discriminator augmentation (Karras et al. 2020a) to build our ILSGAN. We simply adopt a generator with 4-channels output to generate foreground and mask layers at the same time: 3 channels for foreground and image and 1 channel for the mask. This 4-channels generator and background generator each contain a mapping network that maps 512-D latent code  $z$  to a 512-D intermediate latent code. Hyperparameters are consistently set across datasets: non-empty mask loss weight  $\lambda_m = 2$ , mask size threshold

$\eta = 0.25$ , mask binarization loss weight  $\lambda_b = 2$ , and ILS-MI loss weight  $\lambda_{ils} = 1$ . Our ILSGAN is optimized using Adam optimizer with initial learning rate 0.0025 and beta parameters 0 and 0.99 until 8M images have been shown to the discriminator. During training, we center crop the images to get square training inputs. In the testing phase, we select the checkpoint with lowest FID for evaluation. Following DRC (Yu et al. 2021), we crop the images and masks according to ground truth object bounding boxes and resize them to square shape as test data. Reported results in Sec. are conducted at 128×128 resolution while all analytical experiments in Sec. are conducted at 64×64 resolution. Experiments are conducted on single RTX 2080Ti GPU.

## Benchmark Results

Our method is compared with other related unsupervised and weakly-supervised foreground and background segmentation methods concerning segmentation performance and generation quality on Cars, CUB, and Dogs datasets. As GAN-based method is notoriously unstable in the training process, the performance of ILSGAN is reported as the average of three times experiments. Table. 1 and Fig. 3 show the quantitative and qualitative results from which we have the following observations.

(1) ILSGAN achieves superior unsupervised segmentation performance which surpasses competitors with large margins. On Cars and Dogs, ILSGAN outperforms previously the best method, DRC (Yu et al. 2021), by 10.4/6.8 IoU/DICE and 14.0/8.9 IoU/DICE, respectively. On CUB, ILSGAN gains 2.4 IoU improvement over previously the

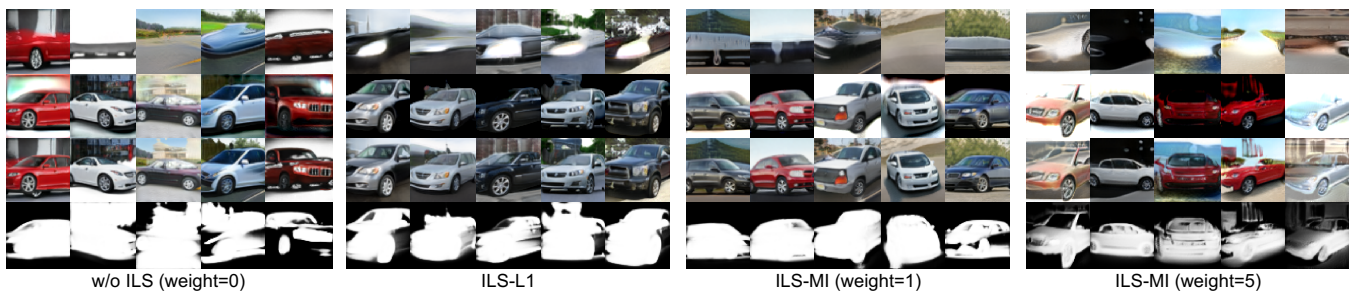


Figure 4: Qualitative ablation results of ILSGAN for settings in Table. 2 on Cars. From top to bottom: background, foreground, composed image, and corresponding mask.

Loss	$\lambda_{\text{ils}}$	Cars				CUB				Dogs			
		FID↓	IoU↑	DICE↑	MI↓	FID↓	IoU↑	DICE↑	MI↓	FID↓	IoU↑	DICE↑	MI↓
w/o ILS	—	7.7	80.6	88.9	7.32	8.9	68.6	80.1	7.40	37.9	81.7	89.6	7.33
ILS-L1	1.0	7.8	82.1	89.9	6.78	<b>8.7</b>	72.5	83.1	6.64	40.1	82.9	90.4	7.43
	0.2	7.7	80.4	88.8	6.61	8.8	70.0	81.2	7.14	38.2	82.8	90.3	7.36
	0.5	8.0	82.2	89.9	6.98	8.8	71.3	82.2	7.03	<b>37.4</b>	<b>84.0</b>	<b>91.1</b>	6.80
ILS-MI	* 1.0	<b>7.3</b>	<b>83.5</b>	<b>90.8</b>	6.09	8.9	72.3	82.9	6.70	38.7	83.8	90.9	<b>6.75</b>
	2.0	11.5	78.5	87.6	<b>6.06</b>	9.0	<b>73.7</b>	<b>84.0</b>	<b>6.35</b>	40.8	80.3	88.7	7.05
	5.0	22.2	61.7	75.9	6.70	17.4	49.3	61.3	7.51	55.3	44.2	60.3	8.74

Table 2: Quantitative evaluations of ILS losses ( $\lambda_{\text{ils}}$  and \* denotes the weight of loss and the default setting respectively).

best method, Impr.LGAN (Yang et al. 2022).

