

Synthetic Data Supervised Salient Object Detection

Zhenyu Wu

State Key Laboratory of Virtual Reality Technology and Systems, Beihang University
Beijing, China

Lin Wang

School of Transportation Science and Engineering, Beihang University
Beijing, China

Wei Wang

School of Computer Science and Technology, Harbin Institute of Technology
Shenzhen, China

Tengfei Shi

State Key Laboratory of Virtual Reality Technology and Systems, Beihang University
Beijing, China

Chenglizhao Chen*

College of Computer Science and Technology, China University of Petroleum (East China)
Qingdao, China

Aimin Hao

State Key Laboratory of Virtual Reality Technology and Systems, Beihang University, Beijing
Peng Cheng Laboratory, Shenzhen
China

Shuo Li

Department of Medical Imaging, Western University
London, Canada

ABSTRACT

Although deep salient object detection (SOD) has achieved remarkable progress, deep SOD models are extremely data-hungry, requiring large-scale pixel-wise annotations to deliver such promising results. In this paper, we propose a novel yet effective method for SOD, coined SODGAN, which can generate infinite high-quality image-mask pairs requiring only a few labeled data, and these synthesized pairs can replace the human-labeled DUTS-TR to train any off-the-shelf SOD model. Its contribution is three-fold. **1)** Our proposed diffusion embedding network can address the manifold mismatch and is tractable for the latent code generation, better matching with the ImageNet latent space. **2)** For the first time, our proposed few-shot saliency mask generator can synthesize infinite accurate image synchronized saliency masks with a few labeled data. **3)** Our proposed quality-aware discriminator can select highquality synthesized image-mask pairs from noisy synthetic data pool, improving the quality of synthetic data. For the first time, our SODGAN tackles SOD with synthetic data directly generated from the generative model, which opens up a new research paradigm for SOD. Extensive experimental results show that the saliency model trained on synthetic data can achieve 98.4% F-measure of the saliency model trained on the DUTS-TR. Moreover, our approach achieves a new SOTA performance in semi/weakly-supervised methods, and even outperforms several fully-supervised SOTA methods. Code is available at <https://github.com/wuzhenyubuaa/SODGAN>

*Corresponding Author: Chenglizhao Chen, cclz123@163.com

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '22, October 10–14, 2022, Lisboa, Portugal.

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9203-7/22/10...\$15.00

<https://doi.org/10.1145/3503161.3547930>

CCS CONCEPTS

• Computer methodologies → Artificial intelligence; Machine learning.

KEYWORDS

Salient object detection, Synthetic data, Semi-supervised learning

ACM Reference Format:

Zhenyu Wu, Lin Wang, Wei Wang, Tengfei Shi, Chenglizhao Chen, Aimin Hao, and Shuo Li. 2022. Synthetic Data Supervised Salient Object Detection. In *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*, October 10–14, 2022, Lisboa, Portugal. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3503161.3547930>

1 INTRODUCTION

Salient object detection (SOD) aims to segment interesting objects that attract human attention in an image. As a fundamental tool, it can be leveraged to various applications including scene understanding [60], semantic segmentation [59] and image editing [5, 15]. Recently, SOD has achieved significant progress [12, 19, 33, 36, 42, 53, 57] due to the development of deep model. However, deep networks are extremely data-hungry, typically requiring pixel-level humanannotated datasets to achieve high performance (see Fig. 1.a). Labeling large-scale datasets with pixel-level annotations for SOD is very time-consuming, e.g., generally more than five people were asked to annotate the same image to guarantee the label consistency and another ten viewers were asked to cross-check the quality of annotations in the SOC dataset [10].

To alleviate the dependency on pixel-wise annotation, many weakly-supervised SOD methods [17, 37, 49] have been devised. Typically, image-level labels (see Fig. 1.c) are utilized in [17, 37] for saliency localization, and then iteratively finetune their models with predicted saliency maps. Additionally, scribble annotations (see Fig. 1.b) has been proposed recently in [52] to reduce the uncertainty of image-level labels. Although these methods are free of pixel-level annotations, they suffer from various disadvantages, including low




Figure 1: The saliency model trained on synthetic data outperforms SOTA weakly-supervised methods, and is even competitive with fully-supervised models.

prediction accuracy, complex training strategy, dedicated network architecture, and extra data information (e.g., edge) to obtain high-quality saliency maps.

In this paper, we propose a new paradigm SODGAN (see Fig. 1.d) for SOD, which can generate infinite high-quality image-mask pairs with a few labeled data to replace the human-labeled DUTS-TR [37] dataset. Concretely, our SODGAN has three stages: *Stage 1*. Learning a few-shot saliency mask generator to synthesize image-synchronous mask, while utilizing the existing generative adversarial networks (BigGAN [3]) to generate realistic images. *Stage 2*. Selecting high-quality image-mask pairs from the synthetic data pool. *Stage 3*. Training a saliency network on these filtered image-mask pairs. However, there are three main challenges with this approach: **1) Lacking pixel-wise labeled data** as the training dataset to learn a segmentor because BigGAN was trained on the ImageNet that was designed to classification tasks without the pixel-level label. **2) Discovering a meaningful direction** in GAN latent space to disentangle foreground saliency objects from backgrounds is nontrivial, which often requires domain knowledge and laborious engineering. **3) Low-quality image-mask pairs** exist in the synthesized datasets.

To tackle these three challenges, **first**, we present a diffusion embedding network (DEN) (see Sec. 3.2) to utilize the existing well-annotated dataset (i.e., DUTS-TR), which can infer the image’s latent code that match with the ImageNet latent code space; thus, the existing labeled DUTS-TR dataset can provide the pixel-wise label for ImageNet. **Second**, in contrast to the existing works [13, 26, 31] focusing on latent space, we propose a few-shot saliency mask generator to automatically discover meaningful directions in the GANs feature space (see Sec. 3.3), which can synthesize infinite high-quality image synchronized saliency masks with a few labeled data. **Third**, we propose a quality-aware discriminator (see Sec. 3.4) to select high-quality synthesized image-mask pairs from the noisy synthetic data pool, improving the quality of synthetic data.

Our SODGAN has several desirable properties. **a) Fewer labels**. Our approach eliminates large-scale pixel-level supervision requiring only a few labeled data, which reduces the annotation costs. **b) High performance**. We demonstrate that the saliency model trained on synthetic data directly generated from GANs achieves an average 98.4% F-measure of the saliency model trained on the DUTS-TR dataset. Moreover, our SODGAN achieves new SOTA performance in semi/weakly-supervised methods, and even outperforms some fully supervised methods. **c) Generality**. The

synthetic data can be used to train any off-the-shelf SOD model without the need of special architectures, showing strong generalization capabilities on the real test datasets. We summarize the key contributions as follows:

- For the first time, our SODGAN tackles SOD with synthetic data directly generated from the generative model, which opens up a new research paradigm for semi-supervised SOD and significantly reduces the annotation costs.
- Our proposed the DEN can address manifold mismatch and is tractable for the latent code generation, better matching with the ImageNet latent space.
- Our lightweight few-shot saliency mask generator can synthesize infinite accurate image-synchronous saliency masks with a few labeled data.
- Our proposed quality-aware discriminator can select high-quality synthesized image-mask pairs from the noisy synthetic data pool, improving the quality of synthetic data.

