



Modality-Guided Subnetwork for Salient Object Detection

Zongwei Wu, Guillaume Allibert, Christophe Stolz, Chao Ma, Cédric Demonceaux

► To cite this version:

Zongwei Wu, Guillaume Allibert, Christophe Stolz, Chao Ma, Cédric Demonceaux. Modality-Guided Subnetwork for Salient Object Detection. 9th International Conference on 3D Vision, Dec 2021, Online, France. hal-03375042v2

HAL Id: hal-03375042

<https://hal.science/hal-03375042v2>

Submitted on 25 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Modality-Guided Subnetwork for Salient Object Detection

Zongwei Wu^{1,3} Guillaume Allibert² Christophe Stolz¹ Chao Ma³ Cédric Demonceaux⁴*

¹ VIBOT EMR CNRS 6000, ImViA, Université Bourgogne Franche-Comté, France

² Université Côte d’Azur, CNRS, I3S, France

³ MoE Key Lab of Artificial Intelligence, AI Institute, Shanghai Jiao Tong University, China

⁴ ImViA, Université Bourgogne Franche-Comté, France

zongwei.wu@etu.u-bourgogne.fr, allibert@i3s.unice.fr, chaoma@sjtu.edu.cn

{christophe.stolz, cedric.demonceaux}@u-bourgogne.fr

Abstract

Recent RGBD-based models for saliency detection have attracted research attention. The depth clues such as boundary clues, surface normal, shape attribute, etc., contribute to the identification of salient objects with complicated scenarios. However, most RGBD networks require multi-modalities from the input side and feed them separately through a two-stream design, which inevitably results in extra costs on depth sensors and computation. To tackle these inconveniences, we present in this paper a novel fusion design named modality-guided subnetwork (MGSnet). It has the following superior designs: 1) Our model works for both RGB and RGBD data, and dynamically estimates depth if not available. Taking the inner workings of depth-prediction networks into account, we propose to estimate the pseudo-geometry maps from RGB input — essentially mimicking the multi-modality input. 2) Our MGSnet for RGB SOD results in real-time inference but achieves state-of-the-art performance compared to other RGB models. 3) The flexible and lightweight design of MGS facilitates the integration into RGBD two-streaming models. The introduced fusion design enables a cross-modality interaction to enable further progress but with a minimal cost.

1. Introduction

In the last decade, RGB-based deep learning models for salient object detection (SOD) [33, 5, 19, 32, 35] achieved significant success thanks to the advances of GPU and CNN. Given an input image, the goal of SOD is to compute the pixel-wise location of the prominent objects that visually attract human attention the most. However, RGB SOD models focus more on photometric information instead of

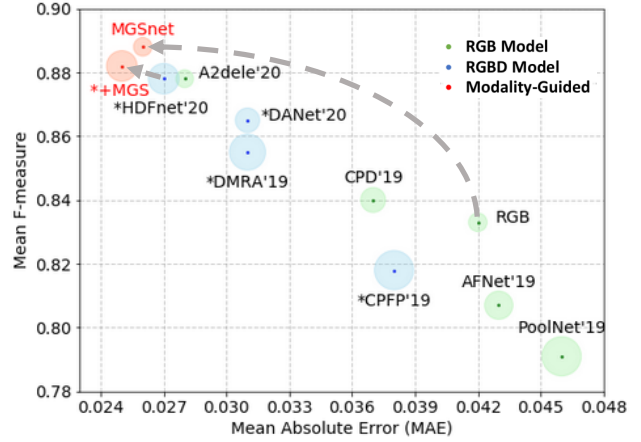


Figure 1: Performance analysis on NLPR dataset [24]. Note that better models are shown in the upper left corner (i.e., with a larger mean F-measure and smaller MAE). The circle size denotes the model size. Our proposed MGSnet for RGB SOD achieves the best performance with the lightest model size. The MGS design can also be embedded to the state-of-the-art RGBD model HDFnet [23] to enable further progress (denoted as * + MGS).

geometry. This is due to the fixed shape and size kernel design of CNN that is not invariant to scale changes and to 3D rotations. By the lack of geometric information on the input side, it is inevitable for RGB models to add additional learning modules in the network to attend to salient objects, resulting in model complexity and computational cost.

Recent RGBD-based SOD has motivated research interest thanks to the accessibility of cross-modal information from the input side. State-of-the-art RGBD models [9, 26, 23, 36] achieve superior performance over the RGB baseline, affirming the effectiveness of learning from two modalities. Most architectures adapt fusion-wise models, such as early fusion [36] where the depth map is fed as

* Corresponding author.