(2) ILSGAN even significantly outperforms other layered GANs, OneGAN (Benny and Wolf 2020) and FineGAN (Singh, Ojha, and Lee 2019), which requires weak supervision such as bounding box annotation. Note that ILSGAN respectively gains 10.0, 16.6, and 15.3 IoU improvement on Cars, CUB, and Dogs compared to OneGAN.

(3) ILSGAN can synthesize images of high quality, achieving 9.0, 8.7, and 49.5 FID on Cars, CUB, and Dogs, respectively, either significantly outperforms or performs on par with other layered GANs such as FineGAN, OneGAN, and Impr.LGAN. Examples in Fig. 3 show the high fidelity and clear foreground-background disentanglement of the generated images.

## Analysis

**The effect of ILS-MI** We study the effect of ILS-MI loss with its weight  $\lambda_{\text{ils}} \in \{0.0, 0.2, 0.5, 1.0, 2.0, 5.0\}$ , where  $\lambda_{\text{ils}} = 0.0$  corresponds to “w/o ILS”. The quantitative and qualitative results are presented in Table. 2 and Fig. 4 respectively. Compared to the baseline method (w/o ILS), our method with ILS-MI ( $\lambda_{\text{ils}} = 1$ ) reduces MI by 1.2, 0.7, and 0.6 on Cars, CUB, and Dogs, respectively, showing improved layer independence. By visualization in Fig. 4, while information leakage emerges in the results of the baseline method (e.g. car wheels and texture are present in the background), the foreground and background layers in the results of our method have less visual confusion, suggesting information leakage issue is effectively suppressed by our ILS-

MI. Accordingly, our ILS-MI brings clear improvement in segmentation performance with 2.9 IoU on Cars, 3.7 IoU on CUB, and 2.1 IoU on Dogs. These results justify that pursuing independent layer synthesis with ILS-MI can suppress the information leakage and thereby significantly improve segmentation performance.

We also notice that raising  $\lambda_{\text{ils}}$  too high has a risk of overwhelming the effect of other losses such as generation quality by adversarial loss and binary mask by binarization loss. It can be observed that the loss weight of ILS-MI greater than 2 leads to significantly degraded generation quality (high FID), non-binary foreground masks (Fig. 4), and accordingly the segmentation performance decrease. Therefore, we use a secure loss weight  $\lambda_{\text{ils}} = 1$  to pursue independent layers without toppling down the synthesis.

**Comparison of ILS-MI and ILS-L1** We further compare the effect of our ILS-MI and the intuitive alternative, ILS-L1. The quantitative and qualitative results of ILS-L1 are also presented in Table. 2 and Fig. 4. We observe that ILS-L1 can achieve similar effects to ILS-MI: it reduces MI (only on Cars and CUB), suppresses information leakage (Fig. 4), and improves the segmentation performance. Despite so, ILS-L1 still shows the following drawbacks compared to ILS-MI. (1) Fig. 4 shows that artifacts (i.e. white block in the center of the background) emerge in the generated background layer of ILS-L1, possibly due to the invisible background being optimized towards the opposite color of visible foreground. In contrast, ILS-MI does not suffer from this prob-

	Dist	Sep?	OptVis?	OptM?	Cars				CUB				Dogs			
					FID↓	IoU↑	DICE↑	MI↓	FID↓	IoU↑	DICE↑	MI↓	FID↓	IoU↑	DICE↑	MI↓
A			w/o $\mathcal{L}_{\text{ILS-MI}}$		7.7	80.6	88.9	7.32	8.9	68.6	80.1	7.40	37.9	81.7	89.6	7.33
*B		✓	✓	✓	<b>7.3</b>	<b>83.5</b>	<b>90.8</b>	<b>6.09</b>	8.9	<b>72.3</b>	<b>82.9</b>	6.70	38.7	83.8	90.9	<b>6.75</b>
C	$\mathcal{L}$	✓	×	✓	7.5	79.6	87.8	6.88	<b>8.6</b>	70.9	82.0	7.05	<b>37.6</b>	<b>84.4</b>	<b>91.3</b>	7.22
D		✓	✓	×	7.9	81.3	89.3	6.47	8.9	70.2	81.4	<b>6.54</b>	39.2	82.9	90.3	7.08
E		×	—	—	8.1	82.0	89.7	6.45	9.1	69.5	80.9	6.92	39.0	82.9	90.4	7.04
F	$\mathcal{N}$	✓	✓	✓	13.5	77.3	86.8	6.71	21.0	59.6	72.5	6.78	52.0	79.0	87.8	7.22