2 RELATED WORK

Semi/Weakly-supervised SOD Approaches. With recent advances in semi/weakly-supervised learning, a few existing works exploit the potential of training saliency detectors on image-level [17, 37, 49], region-level [48, 51, 52], and limited pixel-level [41, 44, 50, 58] labeled data to relax the dependency of manually annotated pixel-level saliency masks. For image-level supervision, these approaches [17, 37, 49] follow the same technical route, *i.e.*, producing initial saliency maps with image-level labels and then further refining it via iterative training. Recently, scribble annotation was proposed in [48, 52], but it requires large-scale scribble annotations (10553 images) and extra data information (e.g., edge) to recover integral object structure. **Differences.** Distinct from all these works, our approach provides a new paradigm for semi-supervised SOD. In particular, we introduce SODGAN, a generative model, which can generate infinite high-quality image-mask pairs requiring minimal manual intervention. These generated pairs can then be used for training any existing SOD approaches.

Latent Interpretability of GANs. The previous works have shown that the GANs latent spaces are endowed with human-interpretable semantic arithmetic. A line of recent works [6, 13, 26, 31, 32, 47] employ explicit human-provided supervision to identify interpretable directions in the latent space. For instance, [13, 31] use the classifiers pretrained on the CelebA [21] dataset to produce pseudo labels for the generated images and their latent codes. Another active line of study on GANs [1, 2, 4, 23, 34, 35, 55] targets the object segmentation task. [1] and [4] are based on the idea of decomposing the generative process in a layer-wise fashion. Other works [2, 23, 35] exploit the idea that the object’s location or appearance can be perturbed without affecting image realism. **Differences.** In contrast to existing works manipulating the latent space, our approach is able to discover interpretable directions in the GANs features space, which allows complete control over the diversity of object categories and can automatically find the expected directions.

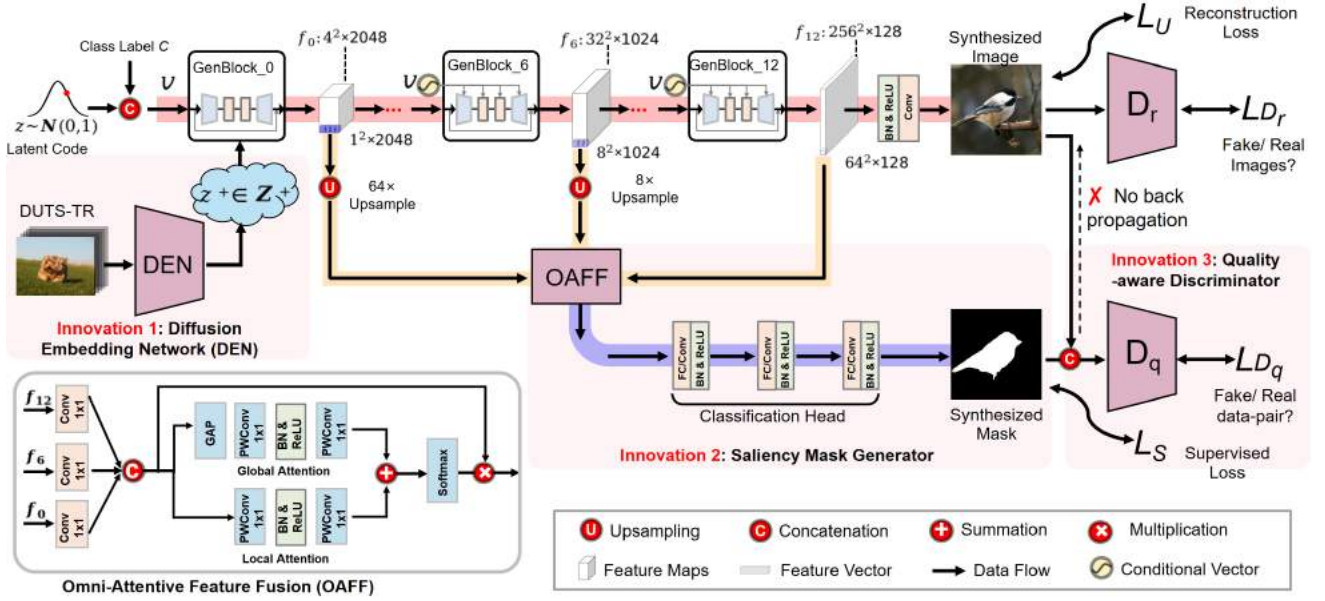


Figure 2: Overview of the proposed SODGAN. Given a latent code $z \sim \mathcal{N}(0, 1)$ and a class label c , we collect hidden semantic feature $\{f_0, f_1, \dots, f_{12}\}$ from $G_{image}(\cdot)$ to disentangle foregrounds from backgrounds. After that, we upsample these collected feature maps to 256×256 resolution and then concatenate these upsampled features together, constructing pixel-wise feature maps for all pixels of the synthesized image. Finally, these pixel-wise features are fed into the proposed saliency mask generator branch to produce image-synchronized saliency mask.

3 METHOD

3.1 Overview

As shown in Fig. 2, our SODGAN is composed of the diffusion embedding network $DEN(\cdot)$, the mask synthesis network $G_{mask}(\cdot)$, the quality-aware discriminator $D_q(\cdot)$, the image synthesis network $G_{image}(\cdot)$, and the image reconstruction discriminator $D_r(\cdot)$. (1) The proposed $DEN(\cdot)$ aims to address the lacking of pixel-wise labels in ImageNet, which is designed for recognition tasks without segmentation groundtruth. Our $DEN(\cdot)$ can utilize existing labeled datasets DUTS-TR [37], and gradually turn it into a unique latent space Z^+ that matches with the ImageNet latent code space. (2) The proposed $G_{mask}(\cdot)$ is to discover meaningful directions in the GANs feature space, synthesizing image synchronized saliency mask. Our $G_{mask}(\cdot)$ is build on top of the $G_{image}(\cdot)$ architecture augmented with a few-shot saliency mask generation branch, (3) Our $D_q(\cdot)$ is designed to select high-quality synthesized image-mask pairs from noisy synthetic data pool. The $G_{image}(\cdot)$ can be any off-the-shelf GANs models, and the $D_r(\cdot)$ is the corresponding real/fake discriminator. Here, we demonstrate our approach using BigGAN [3], a class-conditional GANs trained on ImageNet [8]. In our SODGAN, the proposed $DEN(\cdot)$, $G_{mask}(\cdot)$ and $D_q(\cdot)$ are trainable while the other components remain fixed.

3.2 Diffusion Embedding Network

Our $DEN(\cdot)$ is to address the lacking of pixel-wise label in ImageNet, which is designed for recognition tasks without segmentation groundtruth, better matching with ImageNet latent code space. Previous work [55] addresses this issue by manually labeling a

handful of sampled images, which is labor-consuming. An alternative idea is to utilize the existing labeled datasets (e.g., DUTS-TR) by using variational autoencoder (VAEs). However, the standard VAEs, with a Euclidean latent space, is structurally incapable of capturing topological properties of certain datasets, which is called manifold mismatch [9].