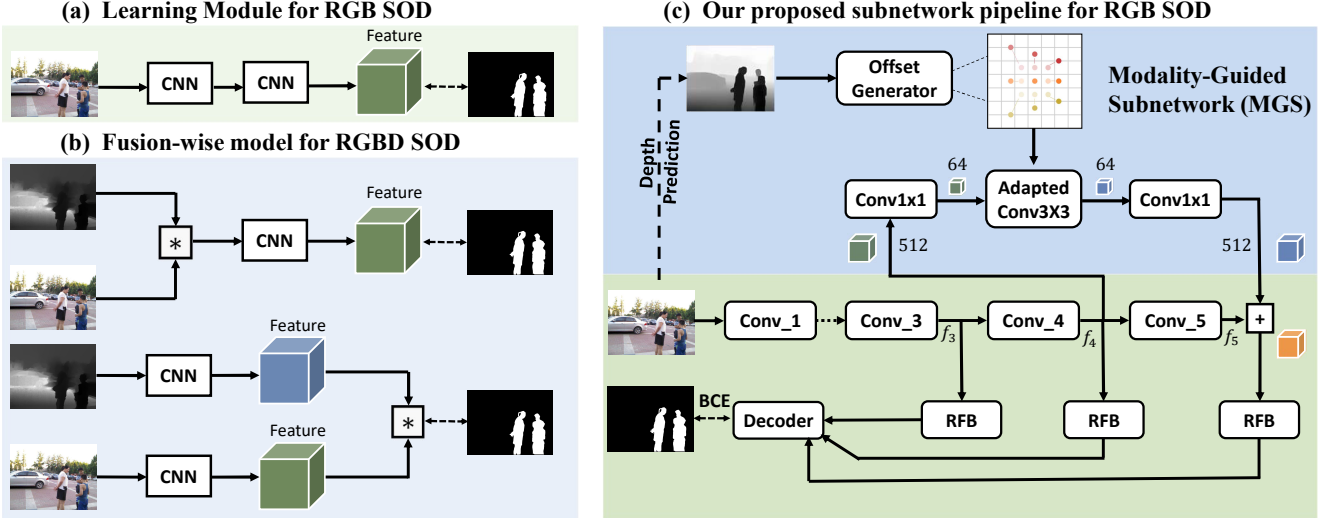


Figure 2: (a) Conventional RGB models [19, 33, 32] insert additional modules to learn geometry-invariant features. (b) RGBD models [9, 36, 23] adopt fusion-wise design to learn both photometric and geometric information. (c) Our proposed MGSnet which takes only RGB image for both training and testing. We use depth prior to guide sampling position on RGB feature map through a subnetwork design to compensate the master streaming.

the fourth channel to RGB image, or multi-scale and late fusion [23] where two-stream networks are adopted. However, early fusion contains more low-level features than semantic ones. Multi-scale or late fusion inevitably requires more learning parameters. As shown in Figure 1, the size of RGBD models is often larger than that of RGB networks.

We explore differently the relationship between depth map and RGB image. Taking human beings as an example, to distinguish salient objects from the 3D world, the input is the visual appearance through human eyes. With the color information and thanks to the depth estimation capability, humans further discover geometric information. This prior guides the understanding of RGB images. It should be the same case for intelligent machines.

To this end, we propose a novel Modality-Guided Subnetwork (MGSnet) which adaptively transforms convolutions by fusing information from one modality to another (e.g., depth to RGB or RGB to depth). Our network matches perfectly both RGB and RGB-D data and dynamically estimates depth if not available by simply applying an off-the-shelf depth prediction model. We design a subnetwork mechanism alongside the master streaming pipeline. The subnetwork can be treated like a light residual-addition branch as the ResNet [10]. It takes one modality map as the master input, e.g. RGB, and enhances its robustness by deforming the convolution kernel with the supervision of the complementary modal prior, e.g. depth, and vice versa.

In summary, the main contributions of this paper are listed as follows :

- By exploiting the nature of CNN sampling position, we propose a novel cross-modal fusion design (MGS) for salient object detection, where we use a subsidiary

modality, i.e., RGB/depth, to guide the main modality streaming, i.e., depth/RGB.

- For RGB-only input, we suggest using an off-the-shelf depth prediction model to mimic the multi-modality input. Our MGSnet enables dramatical performance gain on benchmark datasets and achieves state-of-the-art performance among RGB SOD models.
- The proposed MGS can also be embedded in RGBD two-stream network with the advantage of cross-modality cues while being lightweight.

2. Related Work

RGB SOD: In the past decade, the development of GPU and CNN contributes to the advances of RGB SOD. One core problem is understanding the geometric information from the image. Fully Convolutional Network (FCN) [21] is a pioneering work in leveraging spatial information in CNN. Most recent researches dominating RGB SOD are FCN-based, such as [33] which designs a single stream encoder-decoder system, [15] which adopts a multi-scale network on input, and most currently [5, 19, 32, 35] which fuse multi-level feature map. Some branch designs also have achieved impressive results such as C2S-Net [17] which bridges contour knowledge for SOD. By inserting additional transformation parameters in networks, it contributes to the model performance. Nevertheless, the inference time and computational cost become more significant.

RGBD SOD: The complementary depth map may provide extra clues on the geometry. How to efficiently joint RGB and depth modality is the key challenge for RGBD SOD. One possible solution is to treat the depth map as

an additional channel and adapt a single-stream system as shown in DANet [36]. It further designs a verification process with a depth-enhanced dual attention module. An alternative is to realize multi-stream networks followed by a feature fusion mechanism. PDNet [37] designs a depth-enhanced stream to extract geometric features and further fuses with the RGB features. D3net [7] adopts separate networks to respectively extract features from RGB, depth map, and RGBD four-channel input. A late fusion is further realized. HDFnet [23] adopts two streaming networks for both RGB image and depth map. These features are further fused to generate region-aware dynamic filters. JL-DCF [9] proposes joint learning from cross-modal information through a Siamese network. Generally, RGBD networks achieve superior performance compared to RGB as shown in Figure 1. However, these methods rely on the quality and accessibility of the depth map. A high-quality depth map requires expensive depth sensors and is still sparse compared to an RGB image as suggested in [9, 7]. To this end, DCF [11] proposes to calibrate the raw depth to improve the quality. Nevertheless, the high computational cost due to the two-streaming network requires more development.

Some recent researches [26, 12, 34] propose to learn from RGBD images and tests on RGB. This design enables an RGB CNN to achieve a comparable result with RGBD SOD during testing. Different from it, we propose to firstly discover the hidden geometric modality behind RGB images by simply using an off-the-shelf depth prediction method. With the estimated depth, we further propose a Modality-Guided Subnetwork mechanism to enhance the master RGB network understanding of the contour problem. Our proposed MGSnet achieves state-of-the-art performance with real-time inference speed compared to other RGB models. It can also be embedded in RGBD two-stream models to enable further progress with raw depth.

3. Modality-Guided Subnetwork

3.1. Overview

In Figure 2 (c), our network only takes RGB as input that then estimates the pseudo-depth. Our MGSnet only takes the pseudo-depth to deform the RGB streaming. In other words, only the RGB modality is fed through Conv_4.

Note that our model is not limited by the nature of the modality. It can be a depth-guided RGB convolution as well as an RGB-guided depth convolution. Figure 3 presents our model embedded on an RGBD two-streaming network and Figure 4 illustrates the idea of modality-guided sampling position. We learn the offset from both semantic RGB and depth features to create a cross supervision mechanism.