Table 3: Ablation study of ILS-MI loss with respect to separate region (Sep?), optimizing visible region (OptVis?), and optimizing mask (OptM?).  $\mathcal{N}$  and  $\mathcal{L}$  denote Gaussian and Laplace distribution, respectively. \* denotes the default setting.

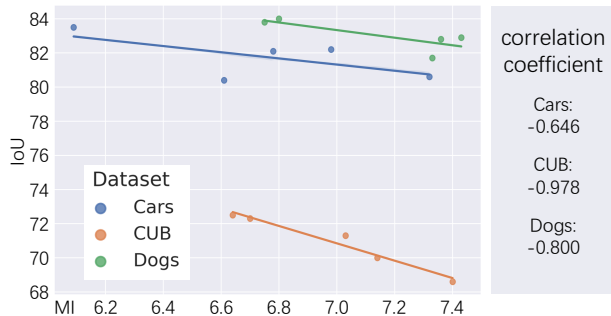


Figure 5: Correlation between IoU and MI. The negative correlation suggests that reducing MI to increase independence contributes to segmentation.

lem. (2) ILS-L1 generally brings less improvement in MI and segmentation performance compared to ILS-MI. These results show that ILS-MI, built upon directly minimizing MI upper bound, achieves more robust and better performance than the intuitive alternative ILS-L1.

**Correlation between MI and IoU** We further plot the MI-IoU data collected from Table. 2 and shows the correlation coefficients between MI and IoU in Fig. 5. The degraded results of which ILS-MI loss weight greater or equal to 2 are not included. It is a clear to see negative correlation between MI and IoU across three datasets, justifying that layer independence can positively contribute to the segmentation performance.

### Ablation Study

We conduct an ablation study about the concrete design choices in our ILS-MI losses. Results are summarized in Table. 3.

**Visible region optimization** We disable the optimization of visible regions by stopping gradients of ILS-MI loss with respect to  $\mathbf{f}_{\text{vis}}$  and  $\mathbf{b}_{\text{vis}}$  in Eq. 11. Results are shown in row C in Table. 3. By comparing row B to row C, it can be seen that optimizing visible regions can decrease MI on all three datasets, has more robust segmentation performance, yet has

no noticeable effect on synthetic fidelity.

**Mask optimization** We ablate the optimization of the mask by stopping the gradients of ILS-MI loss with respect to  $\mathbf{m}$  in Eq. 11. Results are shown in row D in Table. 3. By comparing row B to row D, it is obvious that optimizing the mask can improve segmentation performance on all three datasets and reduce MI on Cars and Dogs, indicating that segmentation is critical to layer independence.

**Separate regions** We ablate the separation of invisible and visible regions in Eq. 11, leading to a form of  $\mathcal{L}_{\text{ILS}} = I(\mathbf{b}; \mathbf{f})$  that computes the MI of the whole foreground and background layers. As shown in row E of Table. 3, separate region, comparing no separation, achieves lower MI and has advantages of both synthetic quality and segmentation performance on all three datasets. This demonstrates the merits of our separation according to mask and visibility.

**Distribution of  $q(\mathbf{b}|\mathbf{f})$**  Besides Laplace distribution, we also try Gaussian distribution to serve as the approximation of  $q(\mathbf{b}|\mathbf{f})$ . As shown in row F of Table. 3, Gaussian distribution completely fails when compared to Laplace and even worse than not using ILS.

## Conclusion

We propose Independent Layer Synthesis GAN, which generates independent foreground and background layers for fully unsupervised foreground-background segmentation. It improves interlayer independence to reduce the semantic and visual confusion caused by information leakage. To enhance independence, we minimize the MI between visible and invisible regions of foreground and background, resulting in more precise masks. Experiments on various single object datasets demonstrate ILSGAN surpasses other unsupervised methods significantly on segmentation performance and synthetic quality. Moreover, the results show that ILSGAN performs more impressively and robustly compared to our intuitive approach, which only reduces visual similarity. We hope our work will motivate future research in unsupervised segmentation, even self-supervised learning (He et al. 2022), *e.g.* foreground masks as attention.