To address these challenges, we developed the diffusion embedding network $DEN(\cdot)$ to utilize the existing labeled datasets with pixel-wise annotation (e.g., DUTS-TR), which allows for an arbitrarily closed manifold as a latent space and captures the underlying geometrical structure. The proposed $DEN(\cdot)$ can gradually turn an image into a unique latent code z^+ that better matches with ImageNet latent code space. Concretely, our $DEN(\cdot)$ are latent variable models of the forms $p_\theta(x_0) = \int p_\theta(x_{0:T}) dx_{1:T}$, where x_1, \dots, x_T are intermediate latent codes and $x_0 \sim q(x_0)$ is the initial image. The joint distribution $p_\theta(x_{0:T})$ is the embedding process, and it is defined as the Markov chain with learned Gaussian transitions $p(x_T) = \mathcal{N}(x_T; 0, 1)$:

$$p_\theta(x_{0:T}) = p(x_T) \prod_{t=1}^T p_\theta(x_{t-1} | x_t), \quad (1)$$

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}\left(x_{t-1}; \mu_\theta(x_t, t), \sum_{\theta} (x_t, t)\right)$$

The difference between our $DEN(\cdot)$ and VAEs is that the approximate posterior $q(x_{1:T}|x_0)$, which is called the diffusion process, is fixed to a Markov chain that progressively adds Gaussian noise to the image in line with variance schedule β_1, \dots, β_T :




Figure 3: Visualizing the omni-attention maps of hidden features, which can locate the salient object masks coarsely.

$$q(x_{1:T} | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}) \quad (2)$$

$$q(x_t | x_{t-1}) = \mathcal{N}(x_t; x_{t-1} \sqrt{1 - \beta_t}, \beta_t)$$

A desirable property of the diffusion process is that it admits sampling x_t at an arbitrary timestep t in closed form:

$$q(x_t | x_0) = \mathcal{N}(x_t; x_0 \sqrt{\alpha_t}, 1 - \alpha_t) \quad (3)$$

where $\alpha = 1 - \beta_t$ and $\hat{\alpha} = \prod_{s=1}^t \alpha_s$. The reconstruction loss is to optimize the variational bound on negative log likelihood:

$$\begin{aligned} \mathcal{L}_U &= \mathbb{E}[-\log p_\theta(x_0)] \leq \mathbb{E}\left[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T} | x_0)}\right] \\ &= \mathbb{E}\left[\sum_{t>1} D_{KL}(q(x_{t-1} | x_t, x_0) \| p_\theta(x_{t-1} | x_t)) - \log_\theta(x_0 | x_1) \right. \\ &\quad \left. + D_{KL}(q(x_T | x_0) \| p(x_T))\right] \end{aligned} \quad (4)$$

where $D_{KL}(\cdot)$ is the KL divergence. The adversarial loss can be defined as:

$$\begin{aligned} \mathcal{L}_{D_r} &= \mathbb{E}_{x \sim \mathcal{D}_{dut}} [\log(D_r(x))] \\ &\quad + \mathbb{E}_{x \sim \mathcal{D}_{dut}} \left[\log \left(1 - D_r \left(G_{image}(E(x)) \right) \right) \right] \end{aligned} \quad (5)$$

where \mathcal{D}_{dut} is the DUTS-TR dataset. Note that our $DEN(\cdot)$ doesn't need special architectures. Here we adopt the MobileNetV3 [14] architecture for the diffusion model. In this way, given an image, our $DEN(\cdot)$ can infer its latent code z^+ that matches with the ImageNet latent code space, and find its groundtruth y in the DUTS-TR.

Summarized Advantages: 1) Gaussian noise has the effect of filling low density regions in the original data distribution; thus, our $DEN(\cdot)$ can obtain more training signal to improve latent distributions that faster converge to the true data distribution. 2) Our $DEN(\cdot)$ is capable of capturing topological properties of certain datasets that better match with ImageNet latent code space.

3.3 Few-shot Saliency Mask Generator

Our $G_{mask}(\cdot)$ is a lightweight few-shot generator trained on a few labeled data, which can synthesize infinite image synchronized accurate saliency masks. The $G_{mask}(\cdot)$ consists of the omni-attentive feature fusion module (OAF), and the classification head, sharing the same feature extractor with the $G_{image}(\cdot)$. Let the $f_i \in \mathbb{R}^{W_i \times H_i \times C_i}$ denote the hidden representation of $G_{image}(\cdot)$.

Omni-Attentive Feature Fusion. Previous work [47] has demonstrated that in GANs feature space, low-level features contain local information like texture and color while high-level features capture global information, such as the style and layout of objects. To fully take advantage of the multi-level features, we proposed a novel omni-attentive feature fusion module, as depicted in Fig. 2 (bottom-left). To ensure the spatial alignment, we first upsample all feature maps $\{f_0, f_1, \dots, f_l\}$ to the highest output resolution 256×256 , and then concatenate them along the channel dimension to obtain an aggregated feature f^* :

$$f^* = \text{Cat} \left(\text{Conv}_{1 \times 1}(U(f_0)), \text{Conv}_{1 \times 1}(U(f_1)), \dots, \text{Conv}_{1 \times 1}(U(f_l)) \right) \quad (6)$$

where $U(\cdot)$ denotes upsample operation, $\text{Conv}_{1 \times 1}(\cdot)$ is 1×1 convolutional operation for reducing channel dimension, and $\text{Cat}(\cdot)$ stands for concatenation. To better fusion the global and local contexts, we introduce the omni-attention $OA(\cdot)$ module, including local attention $LA(\cdot)$ and global attention $GA(\cdot)$:

$$\begin{aligned} OA(f^*) &= LA(f^*) + GA(f^*) \\ LA(f^*) &= PWC(\text{ReLU}(BN(PWC(f^*)))) \\ GA(f^*) &= PWC(\text{ReLU}(BN(PWC(GAP(f^*)))))) \end{aligned} \quad (7)$$

where the $GAP(\cdot)$ is global average pooling, and $PWC(\cdot)$ is the 1×1 point-wise convolution for reducing the parameters. Fig. 3 shows the visualized omni-attention maps. The aggregated feature f' can be obtained by multiplying with the $OA(f^*)$:

$$f' = f^* \otimes OA(f^*) \quad (8)$$

where the \otimes is the element-wise multiplication operator.

Classification Head. After that, the $f' \in \mathbb{R}^{256 \times 256 \times C}$ are flattened into vectors $X \in \mathbb{R}^{65536 \times C}$, constructing pixel-wise feature vectors for all pixels on the synthesized image. Finally, we feed it into the proposed classification head network to produce pixel-wise saliency mask, where the detailed network structure discussed in Sec. 4.1.

To train $G_{mask}(\cdot)$, we need to collect a small training set $\mathcal{D}_m = \{(z_1^+, y_1), \dots, (z_k^+, y_k)\}$, where y_i is selected from the DUTS-TR. Specifically, we use the state-of-the-art (SOTA) image classification model CoAtNet [7] to classify the DUTS-TR, which can be divided into 522 categories. We then randomly select a pair of (z^+, y) for each class, forming a small training set with 522 images. We then train the proposed $G_{mask}(\cdot)$ by using activation features z^+ and the corresponding pixel-wise annotations. The training objective is

$$\mathcal{L}_{G_{mask}} = \mathcal{L}_S + \mathcal{L}_{D_q} \quad (9)$$

\mathcal{L}_S is the supervised loss on labeled images with a combination of cross entropy and dice loss, defined as:

$$\mathcal{L}_S = \frac{1}{HW} \sum_i^H \sum_j^W y_{ij} \log(\hat{y}_{ij}) + \left(1 - \frac{2 \sum_i^H \sum_j^W y_{ij} \hat{y}_{ij}}{\sum_i^H \sum_j^W [y_{ij} + \hat{y}_{ij}]} \right) \quad (10)$$

where the H and W are the height and width of the image respectively, and \hat{y}_{ij} is the prediction probability at position (i, j) . The quality-aware discriminator loss \mathcal{L}_{D_q} is:



Figure 4: With a few labeled data, our SODGAN can generate infinite realistic and diverse (1,000 categories) image-mask pair.

$$\mathcal{L}_{D_q} = \argmin_{G_{\text{mask}}} \mathbb{E}_{(x_{\text{real}}, y_{\text{real}}) \sim \mathcal{D}_m} [\log(D_q(x_{\text{real}}, y_{\text{real}}))] + \mathbb{E}_{z^+ \sim Z^+} [\log(1 - D_q(G_{\text{image}}(E(z^+)), G_{\text{mask}}(E(z^+))))] \quad (11)$$

Summarized Advantages: 1) **Lightweight.** Our $G_{\text{mask}}(\cdot)$ is extremely lightweight yet powerful, which consists of a OAFF and classification head with total 90K parameters and 3.6MB model size. 2) **Fewer labels.** We only need 522 images to train the $G_{\text{mask}}(\cdot)$ because our $G_{\text{mask}}(\cdot)$ is lightweight with only 90K parameters.