For simplicity, we present in the following section a depth-guided subnetwork for RGB features. It contains three parts: a master RGB streaming network, an off-the-

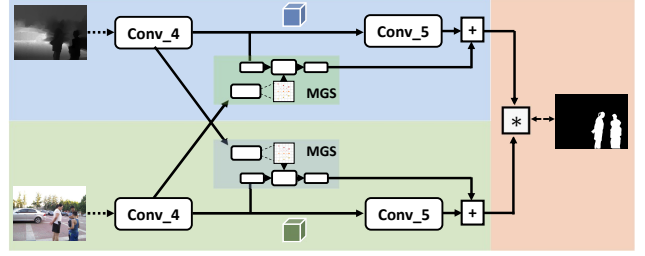


Figure 3: Illustration of embedded MGS on a RGBD two-streaming network.

shelf prediction model to estimate a pseudo-depth map if not available, and a depth-guided subnetwork design. For simplicity, VGG-16 [28] architecture is adopted as our basic convolutional network to extract RGB features for its wide application in SOD. We use RFB [20] on the steamer layers (f_3, f_4, f_5) which contains high level features for SOD as suggested in [26, 23, 9]. We further embed our subnetwork to enhance the edge understanding of the encoder output. We take the same decoder as proposed in [26] and a simple binary cross-entropy (BCE) as the loss.

3.2. Depth-guided Subnetwork

To proceed with the geometric prior, the depth map D and the RGB feature map (output of $Conv_4$) are fed together to our model. We use $f_4 \in \mathbb{R}^{b \times 512 \times h \times w}$ to denote the input RGB feature. The depth prior and RGB feature maps are articulated through an adaptive convolution to compute depth-aware RGB feature maps as output. The last is added to the master RGB stream to form the final feature map.

The subnetwork contains three convolutions of different filter sizes: 1×1 , 3×3 , and 1×1 . It shares the same architecture of plain baseline of ResNet [10] that the 1×1 layers are used for reducing ($512 \rightarrow 64$) and then increasing dimensions ($64 \rightarrow 512$), allowing the 3×3 layer with smaller input/output dimensions. We denote \mathcal{D} and \mathcal{U} for the first and the last 1×1 convolution, which stands for down-sample and up-sample, respectively. This design can significantly reduce the learning parameters, which contributes to the lightweight design of our subnetwork. Different from ResNet that uses the three layers as a bottleneck, we use them as the residual-addition branch which serves as complementary information to the plain network.

A standard convolution is formulated as:

$$\mathbf{y}(p) = \sum_{p_n \in \mathbf{R}(p)} \mathbf{w}(p_n) \cdot \mathbf{x}(p + p_n). \quad (1)$$

with $\mathbf{R}(p)$ the sampling grid for point p . Physically it represents a local neighborhood on input feature map, which conventionally has regular shape with certain dilation Δd , such that :

$$\mathbf{R}(p) = a\vec{u} + b\vec{v} \quad (2)$$

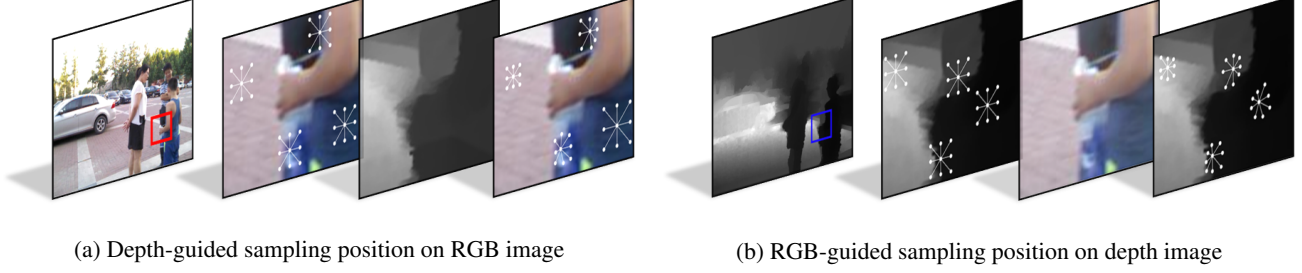


Figure 4: Visual comparison of MGSnet. A pair of RGB and depth images from an RGBD dataset are illustrated on the left. While extracting features through two streaming networks, the cross-modal information beyond the fixed sampling position is not utilized (second left). Our proposed modality-guided sampling position breaks the limit of fixed-local configurations. The new sampling position incorporates supporting modality into the basic function of CNN on the main modality: the fixed sampling position is replaced by relevant neighbors defined by the supporting modality without limitation (right).

where (\vec{u}, \vec{v}) is the pixel coordinate system of input feature map and $(a, b) \in (\Delta d \cdot \{-1, 0, 1\})^2$ for a 3×3 convolution or $(a, b) \in (\{0\})^2$ for 1×1 convolution.

We further replace the conventional 3×3 convolution by deformable convolution (DeformConv) [4], where the kernels are generated with different sampling distributions which is adapted to depth modality. Mathematically, we have:

$$\mathbf{y}(p) = \sum_{p_n \in \mathbf{R}(p)} \mathbf{w}(p_n) \cdot \mathbf{x}(p + p_n + \Delta p_n) \quad (3)$$

The convolution may be operated on the irregular positions $p_n + \Delta p_n$ as the offset Δp_n may be fractional. To address the issue, we use the bilinear interpolation which is the same as that proposed in [4]. The adapted convolution is denoted as \mathcal{A} .

Thanks to the depth input of the subnetwork, the scale and geometric transformation of objects on the RGB feature map can be directly analyzed with the adapted offsets. This process is expressed as:

$$\Delta p_n = \eta(D) \quad (4)$$

We present two types of offset generators according to different plain networks. More details are discussed in the following section. The newly defined sampling position becomes depth-aware and helps to better articulate the RGB feature and geometric information. Finally, the output of MGS is added to the master stream, which serves as complementary depth-aware guidance on RGB features.