## Acknowledgments

This work was supported by the National Key R&D Program of China under Grant 2018AAA0102801.

## References

- Abdal, R.; Zhu, P.; Mitra, N. J.; and Wonka, P. 2021. Labels4free: Unsupervised segmentation using stylegan. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13970–13979.
- Belghazi, M. I.; Baratin, A.; Rajeswar, S.; Ozair, S.; Bengio, Y.; Courville, A.; and Hjelm, R. D. 2018. Mine: mutual information neural estimation. *International Conference on Machine Learning*.
- Benny, Y.; and Wolf, L. 2020. Onegan: Simultaneous unsupervised learning of conditional image generation, foreground segmentation, and fine-grained clustering. In *European Conference on Computer Vision*, 514–530. Springer.
- Bielski, A.; and Favaro, P. 2019. Emergence of object segmentation in perturbed generative models. *Advances in Neural Information Processing Systems*, 32.
- Burgess, C. P.; Matthey, L.; Watters, N.; Kabra, R.; Higgins, I.; Botvinick, M.; and Lerchner, A. 2019. Monet: Unsupervised scene decomposition and representation. *arXiv preprint arXiv:1901.11390*.
- Chen, M.; Artières, T.; and Denoyer, L. 2019. Unsupervised object segmentation by redrawing. *Advances in Neural Information Processing Systems*, 32.
- Chen, X.; Duan, Y.; Houthoofd, R.; Schulman, J.; Sutskever, I.; and Abbeel, P. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. *Advances in Neural Information Processing Systems*, 29.
- Cheng, P.; Hao, W.; Dai, S.; Liu, J.; Gan, Z.; and Carin, L. 2020. Club: A contrastive log-ratio upper bound of mutual information. In *International Conference on Machine Learning*, 1779–1788. PMLR.
- Cho, J. H.; Mall, U.; Bala, K.; and Hariharan, B. 2021. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16794–16804.
- Choudhury, S.; Laina, I.; Rupperecht, C.; and Vedaldi, A. 2021. Unsupervised part discovery from contrastive reconstruction. *Advances in Neural Information Processing Systems*, 34: 28104–28118.
- Donahue, J.; and Simonyan, K. 2019. Large scale adversarial representation learning. *Advances in Neural Information Processing Systems*, 32.
- Ehrhardt, S.; Groth, O.; Monszpart, A.; Engelcke, M.; Posner, I.; Mitra, N.; and Vedaldi, A. 2020. RELATE: Physically plausible multi-object scene synthesis using structured latent spaces. *Advances in Neural Information Processing Systems*, 33: 11202–11213.
- Engelcke, M.; Kosiorek, A. R.; Jones, O. P.; and Posner, I. 2020. Genesis: Generative scene inference and sampling with object-centric latent representations. *International Conference on Learning Representations*.
- Engelcke, M.; Parker Jones, O.; and Posner, I. 2021. Genesis-v2: Inferring unordered object representations without iterative refinement. *Advances in Neural Information Processing Systems*, 34: 8085–8094.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in Neural Information Processing Systems*, 27.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. 2022. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16000–16009.
- He, K.; Gkioxari, G.; Dollár, P.; and Girshick, R. 2017. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, 2961–2969.
- Heusel, M.; Ramsauer, H.; Unterthiner, T.; Nessler, B.; and Hochreiter, S. 2017. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in Neural Information Processing Systems*, 30.
- Hwang, J.-J.; Yu, S. X.; Shi, J.; Collins, M. D.; Yang, T.-J.; Zhang, X.; and Chen, L.-C. 2019. Segsort: Segmentation by discriminative sorting of segments. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7334–7344.
- Ji, X.; Henriques, J. F.; and Vedaldi, A. 2019. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 9865–9874.
- Karras, T.; Aittala, M.; Hellsten, J.; Laine, S.; Lehtinen, J.; and Aila, T. 2020a. Training generative adversarial networks with limited data. *Advances in Neural Information Processing Systems*, 33: 12104–12114.
- Karras, T.; Laine, S.; and Aila, T. 2019. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4401–4410.
- Karras, T.; Laine, S.; Aittala, M.; Hellsten, J.; Lehtinen, J.; and Aila, T. 2020b. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8110–8119.
- Khosla, A.; Jayadevaprakash, N.; Yao, B.; and Li, F.-F. 2011. Novel dataset for fine-grained image categorization: Stanford dogs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshop*.
- Kim, D.; and Hong, B.-W. 2021. Unsupervised segmentation incorporating shape prior via generative adversarial networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7324–7334.
- Kim, H.; Choi, Y.; Kim, J.; Yoo, S.; and Uh, Y. 2021. Exploiting spatial dimensions of latent in gan for real-time image editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 852–861.
- Krause, J.; Stark, M.; Deng, J.; and Fei-Fei, L. 2013. 3d object representations for fine-grained categorization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshop*.