3.4 Quality-aware Discriminator

Our $D_q(\cdot)$ can select high-quality synthesized image-mask pairs from noisy synthetic data pool, providing high-quality synthetic data to train the saliency network. We noticed that the synthetic data fails occasionally for non-rigid objects (e.g., dogs) due to their various poses, resulting in low-quality image-mask pairs. To alleviate this issue, we proposed a quality discriminator $D_q(\cdot)$ adopting the lightweight MobileNetV3 [14] as backbone, which aims to select high-quality synthesized image-mask pair. During training, we feed two pairs to the quality discriminator $D_q(\cdot)$, i.e., $(x_{\text{real}}, y_{\text{real}})$ and $(x_{\text{syn}}, y_{\text{syn}})$. Accordingly, the adversarial training loss for the $D_q(\cdot)$ can be formulated as:

$$\mathcal{L}_{D_q} = \argmax_{D_q} \mathbb{E}_{(x_{\text{real}}, y_{\text{real}}) \sim \mathcal{D}_m} [\log(D_q(x_{\text{real}}, y_{\text{real}}))] + \mathbb{E}_{z^+ \sim Z^+} [\log(1 - D_q(G_{\text{image}}(E(z^+)), G_{\text{mask}}(E(z^+))))] \quad (12)$$

Note that our $D_q(\cdot)$ is different from the typical discriminator, where the discriminator is designed for discriminating real or fake images, while our $D_q(\cdot)$ performs image-mask quality control.

Summarized Advantages: 1) **High-quality image-mask pairs.** Our SODGAN can generate any desired number of high-quality image-mask pairs, which forms our synthetic dataset. The generated image-mask pairs can then be used to train any off-the-shelf SOD architecture just like real datasets are. 2) **Strong generalization capabilities.** Unlike previous works [16, 29, 30, 38], which usually arises significant domain gap between the synthetic (from computer games) and real-world domains, the presented SODGAN can generate realistic images (see Fig. 4) and show strong generalization capabilities on the real test datasets (see Table 3).

4 RESULTS AND ANALYSIS

4.1 Classification Head Architecture

In this section, we provide the detailed implementation regarding two aspects: convolutional neural networks and MLP.

CNN Architecture. We first use a linear embedding layer to reduce the input dimension from C to 128, followed by 3 convolutional layers with kernel size of 3. The corresponding dimensions of the output channels are 128, 32, and 2 (the number of classes). All the layers are followed by a leaky ReLU activation function except for the last output layer. We call this standard version CNN-S. We also introduce CNN-M and CNN-L, where M/L denotes medium/large model size, and the architecture hyper-parameters of these model variants can be seen in the first 2 rows of Table 1.

MLP Architecture. We build our base model, called MLP-S, which consists of 3 fully-connected layers with 128, 32, and 2 hidden nodes, respectively. All layers except the output layer are followed by the BatchNorm layer and ReLU activation function. Similar to CNN-S, we also introduce its variants version MLP-M and MLP-L, and their hyper-parameters can be seen in the last 2 rows of Table 1.

	Layers	Channels
CNN-M	5	{128, 64, 64, 32, 2}
CNN-L	7	{128, 64, 64, 64, 64, 32, 2}
MLP-M	4	{128, 64, 32, 2}
MLP-L	5	{128, 64, 64, 32, 2}

Table 1: Architecture details for the adopted CNN/MLP.

4.2 Synthetic Data VS. Real DUTS-TR

As shown in Fig. 5, we provide analyses of our synthesized datasets compared to the real DUTS-TR datasets in terms of center bias, category distribution, color contrast, and salient object size.

Center bias. We visualize the salient object locations for the synthetic data and the DUTS-TR datasets in Fig. 5.a. Most objects are biased towards the image center for both datasets. Compared to the DUTS-TR, the synthetic data show lower center distributions. **Category distribution.** We use the SOTA classification model CoAtNet [7] to classify the filtered synthetic data and the DUTS-TR, which

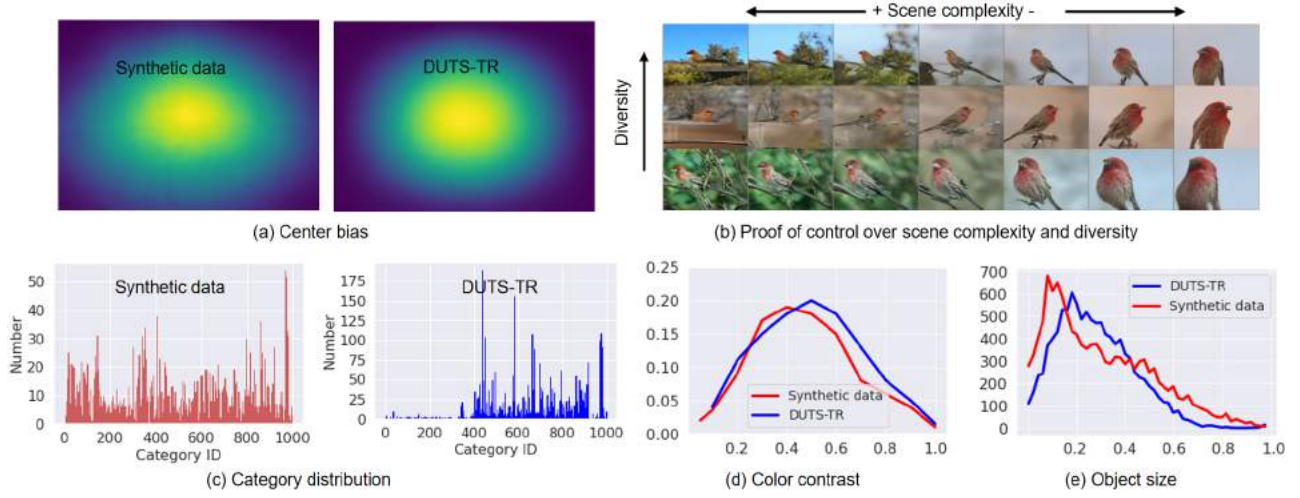


Figure 5: Extensive analysis of synthetic data and the real DUTS-TR dataset show that the synthetic data have many advantages over the DUTS-TR in terms of center bias, category distribution, color contrast, and salient object size. Besides, the synthetic data can control the samples complexity and diversity.