The entire process to compute the modality-guided feature f_M can be formulated as follows:

$$\begin{aligned} f_M &= MGS(f_4, D) \\ &= \mathcal{U}(\mathcal{A}(\mathcal{D}(f_4), \eta(D))) \end{aligned} \quad (5)$$

The output of RGB encoder can be formulated as :

$$out = f_5 + \lambda f_M \quad (6)$$

where λ is the weight parameter.

3.3. Offset generator

We use another modal prior to deform the main stream convolution. When the offset exceeds the input size, the output will be computed as if the zeros padding is applied. For RGB input, the pseudo-depth is used to deform the RGB sampling position. The offset is generated through Z-ACN [31]. It firstly back-projects the 2D conventional grid to form a 3D point cloud according to the depth. Based on the point cloud, it extracts a depth-aware 3D plan and further creates a depth-aware 3D regular grid. Then it projects the 3D regular grid to the image plan to form the deformable sampling position. More details can be found in Z-ACN [31] paper. Different to DeformConv [4] that learns offset from the RGB feature map to deform RGB sampling position, Z-ACN computes offset according to low-level geometric constraint (one-channel depth) and does not require gradient descent, thus perfectly matches our light-weight subnetwork design. The computed offset allows the RGB convolution to be scale and rotation independent. We verify through experiments the superior performance of our model in the ablation study.

For RGBD input, current Sconv [2] suggests learning the RGB offset from a semantic depth feature map. We share the same motivation as Sconv. However, Sconv firstly projects the depth into a high-dimensional feature space and secondly learns a depth-aware offset and mask. Unlike Sconv, we learn the offset from the encoder or high-level features to avoid the additional projection. In other words, in our case, the offset generator η is realized through a simple 3×3 convolution to minimize the computational cost. Furthermore, we adapt to different modalities as input, i.e., it learns offset from both RGB and depth, while Sconv only learns from depth.

3.4. Understand adaptive sampling position

Our model aims to compensate for the single modality streaming. As shown in Figure 4, while extracting features from RGB images, the conventional sampling position is

limited by the lack of capability to include geometry due to the fixed shape. We propose to use the depth prior to accurately locate the sampling position. For RGB input without depth prior, we suggest mimicking the depth map by using a monocular depth estimation model. Some pseudo-depth images may be inaccurate due to the domain gap between SOD and monocular depth estimation. In such a case, the offset will converge to 0 so that the deformation becomes minimal and local. The contribution of the depth-aware RGB feature is further regularized by the weight parameter λ of Eq. 6. In Fig. 5, we show that our method is robust to non-optical depth through several examples.

While extracting features from raw depth, conventional sampling positions may produce sub-optimal results due to some inaccurate measurements. The raw depth maps for SOD are obtained by camera measurements such as Kinect and Light Field cameras, or estimated by classic computer vision algorithms as [29, 18]. Thus, the raw depth images may contain noise and ambiguity. We can visualize several low-quality samples on the third row of Figure 5. To this end, we propose to use the RGB image to deform the depth sampling position. In such a case, the RGB-guided sampling position can make up for the measurement error on geometry.

4. Experiments

4.1. Benchmark Dataset

To verify the effectiveness of our method, we conduct experiments on seven following benchmark RGBD datasets. DES [3]: includes 135 images about indoor scenes captured by Kinect camera. LFSD [16]: contains 100 images collected on the light field with an embedded depth map and human-labeled ground truths. NLPR [24]: contains 1000 natural images captured by Kinect under different illumination conditions. NJUD [13]: contains 1,985 stereo image pairs from different sources such as the Internet, 3D movies, and photographs taken by a Fuji W3 stereo camera and with estimated depth by using optical flow method [29]. SSD [38]: contains 80 images picked up from stereo movies with estimated depth from flow map [29]. STEREO [22]: includes 1000 stereoscopic images downloaded from the Internet where the depth map is estimated by using SIFT flow method [18]. DUT-RGBD [25]: contains 1200 images captured by Lytro camera in real-life scenes.

4.2. Experimental Settings

Our model is implemented basing on the Pytorch toolbox and trained with a GTX 3090Ti GPU. We adopt several generally-recognized metrics for quantitative evaluation: F-measure is a region-based similarity metric that takes into account both Precision (Pre) and Recall (Rec). Mathematically, we have: $F_\beta = \frac{(1+\beta^2) \cdot \text{Pre} \cdot \text{Rec}}{\beta^2 \cdot \text{Pre} + \text{Rec}}$. The value of β^2 is set

to be 0.3 as suggested in [1] to emphasize the precision. In this paper, we report the **maximum F-measure** (F_β) score across the binary maps of different thresholds, the **mean F-measure** (F_β^{mean}) score across an adaptive threshold and the **weighted F-measure** (F_β^w) which focuses more on the weighted precision and weighted recall. **Mean Absolute Error** (MAE) studies the approximation degree between the saliency map and ground-truth map on the pixel level. **S-measure** (S_m) evaluates the similarities between object-aware (S_o) and region-aware (S_r) structure between the saliency map and ground-truth map. Mathematically, we have: $S_m = \alpha \cdot S_o + (1 - \alpha) \cdot S_r$, where α is set to be 0.5. **E-measure** (E_m) studies both image level statistics and local pixel matching information. Mathematically, we have: $E_m = \frac{1}{W \times H} \sum_{i=1}^W \sum_{j=1}^H \phi_{FM}(i, j)$, where $\phi_{FM}(i, j)$ stands for the enhanced-alignment matrix as presented in [6].

4.3. Performance Comparison with RGB Input

We firstly compare with RGB models, including R3Net [5], PoolNet [19], CPD [32], AFnet [8]. All saliency maps are directly provided by authors or computed by authorized codes. For fair comparisons, we adopt the same training set as suggested in [26], which contains 1485 samples from NJUD, 700 samples from NLPR, and 800 samples from the DUT-RGBD dataset. The remaining images of all listed datasets are used for testing. The quantitative comparison is presented in Table 1. Our model is trained with 50 epochs with 256×256 input image size.