- Krishna, R.; Zhu, Y.; Groth, O.; Johnson, J.; Hata, K.; Kravitz, J.; Chen, S.; Kalantidis, Y.; Li, L.-J.; Shamma, D. A.; et al. 2017. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123(1): 32–73.
- Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4681–4690.
- Ling, H.; Kreis, K.; Li, D.; Kim, S. W.; Torralba, A.; and Fidler, S. 2021. Editgan: High-precision semantic image editing. *Advances in Neural Information Processing Systems*, 34: 16331–16345.
- Locatello, F.; Weissenborn, D.; Unterthiner, T.; Mahendran, A.; Heigold, G.; Uszkoreit, J.; Dosovitskiy, A.; and Kipf, T. 2020. Object-centric learning with slot attention. *Advances in Neural Information Processing Systems*, 33: 11525–11538.
- Melas-Kyriazi, L.; Rupprecht, C.; Laina, I.; and Vedaldi, A. 2022. Finding an unsupervised image segmenter in each of your deep generative models. *International Conference on Learning Representations*.
- Mescheder, L.; Geiger, A.; and Nowozin, S. 2018. Which training methods for GANs do actually converge? In *International Conference on Machine Learning*.
- Ouali, Y.; Hudelot, C.; and Tami, M. 2020. Autoregressive unsupervised image segmentation. In *European Conference on Computer Vision*, 142–158. Springer.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer Assisted Intervention*.
- Savarese, P.; Kim, S. S.; Maire, M.; Shakhnarovich, G.; and McAllester, D. 2021. Information-theoretic segmentation by inpainting error maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4029–4039.
- Singh, K. K.; Ojha, U.; and Lee, Y. J. 2019. Finegan: Unsupervised hierarchical disentanglement for fine-grained object generation and discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6490–6499.
- Van Gansbeke, W.; Vandenhende, S.; Georgoulis, S.; and Van Gool, L. 2021. Unsupervised semantic segmentation by contrasting object mask proposals. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10052–10062.
- Voynov, A.; Morozov, S.; and Babenko, A. 2021. Object segmentation without labels with large-scale generative models. In *International Conference on Machine Learning*, 10596–10606. PMLR.
- Wah, C.; Branson, S.; Welinder, P.; Perona, P.; and Belongie, S. 2011. The caltech-ucsd birds-200-2011 dataset. *Computation & Neural Systems Technical Report*.
- Yang, J.; Kannan, A.; Batra, D.; and Parikh, D. 2017. Lrgan: Layered recursive generative adversarial networks for image generation. *International Conference on Learning Representations*.
- Yang, J.; Lu, J.; Lee, S.; Batra, D.; and Parikh, D. 2018. Graph r-cnn for scene graph generation. In *European Conference on Computer Vision*, 670–685.
- Yang, Y.; Bilen, H.; Zou, Q.; Cheung, W. Y.; and Ji, X. 2022. Learning Foreground-Background Segmentation from Improved Layered GANs. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2524–2533.
- Yang, Y.; Cheng, X.; Bilen, H.; and Ji, X. 2021. Learning to annotate part segmentation with gradient matching. In *International Conference on Learning Representations*.
- Yu, P.; Xie, S.; Ma, X.; Zhu, Y.; Wu, Y. N.; and Zhu, S.-C. 2021. Unsupervised foreground extraction via deep region competition. *Advances in Neural Information Processing Systems*, 34: 14264–14279.
- Zellers, R.; Yatskar, M.; Thomson, S.; and Choi, Y. 2018. Neural motifs: Scene graph parsing with global context. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5831–5840.
- Zhang, Y.; Ling, H.; Gao, J.; Yin, K.; Lafleche, J.-F.; Barriuso, A.; Torralba, A.; and Fidler, S. 2021. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10145–10155.