	DUTS-TE			ECSSD		
	$\max F_\beta \uparrow$	S-m \uparrow	MAE \downarrow	$\max F_\beta \uparrow$	S-m \uparrow	MAE \downarrow
VAEs	.8375	.8331	.0644	.9241	.8945	.0466
$DEN(\cdot)$.8557	.8507	.0530	.9377	.9129	.0389
w/o OAFF	.8333	.8320	.0622	.9187	.8945	.0476
w/ $GA(\cdot)$.8437	.8413	.0594	.9285	.9024	.0429
w/ OAFF	.8557	.8507	.0530	.9377	.9129	.0389
CNN-S	.8416	.8400	.0562	.9229	.8995	.0464
CNN-M	.8424	.8478	.0479	.9274	.9025	.0432
CNN-L	.8208	.8098	.0787	.9206	.8905	.0529
MLP-S	.8557	.8507	.0530	.9377	.9129	.0389
MLP-M	.8382	.8327	.0625	.9349	.9118	.0399
MLP-L	.8510	.8476	.0551	.9375	.9114	.0394
w/o D_q	.8406	.8370	.0603	.9306	.9038	.0437
w/ D_q	.8557	.8507	.0530	.9377	.9129	.0389

Table 2: Comparisons of different network structures for mask generator on DUTS-TE and ECSSD datasets.

can be divided into 764 and 522 categories, respectively. As shown in Fig. 5.c, our synthetic data contains more object categories than the DUTS-TR. **Color contrast & Object size.** Since the DUTS-TR was designed for SOD tasks, the DUTS-TR’s images containing at least one salient object are higher color contrast than randomly generated synthetic data (see Fig. 5.d). Besides, we also statistics the object size of the DUTS-TR and our synthetic data in Fig. 5.e. As we can see, the synthetic data also contains smaller objects than the DUTS-TR. Additionally, BigGAN introduced the “truncation coefficient” λ , allowing explicit, fine-grained control of the trade-off between sample variety and complexity (see Fig. 5.b).

4.3 Ablation Study of Our Innovations

Effects of the proposed $DEN(\cdot)$. To demonstrate the effects of our $DEN(\cdot)$, we compared the proposed $DEN(\cdot)$ with commonly used VAEs. As shown in Table 2, the proposed $DEN(\cdot)$ improved by 1.8% compared to the VAEs in terms of S-measure, which shows the effectiveness of the proposed diffusion model.




Figure 6: Left: Results on varying amount of synthesized data. Right: The effects of different truncation λ .

Effects of the proposed OAFF. In Table 2, we evaluate 3 settings of OAFF: 1) $G_{mask}(\cdot)$ without using the OAFF; 2) $G_{mask}(\cdot)$ only using the global attention $GA(\cdot)$; 3) the $G_{mask}(\cdot)$ with the OAFF. As we can see, the $G_{mask}(\cdot)$ with the $GA(\cdot)$ achieves better performance than the plain version, and the performance can be further improved by using OAFF, demonstrating the contribution of the OAFF to the final results.

The choice of classification head architecture. We evaluate 2 architectures on the proposed classification head network, i.e., CNN and MLP, with small (S), medium (M), and large (L) networks described in Sec. 4.1. As shown in Table 2, we notice that the MLP-S outperforms all the three CNN networks. Besides, we also notice that smaller networks obtain better performance due to the limited training data. Therefore, we take the MLP-S with channel dimension {128, 32, 2} as our classification head.

Effects of the proposed $D_q(\cdot)$. To illustrate the effectiveness of the proposed $D_q(\cdot)$, we implement 2 different settings, i.e., our SODGAN with/without using the $D_q(\cdot)$. As shown in Table 2, the performance can be improved by 1.5% in terms of F-measure on the DUTS-TE dataset by using the $D_q(\cdot)$, verifying the contribution of our $D_q(\cdot)$ to the final results.

Impacts of the amount of synthesized data. We further explore the number of synthesized data how to influence the saliency performance. As shown in left of Fig. 6, when the number of synthesized images is insufficient ($< 12k$), model performance can benefit

	Metric	Fully-Supervised Models								Semi/Weakly-Supervised Models							Ours
		DGRL [39]	PAGR [54]	BAS [27]	CPD [43]	MINet [24]	F3Net [40]	SAMN [20]	PFSN [22]	MWS [49]	ENDS [51]	WS ³ A [52]	SCWS [48]	FCS [50]	MFNet [25]		
DUTS-OM.	maxF \uparrow	.7742	.7709	.8053	.7966	.8098	.8133	.8026	.8233	.7176	.7581	.7532	.7827	.7170	.7062	.7930	
	S-m \uparrow	.8059	.7751	.8362	.8248	.8329	.8385	.8299	.8425	.7559	.7832	.7848	.8019	.7448	.7418	.8022	
	MAE \downarrow	.0618	.0709	.0565	.0560	.0555	.0526	.0652	.0545	.1086	.0759	.0684	.0602	.0656	.0867	.0768	
	AUC \uparrow	.8821	.8983	.9262	.9378	.9396	.9413	.9573	.9496	.9413	.9506	.9182	.8822	.8381	.9090	.9565	
	avgF \uparrow	.7656	.7354	.7875	.7770	.7907	.7957	.7655	.8069	.6777	.7246	.7386	.7602	.7073	.6816	.7689	
DUTS-TE	maxF \uparrow	.8287	.8545	.8591	.8654	.8835	.8905	.8360	.8949	.7686	.8173	.7889	.8448	.8296	.7707	.8557	
	S-m \uparrow	.8410	.8369	.8649	.8684	.8834	.8881	.8479	.8916	.7573	.8190	.8021	.8391	.8206	.7728	.8507	
	MAE \downarrow	.0500	.0562	.0480	.0438	.0375	.0358	.0582	.0359	.0920	.0657	.0628	.0493	.0459	.0772	.0530	
	AUC \uparrow	.9137	.9540	.9451	.9627	.9714	.9726	.9708	.9739	.9539	.9645	.9312	.8967	.9000	.9343	.9720	
	avgF \uparrow	.8209	.8108	.8261	.8357	.8566	.8647	.7920	.8714	.7311	.7743	.7715	.8326	.8085	.7415	.8374	
ECSSD	maxF \uparrow	.9224	.9268	.9424	.9392	.9475	.9453	.9279	.9523	.8778	.9002	.8880	.9145	.9108	.8796	.9377	
	S-m \uparrow	.9028	.8892	.9162	.9181	.9249	.9242	.9071	.9298	.8275	.8707	.8655	.8818	.8787	.8345	.9129	
	MAE \downarrow	.0407	.0609	.0370	.0371	.0334	.0333	.0501	.0309	.0963	.0676	.0590	.0489	.0471	.0843	.0389	
	AUC \uparrow	.9505	.9685	.9666	.9812	.9845	.9846	.9857	.9860	.9771	.9776	.9531	.9268	.9478	.9497	.9868	
	avgF \uparrow	.9122	.8944	.8970	.9216	.9295	.9272	.8985	.9346	.8430	.8730	.8733	.9003	.8951	.8490	.9137	
HKU-IS	maxF \uparrow	.9105	.9176	.9285	.9251	.9351	.9368	.9147	.9428	.8560	.9041	.8805	.9085	.8992	.8766	.9320	
	S-m \uparrow	.8945	.8873	.9090	.9055	.9190	.9173	.8983	.9244	.8182	.8838	.8649	.8820	.8718	.8465	.9092	
	MAE \downarrow	.0356	.0475	.0322	.0342	.0285	.0280	.0449	.0259	.0843	.0461	.0470	.0375	.0389	.0585	.0324	
	AUC \uparrow	.9475	.9704	.9650	.9765	.9833	.9817	.9852	.9834	.9774	.9826	.9564	.9282	.9401	.9671	.9861	
	avgF \uparrow	.8968	.8904	.9046	.9004	.9172	.9177	.8856	.9256	.8291	.8801	.8677	.8945	.8836	.8535	.9054	
PASCAL-S	maxF \uparrow	.8808	.8691	.8757	.8841	.8894	.8948	.8568	.8986	.8140	.8706	.8374	.8660	.8742	.8202	.8924	
	S-m \uparrow	.8278	.7925	.8194	.8277	.8333	.8404	.8027	.8431	.7532	.8025	.7805	.7936	.8102	.7489	.8422	
	MAE \downarrow	.0823	.1149	.0924	.0890	.0828	.0799	.1130	.0790	.1509	.1144	.1106	.1000	.0849	.1379	.0743	
	AUC \uparrow	.8988	.9162	.9113	.9316	.9339	.9428	.9348	.9438	.9494	.9588	.9062	.8643	.9071	.9032	.9497	
	avgF \uparrow	.8528	.8148	.8100	.8439	.8512	.8580	.8054	.8614	.7566	.8222	.8054	.8321	.8387	.7739	.8542	