For the RGB model, we can conclude from Table 1 that the improvement on the saliency map is attributed to different learning modules, which results in high computational cost (size). Different from traditional RGB models which do not exploit the depth information, we propose to take full advantage of the pseudo-geometry estimated with an existing monocular depth estimation method.

We re-train two RGB-D SOD network (HDFnet [23], CoNet [12]) with the additional estimated pseudo-depth. We observe a significant performance gap between the recent RGB-D models and the previous RGB models. The main reason is the quality of depth estimation: the domain gap between the depth estimation dataset and the SOD dataset leads to some failure depth maps. This can be noticed in the poor performance of HDFnet that extracts features from both RGB and depth images. CoNet, however, is more robust to the depth quality since the depth map is only used to supervise the feature extraction on RGB images. Our model shares the same motivation as CoNet to use depth prior to guide SOD but in a completely different manner. In our model, we directly learn a geometric-aware offset from the depth map to the sampling position on the RGB image. Our model achieves consistent superior performance compared with other models.

Dataset		DES			NLPR			NJUD			STEREO			DUT-RGBD		
Metric	Size ↓	MAE ↓	$F_{\beta}^{mean} \uparrow$	$F_{\beta}^w \uparrow$	MAE ↓	$F_{\beta}^{mean} \uparrow$	$F_{\beta}^w \uparrow$	MAE ↓	$F_{\beta}^{mean} \uparrow$	$F_{\beta}^w \uparrow$	MAE ↓	$F_{\beta}^{mean} \uparrow$	$F_{\beta}^w \uparrow$	MAE ↓	$F_{\beta}^{mean} \uparrow$	$F_{\beta}^w \uparrow$
RGB input																
<i>R³Net</i> ₁₈	225	.066	.728	.693	.101	.649	.611	.092	.775	.736	.084	.800	.752	.113	.781	.709
<i>PoolNet</i> ₁₉	279	.031	.852	.814	.046	.791	.771	.057	.850	.816	.045	.877	.849	.049	.871	.836
<i>CPD</i> ₁₉	112	.028	.860	.841	.037	.840	.829	.059	.853	.821	.046	.880	.851	.055	.872	.835
<i>AFNet</i> ₁₉	144	.034	.840	.816	.043	.807	.796	.056	.857	.832	.046	.876	.850	.064	.851	.817
<i>EGNet</i> ₁₉	412	.035	.831	.797	.047	.800	.774	.060	.846	.808	.049	.876	.835	.059	.866	.805
+ Pseudo Depth (86 Mb extra model size)																
<i>HDFnet</i> ₂₀	177 (+86)	.070	.721	.664	.062	.758	.741	.124	.716	.656	.106	.743	.684	-	-	-
<i>CoNet</i> ₂₀	171 (+86)	.037	.820	.808	.049	.744	.835	.068	.827	.795	.050	.848	.825	.045	.865	.847
<i>Ours</i>	62 (+86)	.028	.871	.837	.025	.888	.874	.047	.882	.856	.041	.881	.857	.037	.906	.889

Table 1: Quantitative comparisons of with RGB input. The off-the-shelf depth estimation is realized with MiDaS [27] which presents 86Mb model size. \uparrow & \downarrow denote larger and smaller is better, respectively. (red: best, blue: second best).

4.4. Performance Comparison with RGB-D Input

We also compare with state-of-the-art RGBD models with raw depth input in the Table 2, including CoNet [12], A2dele [26], DANet [36], cmMS [14], HDFnet [23], and DSA2F [30]. For fair comparisons, all saliency maps and the FPS are directly provided by authors or computed by authorized codes. Note that the FPS depends on the GPU for inference. Thus, only the FPS of HDFnet is tested on the same GPU as ours.

While depth is only used as supervision during training and only RGB image is required during testing, our model surpasses existing efficient A2dele significantly on performance with only an + around 5Mb model size. Compared to CoNet, the model size is minimized by 63% and achieves a comparable result. As presented in Figure 4, our proposed module can take advantage of cross-modality cues while being lightweight. Thus, we further incorporate with the HDFnet [23] to show the performance gain by integrating our approach. It achieves the state-of-the-art (SOTA) performance on VGG16 based models (*HDF* + *Ours*). To better demonstrate the superiority of the proposed method, we also use a larger backbone (VGG19) to compare with the plain version HDFnet and the SOTA method DSA2F. Note that DSA2F uses neural architecture search to automate the model architecture while ours is hand-designed. Our model enables significant gains on the plain version with minimal cost (+ around 1 Mb on model size) and achieves comparable results with the DSA2F.

4.5. Qualitative Evaluation

We present the qualitative result with some challenging cases in Figure 5: low density (1st columns), similar visual appearance between foreground and background (2nd – 5th columns), small objects (6th columns), far objects (7th – 9th columns), human in scene (10th columns), and similar and low contrast on depth map (11th – 13th columns). It can be

seen that our MGSnet yields the results closer to the ground truth mask in various challenging scenarios, especially for the last three columns with low-quality depth clues. Different from two-stream networks that tend to treat sub-optimal depth equally as RGB input, MGSnet extracts features from RGB images while the depth map serves only as complementary guidance, thus becoming robust to depth bias. By analyzing the response on HDFnet (sixth row) and HDFnet with embedded MGS (seventh row), we observe that our approach enables the plain network better discrimination of salient objects from the background.

5. Ablation Study

Effect of Modality-Guided Sampling Position: Our modality-guided sampling position aims to incorporate multi-modal information through the basic function of CNN - the sampling position of convolution. This pattern is integrated in Eq. 3 and Eq. 4. To verify the effectiveness of the proposed modality-guided sampling position, a series of experiments with different learning strategies are realized.

(1) - (4) are experiments on RGB model: (1) RGB Baseline. (2) Self-guided deformable sampling position. We learn the offset from the RGB feature map. (3) RGB pseudo-depth early fusion. We form a four-channel input with pseudo depth. (4) Depth-guided deformable position. We compute an offset from pseudo-depth using Z-ACN to guide RGB streaming. (5) - (7) are experiments on RGBD model: (5) Baseline. We use the same architecture as HDFnet. (6) Self-guided deformable sampling position. The offset applied to RGB streaming is learned from the RGB feature. Idem for depth streaming. (7) Cross modality-guided deformable position. We learn an offset from depth to guide RGB streaming, and vice versa.