Table 3: Extensive experiments demonstrate that our approach achieves a new SOTA performance in terms of semi/weakly supervised methods, and even outperforms several fully-supervised SOTA methods. The top 2 results are highlighted in red and blue respectively. The “DUTS-OM.” denotes DUT-OMRON dataset. The detailed training data setting can be found in Table 4

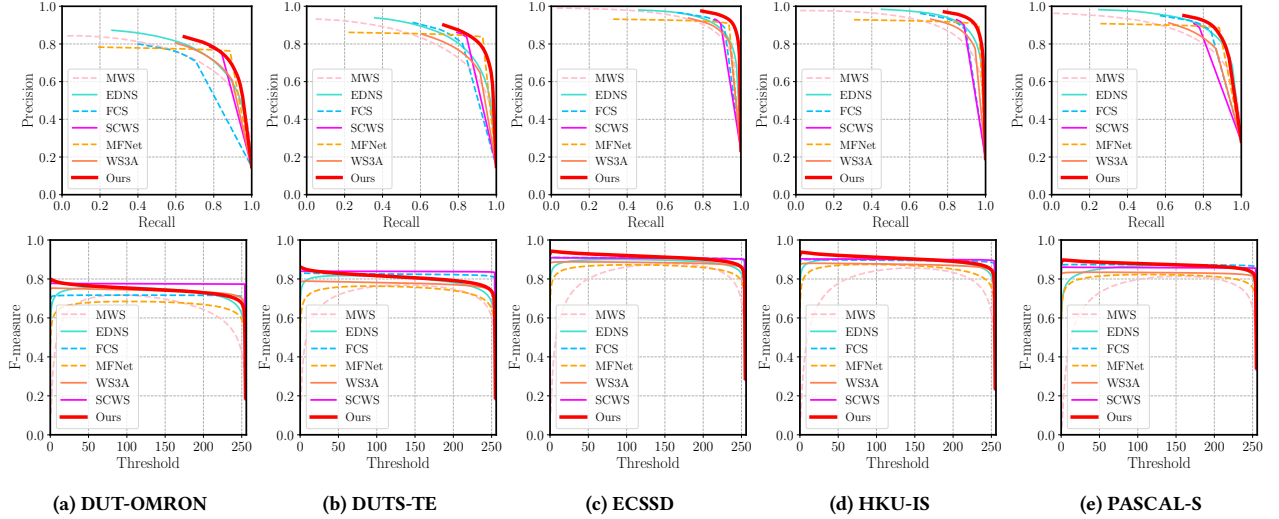


Figure 7: The performance on PR and F-measure curves also show the superiority of our method comparison with 6 SOTA weakly/semi-supervised methods over 5 datasets.

substantially from the increased synthesized data. However, when the training set is large enough ($> 12k$), the application of more synthesized data does not necessarily lead to better performance. In this paper, unless otherwise specified, the reported SOD results were obtained by training on 12k synthetic image-mask pairs. Besides, to study the effects of λ , we vary the truncation coefficient $\lambda \in \{0.2, 0.4, 0.6, 0.8, 1\}$. The results are shown in the right of Fig. 6. We observed that the saliency performance is inversely proportional to λ when $\lambda > 0.4$, and the optimal setting is $\lambda = 0.4$.

4.4 Synthetic Data for SOD

Setup. In this work, we do not focus on SOD network architecture design, so in our experiments, we adopt F3Net [40] as our saliency network by considering effectiveness and computational cost. Different from the previous works trained on the human wellannotated DUTS-TR [37] dataset (the detailed training data setting can be found in Table 4), we train our model on the SODGAN’s generated images-mask pairs (12k).

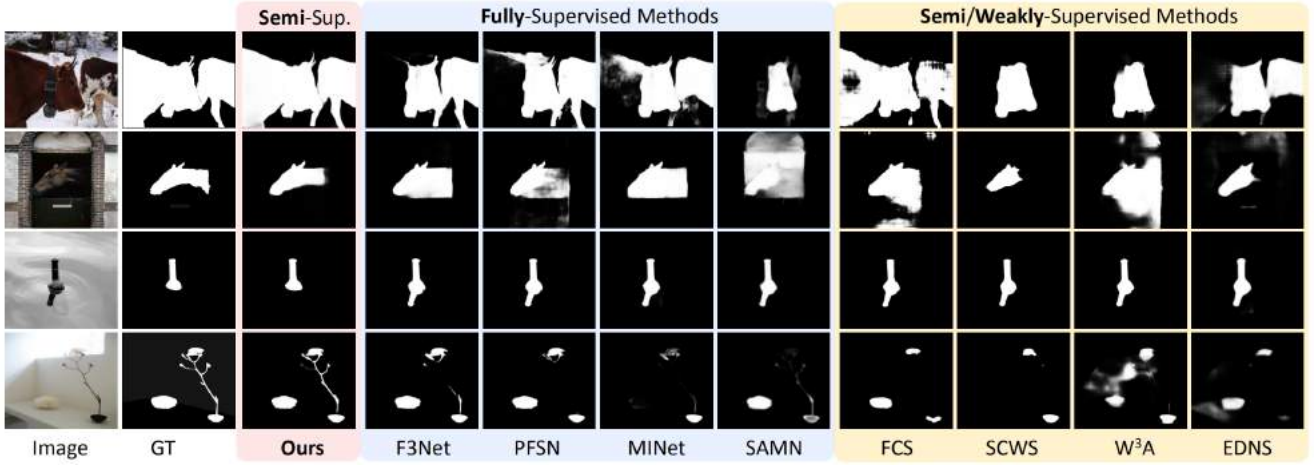


Figure 8: The visual comparison of the proposed model and existing SOTA methods also show that our model can generate more complete and accurate saliency maps than other semi/weakly-supervised, even outperforms fully-supervised SOD models.

Method	Sup.	Training dataset	Annotations	Number
All-S	F	DUTS-TR	Pixel-wise	10,553
MWS [49]	W	ImageNet+COCO	Image-level	1.3M
MFNet [25]	W	ImageNet+DUTS-TR	Image-level	1.01M
EDNS [51]	W	DUTS-TR	Pseudo	10,553
WS ³ A [52]	W	DUTS-TR	Scribble	10,553
SCWS [48]	W	DUTS-TR	Scribble	10,553
FCS [50]	S	DUTS-TR	Pixel-wise	1,000
Ours	S	DUTS-TR	Pixel-wise	522

Table 4: Statistics of popular SOD training dataset. “ALL-S” denotes all supervised models in Table 3. “Sup.” stands for level of supervision. “F, W and S” denote the fully-, weakly- and semi- supervised learning respectively.

Datasets. We evaluate the performance of the proposed method on 5 commonly used benchmark datasets, including DUTS-TE [37], DUT-OMRON [46], ECSSD [45], HKU-IS [56], and PASCAL-S [18]. **Evaluation metrics.** We adopt several widely-used metrics to evaluate our method, including the Precision-Recall (PR) curves, the F-measure curves, Mean Absolute Error (MAE), max and mean F-measure [28], S-measure [11] and Area Under Curve (AUC).

Competitors. We compare the proposed approach with 13 SOTA SOD models, including MWS [49], EDNS [51], WS³A [52], SCWS [48], FCS [50], MFNet [25], DGRL [39], PAGR [54], BAS [27], CPD [43], MINet [24], F3Net [40], PFSN [22], and SAMN [20]. For fair comparison, we evaluate these SOTA models by using the same metric code with the authors provided saliency maps.

Quantitative comparison. In Table 3, we compare our results with SOTA saliency methods. As indicated in Table 3, our method consistently achieves significant improvement compared with semi- and weakly- supervised methods in terms of 5 evaluation metrics. Concretely, our method improved by 1.13%, 1.09%, 2.32%, 2.35%, and 1.82% on average compared to the second-best method in max F-measure on 5 datasets. Moreover, our saliency model even outperforms fully-supervised saliency models, such as CPD [43], BAS [27] and SAMN [20], on ECSSD, HKU-IS and PASCAL-S datasets. Our approach trained on synthetic data achieves comparable or superior to the fully supervised F3Net (0.8422 vs. 0.8404 in terms

of S-measure on the PASCAL-S) trained on more than 10k well-annotated image-label pairs. Besides, we also provide the PR and F-measure curves in Fig. 7, which also demonstrate the effectiveness of the synthesized high-quality image-mask pairs for saliency detection.

Qualitative comparison. As demonstrated in Fig. 8, our synthetic data supervised saliency model has better visual superiority than other SOTA models. Concretely, our model excels in dealing with various challenging scenarios, including cluttered backgrounds (the 1st row), low contrast objects (the 2nd row), inverted reflection in the water (the 3rd row), and small objects (the 4th row).

4.5 Conclusion

In this paper, we present a simple but powerful approach, namely SODGAN, to explore the potential of synthetic data for SOD. It opens up a new research paradigm for semi-supervised SOD, and shows that promising segmentation accuracy can be achieved by using controllable synthesized data. Our major novelty is to discover the interpretable direction that can disentangle the foreground object from the background in GANs feature space with only a few annotated images. Our work expands the application of the generative model to salient object detection tasks. We believe this is only the first step by utilizing synthetic data to train saliency deep networks. In the future, we plan to extend SODGAN to handle other tasks such as medical image segmentation, semantic segmentation, and object detection.

ACKNOWLEDGMENTS

This research is supported in part by the National Natural Science Foundation of China (No. 62172437 and 62172246), the Open Project Program of State Key Laboratory of Virtual Reality Technology and Systems (VRLAB2021A05), the Youth Innovation and Technology Support Plan of Colleges and Universities in Shandong Province (2021KJ062), and the Science and Technology Innovation Committee of Shenzhen Municipality (No. JCYJ20210324131800002 and RCBS20210609103820029).