Table 3 (1) and (3) compare the performance of the baseline RGB three-channel input and mimicked RGBD four-channel input with pseudo-depth, respectively. The mim-

		Extract RGB feature			Extract RGBD feature						
		<i>CoNet</i> ₂₀	<i>A2dele</i> ₂₀	<i>Ours</i>	<i>DANet</i> ₂₀	<i>cmMS</i> ₂₀	<i>HDFnet</i> ₂₀	<i>HDF</i> ⁺ <i>Ours</i>	<i>DSA2F</i> ₂₁	<i>HDFnet</i> ₂₀	<i>HDF</i> ⁺ <i>Ours</i>
Backbone		Resnet101	– VGG16 –		– VGG16 –				– VGG19 –		
	<i>Size</i> ↓	167	<u>57</u>	62	<u>102</u>	430	<u>177</u>	178	-	<u>220</u>	221
	<i>FPS</i> ↑	-	120	150	32		<u>62</u>	58	-	-	-
DES	<i>MAE</i> ↓	<u>.027</u>	.029	.028	.023	-	.030	.019	.021	.017	.017
	<i>F_β^{mean}</i> ↑	.862	.870	.871	.887	-	.843	.920	.896	.918	.923
	<i>S_m</i> ↑	<u>.910</u>	.881	.882	.904	-	.899	.935	.920	.937	.937
	<i>E_m</i> ↑	<u>.945</u>	.918	.922	.967	-	.944	.979	.962	.976	.979
NLPR	<i>MAE</i> ↓	.031	.031	.025	.028	.027	.027	.025	<u>.024</u>	.027	.025
	<i>F_β^{mean}</i> ↑	.848	.871	.888	.871	.869	.878	.885	<u>.891</u>	.883	.882
	<i>S_m</i> ↑	.908	.889	.908	.915	.899	.898	.918	.918	.915	.918
	<i>E_m</i> ↑	.934	.937	.952	.949	.945	.948	.954	.950	.951	.951
NUJ	<i>MAE</i> ↓	.047	.052	.047	.045	.044	.039	.037	.039	.038	.035
	<i>F_β^{mean}</i> ↑	.872	.873	.882	.871	.886	.887	.893	.898	.887	.898
	<i>S_m</i> ↑	<u>.895</u>	.867	.879	.899	.900	.907	.911	.903	.911	.912
	<i>E_m</i> ↑	.911	.914	.928	.922	.914	.931	.935	.923	.932	.942
STEREO	<i>MAE</i> ↓	<u>.037</u>	.044	.041	.047	.043	.042	.039	.039	.040	.039
	<i>F_β^{mean}</i> ↑	<u>.885</u>	.875	.881	.858	<u>.879</u>	.864	.864	<u>.893</u>	.875	.878
	<i>S_m</i> ↑	<u>.908</u>	.878	.887	.901	.895	.900	.904	.897	.903	.902
	<i>E_m</i> ↑	.928	.929	.936	.914	.922	.929	.937	.933	.934	.938

Table 2: Quantitative comparisons of with recent RGBD models. \uparrow & \downarrow denote larger and smaller is better, respectively. MGS can also be embedded to the HDFnet [23] to enable further progress. The scores/numbers better than ours are underlined (extracting RGB feature, extracting RGBD feature with VGG16, and extracting RGBD feature with VGG19 models are labeled separately). More details on all RGB-D dataset can be found in the Supplementary Material.