REFERENCES

- [1] Rameen Abdal, Peihao Zhu, Niloy Mitra, and Peter Wonka. 2021. Labels4Free: Unsupervised Segmentation using StyleGAN. In *ICCV*.
- [2] Adam Jakub Bielski and Paolo Favaro. 2019. Emergence of Object Segmentation in Perturbed Generative Models. In *NeurIPS*, Vol. 32.
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. 2018. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *ICLR*.
- [4] Mickaël Chen, Thierry Artières, and Ludovic Denoyer. 2019. Unsupervised Object Segmentation by Redrawing. In *NeurIPS*. 12705–12716.
- [5] Ming-Ming Cheng, Fang-Lue Zhang, Niloy J Mitra, Xiaolei Huang, and Shi-Min Hu. 2010. Repfinder: finding approximately repeated scene elements for image editing. *ACM TOG* 29, 4 (2010), 1–8.
- [6] Anton Cherepkov, Andrey Voynov, and Artem Babenko. 2021. Navigating the gan parameter space for semantic image editing. In *CVPR*. 3671–3680.
- [7] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. 2021. Coatnet: Marrying convolution and attention for all data sizes. *NeurIPS* 34 (2021), 3965–3977.
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*. 248–255.
- [9] Luca Falorsi, Pim de Haan, Tim R Davidson, Nicola De Cao, Maurice Weiler, Patrick Forré, and Taco S Cohen. 2018. Explorations in homeomorphic variational auto-encoding. *arXiv preprint arXiv:1807.04689* (2018).
- [10] Deng-Ping Fan, Ming-Ming Cheng, Jiang-Jiang Liu, Shang-Hua Gao, Qibin Hou, and Ali Borji. 2018. Salient objects in clutter: Bringing salient object detection to the foreground. In *ECCV*. 186–202.
- [11] Deng-Ping Fan, Ming-Ming Cheng, Yun Liu, Tao Li, and Ali Borji. 2017. Structure-measure: A new way to evaluate foreground maps.. In *CVPR*. 4548–4557.
- [12] Hao Fang, Daoxin Zhang, Yi Zhang, Minghao Chen, Jiawei Li, Yao Hu, Deng Cai, and Xiaofei He. 2021. Salient Object Ranking With Position-Preserved Attention. In *ICCV*. 16331–16341.
- [13] Lore Goetschalckx, Alex Andonian, Aude Oliva, and Phillip Isola. 2019. Ganalyze: Toward visual definitions of cognitive image properties. In *ICCV*. 5744–5753.
- [14] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. 2019. Searching for mobilenetv3. In *ICCV*. 1314–1324.
- [15] Lai Jiang, Mai Xu, Xiaofei Wang, and Leonid Sigal. 2021. Saliency-Guided Image Translation. In *CVPR*. 16509–16518.
- [16] Amlan Kar, Aayush Prakash, Ming-Yu Liu, Eric Cameracci, Justin Yuan, Matt Rusiniak, David Acuna, Antonio Torralba, and Sanja Fidler. 2019. Meta-sim: Learning to generate synthetic datasets. In *ICCV*. 4551–4560.
- [17] Guanbin Li, Yuan Xie, and Liang Lin. 2018. Weakly supervised salient object detection using image labels. In *AAAI*.
- [18] Yin Li, Xiaodi Hou, Christof Koch, James M Rehg, and Alan L Yuille. 2014. The secrets of salient object segmentation. In *CVPR*. 280–287.
- [19] Nian Liu, Wangbo Zhao, Dingwen Zhang, Junwei Han, and Ling Shao. 2021. Light Field Saliency Detection With Dual Local Graph Learning and Reciprocal Guidance. In *ICCV*. 4712–4721.
- [20] Yun Liu, Xin-Yu Zhang, Jia-Wang Bian, Le Zhang, and Ming-Ming Cheng. 2021. SAMNet: Stereoscopically Attentive Multi-Scale Network for Lightweight Salient Object Detection. *IEEE TIP* 30 (2021), 3804–3814.
- [21] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. 2015. Deep learning face attributes in the wild. In *ICCV*. 3730–3738.
- [22] Mingcan Ma, Changqun Xia, and Jia Li. 2021. Pyramidal Feature Shrinking for Salient Object Detection. In *AAAI*.
- [23] Luke Melas-Kyriazi, Christian Rupprecht, Iro Laina, and Andrea Vedaldi. 2021. Finding an Unsupervised Image Segmenter in Each of Your Deep Generative Models. *arXiv preprint arXiv:2105.08127* (2021).
- [24] Youwei Pang, Xiaoqi Zhao, Lihe Zhang, and Huchuan Lu. 2020. Multi-Scale Interactive Network for Salient Object Detection. In *CVPR*. 9413–9422.
- [25] Yongri Piao, Jian Wang, Miao Zhang, and Huchuan Lu. 2021. MFNet: Multi-filter Directive Network for Weakly Supervised Salient Object Detection. In *ICCV*. 4136–4145.
- [26] Antoine Plummerault, Hervé Le Borgne, and Céline Hudelot. 2019. Controlling generative models with continuous factors of variations. In *ICLR*.
- [27] Xuebin Qin, Zichen Zhang, Chenyang Huang, Chao Gao, Masood Dehghan, and Martin Jagersand. 2019. BASNet: Boundary-Aware Salient Object Detection. In *CVPR*. 7479–7489.
- [28] Margolin Ran, Zelnik-Manor Lihi, and Tal Ayallet. 2014. How to evaluate foreground maps?. In *CVPR*. 248–255.
- [29] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. 2016. Playing for data: Ground truth from computer games. In *ECCV*. 102–118.
- [30] German Ros, Laura Sellart, Joanna Materzynska, David Vazquez, and Antonio M Lopez. 2016. The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In *CVPR*. 3234–3243.
- [31] Yujun Shen, Jinjin Gu, Xiaoou Tang, and Bolei Zhou. 2020. Interpreting the latent space of gans for semantic face editing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 9243–9252.
- [32] Yujun Shen and Bolei Zhou. 2021. Closed-form factorization of latent semantics in gans. In *CVPR*. 1532–1540.
- [33] Lv Tang, Bo Li, Yijie Zhong, Shouhong Ding, and Mofei Song. 2021. Disentangled High Quality Salient Object Detection. In *ICCV*. 3580–3590.
- [34] Nontawat Tritrong, Pitchaporn Rewatbowornwong, and Supasorn Suwajanakorn. 2021. Repurposing gans for one-shot semantic part segmentation. In *CVPR*. 4475–4485.
- [35] Andrey Voynov, Stanislav Morozov, and Artem Babenko. 2021. Object Segmentation Without Labels with Large-Scale Generative Models. In *ICML*. 10596–10606.
- [36] Guotao Wang, Chenglizhao Chen, Deng-Ping Fan, Aimin Hao, and Hong Qin. 2021. From Semantic Categories to Fixations: A Novel Weakly-Supervised Visual-Auditory Saliency Detection Approach. In *CVPR*. 15119–15128.
- [37] Lijun Wang, Huchuan Lu, Yifan Wang, Mengyang Feng, Dong Wang, Baocai Yin, and Xiang Ruan. 2017. Learning to detect salient objects with image-level supervision. In *CVPR*. 136–145.
- [38] Qi Wang, Junyu Gao, Wei Lin, and Yuan Yuan. 2019. Learning from synthetic data for crowd counting in the wild. In *CVPR*. 8198–8207.
- [39] Tiantian Wang, Lihe Zhang, Shuo Wang, Huchuan Lu, Gang Yang, Xiang Ruan, and Ali Borji. 2018. Detect globally, refine locally: A novel approach to saliency detection. In *CVPR*. 3127–3135.
- [40] Jun Wei, Shuhui Wang, and Qingming Huang. 2020. F³Net: Fusion, Feedback and Focus for Salient Object Detection. In *AAAI*. 12321–12328.
- [41] Zhenyu Wu, Shuai Li, Chenglizhao Chen, Aimin Hao, and Hong Qin. 2020. A deeper look at image salient object detection: Bi-stream network with a small training dataset. *IEEE TMM* (2020).
- [42] Zhenyu Wu, Shuai Li, Chenglizhao Chen, Aimin Hao, and Hong Qin. 2022. Recursive multi-model complementary deep fusion for robust salient object detection via parallel sub-networks. *Pattern Recognition* 121 (2022), 108212.
- [43] Zhe Wu, Li Su, and Qingming Huang. 2019. Cascaded Partial Decoder for Fast and Accurate Salient Object Detection. In *CVPR*. 3907–3916.
- [44] Pengxiang Yan, Guanbin Li, Yuan Xie, Zhen Li, Chuan Wang, Tianshui Chen, and Liang Lin. 2019. Semi-supervised video salient object detection using pseudo-labels. In *ICCV*. 7284–7293.
- [45] Qiong Yan, Li Xu, Jianping Shi, and Jiaya Jia. 2013. Hierarchical saliency detection. In *CVPR*. 1155–1162.
- [46] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. 2013. Saliency detection via graph-based manifold ranking. In *CVPR*. 3166–3173.
- [47] Huiting Yang, Liangyu Chai, Qiang Wen, Shuang Zhao, Zixun Sun, and Shengfeng He. 2021. Discovering Interpretable Latent Space Directions of GANs Beyond Binary Attributes. In *CVPR*. 12177–12185.
- [48] Siyue Yu, Bingfeng Zhang, Jimin Xiao, and Eng Gee Lim. 2021. Structure-consistent weakly supervised salient object detection with local saliency coherence. In *AAAI*.
- [49] Yu Zeng, Yunzhi Zhuge, Huchuan Lu, Lihe Zhang, Mingyang Qian, and Yizhou Yu. 2019. Multi-source weak supervision for saliency detection. In *CVPR*. 6074–6083.
- [50] Dingwen Zhang, Haibin Tian, and Jungong Han. 2020. Few-cost salient object detection with adversarial-paced learning. In *NeurIPS*.
- [51] Jing Zhang, Jianwen Xie, and Nick Barnes. 2020. Learning noise-aware encoder-decoder from noisy labels by alternating back-propagation for saliency detection. In *ECCV*. 349–366.
- [52] Jing Zhang, Xin Yu, Aixuan Li, Peipei Song, Bowen Liu, and Yuchao Dai. 2020. Weakly-supervised salient object detection via scribble annotations. In *CVPR*. 12546–12555.
- [53] Miao Zhang, Tingwei Liu, Yongri Piao, Shunyu Yao, and Huchuan Lu. 2021. Auto-msfnet: Search multi-scale fusion network for salient object detection. In *ACM Multimedia*. 667–676.
- [54] Xiaoning Zhang, Tiantian Wang, Jinqing Qi, Huchuan Lu, and Gang Wang. 2018. Progressive attention guided recurrent network for salient object detection. In *CVPR*. 714–722.
- [55] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. 2021. Datasetgan: Efficient labeled data factory with minimal human effort. In *CVPR*. 10145–10155.
- [56] Rui Zhao, Wanli Ouyang, Hongsheng Li, and Xiaogang Wang. 2015. Saliency detection by multi-context deep learning. In *CVPR*. 1265–1274.
- [57] Zhirui Zhao, Changqun Xia, Chenxi Xie, and Jia Li. 2021. Complementary Trilateral Decoder for Fast and Accurate Salient Object Detection. In *ACM Multimedia*. 4967–4975.
- [58] Yuan Zhou, Shuwei Huo, Wei Xiang, Chunping Hou, and Sun-Yuan Kung. 2018. Semi-supervised salient object detection using a linear feedback control system model. *IEEE TCYB* 49, 4 (2018), 1173–1185.
- [59] Yanzhao Zhou, Xin Wang, Jianbin Jiao, Trevor Darrell, and Fisher Yu. 2020. Learning Saliency Propagation for Semi-Supervised Instance Segmentation. In *CVPR*.
- [60] Zikun Zhou, Wenjie Pei, Xin Li, Hongpeng Wang, Feng Zheng, and Zhenyu He. 2021. Saliency-Associated Object Tracking. In *ICCV*. 9866–9875.