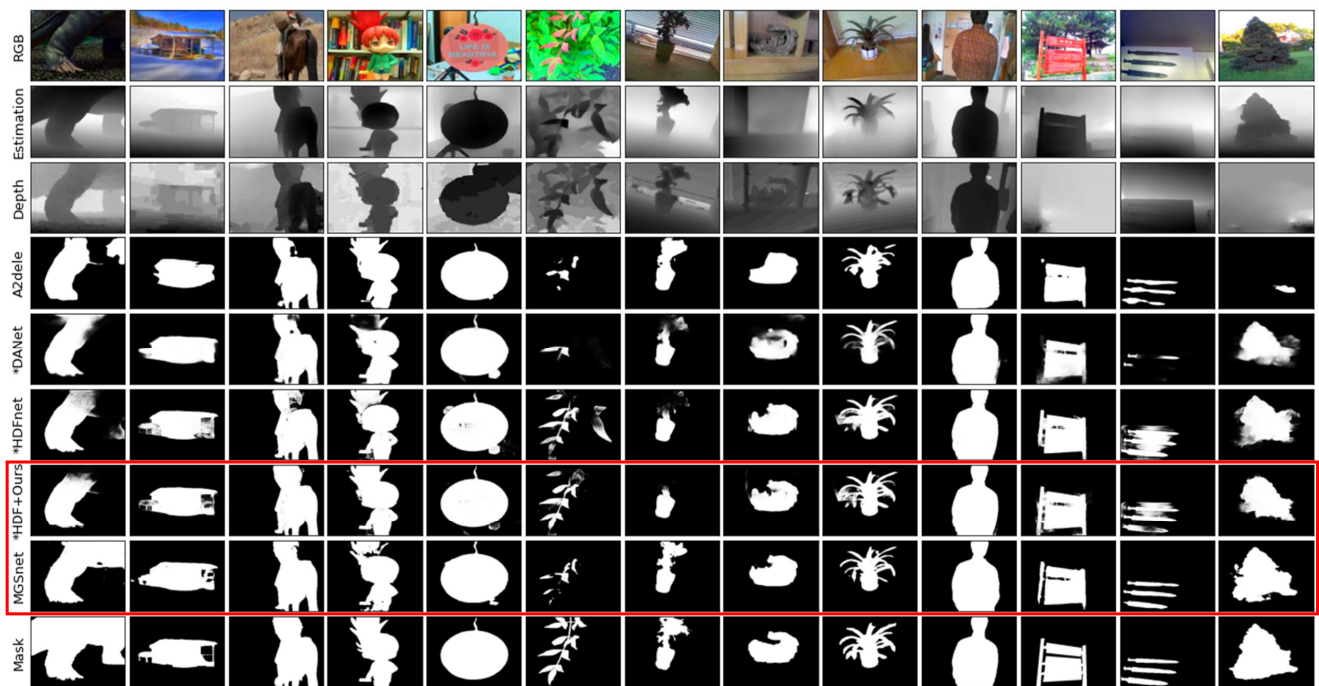


Figure 5: Visual comparison between our proposed MGSnet and the state-of-the-art RGB/RGBD methods. * denotes that the ground truth depth is used during testing. We also embed MGS on the HDFnet [23] to enable further improvement, denoted as **HDF + Ours*. More examples can be found in the Supplementary Material.

Dataset	DES		LFSD		NLPR		NJUD		SSD		STEREO	
# Metric	MAE ↓	F_{\max} ↑	MAE ↓	F_{\max} ↑	MAE ↓	F_{\max} ↑	MAE ↓	F_{\max} ↑	MAE ↓	F_{\max} ↑	MAE ↓	F_{\max} ↑
1 RGB Baseline	.036	.880	.104	.835	.037	.883	.059	.877	.082	.785	.054	.882
2 RGB + Self Deform	.042	.860	.110	.798	.031	.885	.052	.879	.062	.814	.046	.882
3 RGB pseudo D Fusion	.032	.888	.093	.819	.029	.893	.061	.863	.077	.776	.049	.878
4 RGB + Depth-Deform	.028	.899	.078	.849	.025	.905	.047	.897	.063	.801	.041	.898
5 RGB-D Baseline (B)	.020	.933	.089	.856	.028	.921	.039	.922	.054	.867	.042	.911
6 B + Self Deform	.019	.936	.089	.853	.026	.916	.038	.925	.053	.878	.044	.895
7 B + Cross-Modal Deform	.019	.936	.079	.871	.025	.921	.037	.925	.049	.867	.039	.917

Table 3: Ablation study of modality-guided sampling position




Figure 6: Visual analysis of embedded depth with MGSNet.

icked multi-modality early fusion achieves better performance, indicating that the pseudo-depth provides additional semantic. However, by comparing (3) and (4), we observe that the proposed depth-guided deformable sampling position can better use the complementary information to supervise RGB streaming, compared with early fusion. By comparing (2) and (4), we show that the depth-guided deformable position is more accurate on saliency compared to that of the self/RGB-guided. This verifies the assumption that depth cues can help the RGB model to better distinguish the foreground and background. Note that in (4) we only extract features from RGB images. The additional awareness of the geometry is only treated as a 2D offset to better locate the sampling position. This new integration design contributes to the model performance with minimal cost. For better understanding, the qualitative result presented in Figure 6 shows that our approach provides more accurate saliency maps with better contrast. On the RGBD model (5-7), we also observe the superior performance with the cross-modality deformable sampling position achieves as it directly compensates for the single modal streaming.

Performance with different depth qualities: We also conduct an experiment to show the impact of depth quality. We choose the HDFnet [23] as the baseline and further embed it with our method. We present the average metric on all testing datasets in Table 4 with pseudo-depth (estimated) and raw depth from the RGBD dataset. Results obtained with pseudo-depth are denoted with *.

It shows that the quality of depth has an important in-

AvgMetric	HDFnet*	+Ours*	HDFnet	+Ours
MAE ↓	.1053	.0758	.0405	.0375
F_{β} ↑	.8410	.8599	.9121	.9166
F_{β}^{mean} ↑	.7326	.7868	.8730	.8831
F_{β}^w ↑	.6789	.7488	.8569	.8672
S_m ↑	.8010	.8390	.9013	.9053
E_m ↑	.8359	.8797	.9312	.9377

Table 4: Performance variation with different depth qualities. (*) denotes results obtained with pseudo-depth.

fluence on performance. Features extracted from raw depth describe better the salient object and were in line with our expectations. However, in both cases, our MGS can significantly enable progress compared to the plain networks. For pseudo-depth, the contribution of our MGS is more significant, which can be explained by the effectiveness of our RGB-guided sampling position for depth streaming. It can efficiently help to alleviate depth errors.

6. Conclusions

In this paper, we propose a modality-guided module (MGSnet) for RGB-D salient object detection models. The depth channel can either be the input or be estimated using a state-of-the-art monodepth network. Our model adaptively transforms convolutions that their size and shape are built by fusing information from one modality to another (e.g., depth to RGB and/or RGB to depth), thus enabling a cross-modality interaction. Extensive experiments against RGB baselines demonstrate the performance gains of the proposed module, and the addition of the proposed module to existing RGB-D models further improved results.

Acknowledgements

We gratefully acknowledge Zhuyun Zhou for her support and proofreading. We also thank Jilai Zheng, Zhongpai Gao, and Yonglin Zhang for the discussion. This research is supported by the French National Research Agency through ANR CLARA (ANR-18-CE33-0004) and financed by the French Conseil Régional de Bourgogne-Franche-Comté.

References

- [1] Radhakrishna Achanta, Sheila Hemami, Francisco Estrada, and Sabine Susstrunk. Frequency-tuned salient region detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 5
- [2] Lin-Zhuo Chen, Zheng Lin, Ziqin Wang, Yong-Liang Yang, and Ming-Ming Cheng. Spatial information guided convolution for real-time RGBD semantic segmentation. *IEEE Transactions on Image Processing (TIP)*, 2021. 4
- [3] Yupeng Cheng, Huazhu Fu, Xingxing Wei, Jiangjian Xiao, and Xiaochun Cao. Depth enhanced saliency detection method. In *Proceedings of international conference on internet multimedia computing and service*, 2014. 5
- [4] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision (ICCV)*, 2017. 4
- [5] Zijun Deng, Xiaowei Hu, Lei Zhu, Xuemiao Xu, Jing Qin, Guoqiang Han, and Pheng-Ann Heng. R3net: Recurrent residual refinement network for saliency detection. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018. 1, 2, 5
- [6] Deng-Ping Fan, Cheng Gong, Yang Cao, Bo Ren, Ming-Ming Cheng, and Ali Borji. Enhanced-alignment measure for binary foreground map evaluation. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 2018. 5
- [7] Deng-Ping Fan, Zheng Lin, Zhao Zhang, Menglong Zhu, and Ming-Ming Cheng. Rethinking RGB-D salient object detection: Models, data sets, and large-scale benchmarks. *IEEE Transactions on Neural Networks and Learning Systems*, 2020. 3
- [8] Mengyang Feng, Huchuan Lu, and Errui Ding. Attentive feedback network for boundary-aware salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 5
- [9] Keren Fu, Deng-Ping Fan, Ge-Peng Ji, and Qijun Zhao. JL-DCF: Joint learning and densely-cooperative fusion framework for RGB-D salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 2, 3
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, 2016. 2, 3
- [11] Wei Ji, Jingjing Li, Shuang Yu, Miao Zhang, Yongri Piao, Shunyu Yao, Qi Bi, Kai Ma, Yefeng Zheng, Huchuan Lu, et al. Calibrated RGB-D salient object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 3
- [12] Wei Ji, Jingjing Li, Miao Zhang, Yongri Piao, and Huchuan Lu. Accurate RGB-D salient object detection via collaborative learning. In *European Conference on Computer Vision (ECCV)*, 2020. 3, 5, 6
- [13] Ran Ju, Ling Ge, Wenjing Geng, Tongwei Ren, and Gangshan Wu. Depth saliency based on anisotropic center-surround difference. In *2014 IEEE International Conference on Image Processing (ICIP)*, 2014. 5
- [14] Chongyi Li, Runmin Cong, Yongri Piao, Qianqian Xu, and Chen Change Loy. RGB-D salient object detection with cross-modality modulation and selection. In *European Conference on Computer Vision (ECCV)*, 2020. 6
- [15] Guanbin Li and Yizhou Yu. Deep contrast learning for salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 2
- [16] Nianyi Li, Jinwei Ye, Yu Ji, Haibin Ling, and Jingyi Yu. Saliency detection on light field. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 5
- [17] Xin Li, Fan Yang, Hong Cheng, Wei Liu, and Dinggang Shen. Contour knowledge transfer for salient object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 2
- [18] Ce Liu, Jenny Yuen, and Antonio Torralba. Sift flow: Dense correspondence across scenes and its applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2010. 5
- [19] Jiang-Jiang Liu, Qibin Hou, Ming-Ming Cheng, Jiashi Feng, and Jianmin Jiang. A simple pooling-based design for real-time salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5
- [20] Songtao Liu, Di Huang, et al. Receptive field block net for accurate and fast object detection. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. 3
- [21] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [22] Yuzhen Niu, Yujie Geng, Xueqing Li, and Feng Liu. Leveraging stereopsis for saliency analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 5
- [23] Youwei Pang, Lihe Zhang, Xiaoqi Zhao, and Huchuan Lu. Hierarchical dynamic filtering network for RGB-D salient object detection. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3, 5, 6, 7, 8
- [24] Houwen Peng, Bing Li, Weihua Xiong, Weiming Hu, and Rongrong Ji. RGBD salient object detection: a benchmark and algorithms. In *European Conference on Computer Vision (ECCV)*, 2014. 1, 5
- [25] Yongri Piao, Wei Ji, Jingjing Li, Miao Zhang, and Huchuan Lu. Depth-induced multi-scale recurrent attention network for saliency detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 5
- [26] Yongri Piao, Zhengkun Rong, Miao Zhang, Weisong Ren, and Huchuan Lu. A2dele: Adaptive and attentive depth distiller for efficient RGB-D salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 3, 5, 6
- [27] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset

- transfer. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2019. 6
- [28] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *3rd International Conference on Learning Representations, (ICLR)*, 2015. 3
 - [29] Deqing Sun, Stefan Roth, and Michael J Black. Secrets of optical flow estimation and their principles. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 5
 - [30] Peng Sun, Wenhui Zhang, Huanyu Wang, Songyuan Li, and Xi Li. Deep RGB-D saliency detection with depth-sensitive attention and automatic multi-modal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6
 - [31] Zongwei Wu, Guillaume Allibert, Christophe Stolz, and Cédric Demonceaux. Depth-Adapted CNN for RGB-D cameras. In *Asian conference on computer vision (ACCV)*, 2020. 4
 - [32] Zhe Wu, Li Su, and Qingming Huang. Cascaded partial decoder for fast and accurate salient object detection. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 5
 - [33] Pingping Zhang, Dong Wang, Huchuan Lu, Hongyu Wang, and Baocai Yin. Learning uncertain convolutional features for accurate saliency detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017. 1, 2
 - [34] Jiawei Zhao, Yifan Zhao, Jia Li, and Xiaowu Chen. Is depth really necessary for salient object detection? In *Proceedings of the 28th ACM International Conference on Multimedia (ACM MM)*, 2020. 3
 - [35] Jia-Xing Zhao, Jiang-Jiang Liu, Deng-Ping Fan, Yang Cao, Jufeng Yang, and Ming-Ming Cheng. EGNet: Edge Guidance Network for salient object detection. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2
 - [36] Xiaoqi Zhao, Lihe Zhang, Youwei Pang, Huchuan Lu, and Lei Zhang. A single stream network for robust and real-time RGB-D salient object detection. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 2, 3, 6
 - [37] Chunbiao Zhu, Xing Cai, Kan Huang, Thomas H Li, and Ge Li. PDNet: Prior-model guided depth-enhanced network for salient object detection. In *IEEE International Conference on Multimedia and Expo (ICME)*, 2019. 3
 - [38] Chunbiao Zhu and Ge Li. A three-pathway psychobiological framework of salient object detection using stereoscopic technology. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, 2017. 5